

Edited by
Paul Égré and
Nathan Klinedinst



Vagueness and Language Use

Palgrave Studies in Pragmatics, Language and Cognition



Palgrave Studies in Pragmatics, Language and Cognition

Series Editors: **Richard Breheny** and **Uli Sauerland**

Series Advisors: **Kent Bach**, **Anne Bezuidenhout**, **Noël Burton-Roberts**,
Robyn Carston, **Sam Glucksberg**, **Francesca Happé**, **François Recanati**,
Deirdre Wilson

Palgrave Studies in Pragmatics, Language and Cognition is a series of high quality research monographs and edited collections of essays focusing on the human pragmatic capacity and its interaction with natural language semantics and other faculties of mind. A central interest is the interface of pragmatics with the linguistic system(s), with the 'theory of mind' capacity and with other mental reasoning and general problem-solving capacities. Work of a social or cultural anthropological kind is included if firmly embedded in a cognitive framework. Given the interdisciplinarity of the focal issues, relevant research will come from linguistics, philosophy of language, theoretical and experimental pragmatics, psychology and child development. The series aims to reflect all kinds of research in the relevant fields – conceptual, analytical and experimental.

Titles include:

Anton Benz, Gerhard Jäger and Robert van Rooij (*editors*)

GAME THEORY AND PRAGMATICS

Reinhard Blutner and Henk Zeevat (*editors*)

OPTIMALITY THEORY AND PRAGMATICS

Paul Égré and Nathan Klinedinst (*editors*)

VAGUENESS AND LANGUAGE USE

María J. Frápolli (*editor*)

SAYING, MEANING AND REFERRING

Essays on François Recanati's Philosophy of Language

Corinne Iten

LINGUISTIC MEANING, TRUTH CONDITIONS AND RELEVANCE

The Case of Concessives

Mark Jary

ASSERTION

Ira Noveck and Dan Sperber (*editors*)

EXPERIMENTAL PRAGMATICS

Klaus Petrus

MEANING AND ANALYSIS

New Essays on Grice

George Powell (*editor*)

LANGUAGE, THOUGHT AND REFERENCE

Uli Sauerland and Penka Stateva (*editors*)
PRESUPPOSITION AND IMPLICATURE IN COMPOSITIONAL SEMANTICS

Uli Sauerland and Kazuko Yatsushiro (*editors*)
SEMANTICS AND PRAGMATICS
From Experiment to Theory

Hans-Christian Schmitz
ACCENTUATION AND INTERPRETATION

Belén Soria and Esther Romero (*editors*)
EXPLICIT COMMUNICATION
Robyn Carston's Pragmatics

Christoph Unger
GENRE, RELEVANCE AND GLOBAL COHERENCE
The Pragmatics of Discourse Type

Palgrave Studies in Pragmatics, Language and Cognition Series
Series Standing Order ISBN 978-0-333-99010-0 Hardback
978-0-333-98584-7 Paperback
(*outside North America only*)

You can receive future titles in this series as they are published by placing a standing order. Please contact your bookseller or, in case of difficulty, write to us at the address below with your name and address, the title of the series and the ISBN quoted above.

Customer Services Department, Macmillan Distribution Ltd, Houndmills,
Basingstoke,

Vagueness and Language Use

Edited by

Paul Égré

Institut Jean-Nicod, Paris

and

Nathan Klinedinst

University College London, UK

palgrave
macmillan



Selection and editorial matter © Paul Égré and Nathan Klinedinst 2011
Chapters © their individual authors 2011

All rights reserved. No reproduction, copy or transmission of this publication may be made without written permission.

No portion of this publication may be reproduced, copied or transmitted save with written permission or in accordance with the provisions of the Copyright, Designs and Patents Act 1988, or under the terms of any licence permitting limited copying issued by the Copyright Licensing Agency, Saffron House, 6–10 Kirby Street, London EC1N 8TS.

Any person who does any unauthorized act in relation to this publication may be liable to criminal prosecution and civil claims for damages.

The authors have asserted their rights to be identified as the authors of this work in accordance with the Copyright, Designs and Patents Act 1988.

First published 2011 by
PALGRAVE MACMILLAN

Palgrave Macmillan in the UK is an imprint of Macmillan Publishers Limited, registered in England, company number 785998, of Houndmills, Basingstoke, Hampshire RG21 6XS.

Palgrave Macmillan in the US is a division of St Martin's Press LLC, 175 Fifth Avenue, New York, NY 10010.

Palgrave Macmillan is the global academic imprint of the above companies and has companies and representatives throughout the world.

Palgrave® and Macmillan® are registered trademarks in the United States, the United Kingdom, Europe and other countries.

ISBN 978–0–230–23861–9 hardback

This book is printed on paper suitable for recycling and made from fully managed and sustained forest sources. Logging, pulping and manufacturing processes are expected to conform to the environmental regulations of the country of origin.

A catalogue record for this book is available from the British Library.

Library of Congress Cataloging-in-Publication Data

Vagueness and language use / edited by Paul Égré, Nathan Klinedinst.
p. cm.—(Palgrave studies in pragmatics, language and cognition)
ISBN 978–0–230–23861–9

1. Semantics. 2. Ambiguity. 3. Vagueness (Philosophy).
4. Language and languages—Philosophy. 5. Semantics. 6. Ambiguity.
I. Égré, Paul. II. Klinedinst, Nathan, 1977–
P325.V25 2010
401'.43—dc22

2010027503

10 9 8 7 6 5 4 3 2 1
20 19 18 17 16 15 14 13 12 11

Printed and bound in Great Britain by
CPI Antony Rowe, Chippenham and Eastbourne

Contents

<i>List of Figures and Tables</i>	vii
<i>Acknowledgements</i>	viii
<i>Notes on the Contributors</i>	x
1 Introduction: Vagueness and Language Use	1
<i>Paul Égré and Nathan Klinedinst</i>	
Part I Measurement and Comparison	
2 Vagueness and Scales	25
<i>Scott Fults</i>	
3 Implicit versus Explicit Comparatives	51
<i>Robert van Rooij</i>	
4 Vagueness and Comparison	73
<i>Christopher Kennedy</i>	
5 The Inhabitants of Vagueness Models	98
<i>Galit W. Sassoon</i>	
Part II Approximators and Intensifiers	
6 Two Types of Vagueness	121
<i>Uli Sauerland and Penka Stateva</i>	
7 Degree Modifiers and Monotonicity	146
<i>Rick Nouwen</i>	
8 Clarity as Objectivized Belief	165
<i>Lavi Wolf and Ariel Cohen</i>	
9 Commentary on Wolf and Cohen: Reasoning about Public Evidence	191
<i>Chris Barker</i>	

Part III The Sorites Paradox

10	Supervaluationism and Fara's Argument Concerning Higher-Order Vagueness	207
	<i>Pablo Cobreros</i>	
11	Truth in a Region	222
	<i>Delia Graff Fara</i>	
12	Vagueness and Practical Interest	249
	<i>Paula Sweeney and Elia Zardini</i>	
13	Vagueness and Domain Restriction	283
	<i>Peter Pagin</i>	
	<i>Author Index</i>	308
	<i>Subject Index</i>	311

List of Figures and Tables

Figures

2.1	Hypothesized results of experiment	31
2.2	Shape of the buffer zone	32
2.3	Gaussian distributions representing mental magnitudes: discrimination is determined by overlap	36
2.4	Analog magnitudes with scalar variability	37
4.1	Uranus (51,118 km) vs Venus (12,100 km)	74
4.2	Uranus (51,118 km) vs Neptune (49,500 km)	77
4.3	Saturn, Uranus, Neptune and Earth	92
6.1	Alignment of length terms with scale intervals	126
10.1	Fara's argument	212
10.2	Failure of \mathcal{D} -intro	215
10.3	Regional consistency of gap principles	218
13.1	Union of gaps	295

Tables

6.1	Relative frequencies of approximator–adjective strings	134
-----	--	-----

Acknowledgements

The origin of this book was the conference “Vagueness and Language Use” held in Paris at the École Normale Supérieure on April 7–9, 2008. The aim of the conference was to bring together philosophers and linguists with recent contributions on vagueness and mutual interest in interdisciplinary dialogue. The conference, which included both invited and contributed papers, proved to be a success, and soon thereafter Uli Sauerland encouraged us to submit a book proposal to Palgrave Macmillan. We are very grateful to him as well as to Richard Breheny for their support as series editors, and to Palgrave Macmillan for welcoming our book in their series in *Pragmatics, Language and Cognition*.

Special thanks go to all the contributors in this volume, not only for their essays, but also for their collaborative spirit. All chapters in this volume received extended reviews on the basis of either blind cross-refereeing between authors, or reviews by the editors, or blind reviews from external reviewers. We are pleased to stress this point, as essays featuring in volumes tend to receive less credit than papers that appear in journals, and so we thank all our authors for their efforts in reviewing as well as improving the first version of these chapters.

Not all of the good papers that were presented in Paris at the Vagueness and Language Use conference could be solicited to appear in the present volume. However, we would like to take this opportunity to thank all participants in that event, as well as the members of our scientific committee on that occasion, namely Anouk Barberousse, Chris Barker, Denis Bonnay, Richard Dietz, Patrick Greenough, Rosanna Keefe, Pascal Ludwig, Chris Kennedy, Peter Pagin, Agustín Rayo, Robert van Rooij, Benjamin Spector and Timothy Williamson. Special thanks, moreover, go to David Ripley for his expertise and assistance in the volume as an additional reviewer.

We also thank the Agence Nationale de la Recherche and the ANR project “Cognitive Origins of Vagueness” (ANR-07-JCJC-0070) for funding and support in this project, as well as the Institut Jean-Nicod, the Department of Cognitive Studies at the École Normale Supérieure (DEC-ENS), and the Department of Linguistics at University College London.

Last but not least, thanks go to our partners, families and friends who made room both for needed gluts and pleasant gaps between the academic and non-academic moments spent around the production of this volume.

PAUL ÉGRÉ AND NATHAN KLINEDINST

Notes on the Contributors

Chris Barker is Professor of Linguistics at New York University. His recent published work includes discussions of vagueness, scope, parasitic scope (leading to a compositional treatment of “same” and “different”), continuations, and substructural logics for natural language reasoning about permission.

Pablo Cobreros is Lecturer at the Department of Philosophy, University of Navarra (Spain). His research interests include philosophy of language, philosophical logic and philosophy of science. He is the author of “Supervaluationism and Logical Consequence: a Third Way” (*Studia Logica*, 2008) and “Varzi on Supervaluationism and Logical Consequence” (*Mind*, 2010).

Ariel Cohen is Senior Lecturer of Linguistics, Ben-Gurion University of the Negev, Israel. He has published widely on topics in semantics, pragmatics and computational linguistics. Dr Cohen is currently writing a book on the semantics and pragmatics of implicit quantification.

Paul Égré is a CNRS Research Fellow at Institut Jean-Nicod and a member of the Department of Cognitive Studies at the École Normale Supérieure in Paris. His main research areas are the philosophy of language, philosophical logic, and epistemology. Since 2008, Paul Égré has been leading an ANR project on the cognitive sources of vagueness. Paul Égré is an associate editor of the *Journal of Philosophical Logic*, of *Disputatio*, and one of the founders and executive editors of the *Review of Philosophy and Psychology*.

Delia Fara is Associate Professor, Department of Philosophy, Princeton University, USA.

Scott Fults earned a PhD in linguistics from the University of Maryland, College Park. He is currently a postdoc in the Computer Science Department at Maryland.

Christopher Kennedy is Professor and Chair, Department of Linguistics, University of Chicago, USA. His research addresses questions in semantics, pragmatics and philosophy of language, primarily through an exploration of the grammar of comparison, amount, degree and vagueness.

Nathan Klinedinst earned a PhD in linguistics from UCLA, and is now a Lecturer in Linguistics at University College London. His areas of research interest are formal semantics and pragmatics, with a focus on conditionals and modality.

Rick Nouwen is Senior Research Fellow at the Utrecht Institute for Linguistics OTS, Utrecht University. He has published many papers on the formal semantics and pragmatics of quantification and scalarity. He is currently leading a research project on vagueness and gradability.

Peter Pagin is Professor, Department of Philosophy, Stockholm University. Peter Pagin has worked in many subareas of the philosophy of language. His latest publications are “Vagueness and Central Gaps,” in R. Dietz and S. Moruzzi (eds), *Cuts and Clouds. Vagueness, its Nature and its Logic* (2010); and “Compositionality I: Definitions and Variants, Compositionality II: Arguments and Problems,” both in *Philosophy Compass* 5 (2010): 250–64 and 265–82, respectively.

Galit Sassoon is currently a postdoc at the Institute of Logic, Language and Computation in Amsterdam University, working on vagueness and gradability in natural language and on the interface between psychology and semantics. She has recently published the papers “The Degree Functions of Negative Adjectives” (*Natural Language Semantics* 18(2), 2010) and “Measurement Theory in Linguistics” (*Synthese* 174(1), 2010).

Uli Sauerland leads an Emmy-Noether Research Group at the Centre for General Linguistics (ZAS) Berlin, Germany, and also the Eurocores research project VAAG that investigates vagueness, approximation and granularity. He recently published the volume *Semantics and Pragmatics: From Experiment to Theory* in this series.

Penka Stateva is Associate Professor of Linguistics at the University of Nova Gorica, Slovenia. Her published work focuses on the semantics of comparison, vagueness, plurals, possessives, and the relation between the syntactic and semantic grammar components.

Paula Sweeney is Research and Operations Officer, Department of Philosophy, University of Aberdeen, UK. Her research interests lie within philosophy of language, in particular with issues concerning context-dependence.

Robert van Rooij is a senior staff member of the Institute for Logic, Language, and Computation (ILLC), stationed at the Department of Philosophy of the University of Amsterdam. His research interests

include formal semantics, game-theoretical pragmatics, and the philosophy of language. In these areas he has published several books, numerous journal articles (e.g. *Linguistics and Philosophy*, *Journal of Semantics*, *Journal of Philosophical Logic*), and articles in books and conference proceedings. He is a member of the editorial board of the *Journal of Semantics* and the e-journal *Semantics and Pragmatics*, and editor of the language section of a philosophical journal (*Synthese*) and of a book series (Texts in Logic and Games).

Lavi Wolf is a PhD student at the Department of Foreign Literatures and Linguistics at Ben-Gurion University under the supervision of Dr Ariel Cohen. His current research concerns the effects of assertion on the common ground, addressing topics such as modal adverbs and adjectives, predicates of personal taste, evidentials and clarity.

Elia Zardini is AHRC Postdoctoral Research Fellow at the Northern Institute of Philosophy, University of Aberdeen, UK.

1

Introduction: Vagueness and Language Use

Paul Égré and Nathan Klinedinst

1.1 Defining vagueness

1.1.1 Main criteria

Terms like “red”, “bald” and “young” are *vague*. In formal and scientific approaches to the study of language and mind, this claim is standardly taken to point to three particular (related) features of their meaning or use (see for instance Keefe 2000, Kennedy 2007, Smith 2008). First, such predicates admit of borderline cases, cases to which they neither clearly apply, nor clearly fail to apply. For instance, we would readily admit that a 16-year-old man is young, and that a 90-year-old man is not young, but what about a man of 32? Of 37? Of 41? For cases like those, we may hesitate whether to declare the individual young or not.

The sources of this uncertainty are distinctive and worth pondering. Obviously, there are many predicates that we fail to know the complete extension of; take for example “prime number.” We are often uncertain as to whether a given number is prime or not. However, we seem to know that the extension of “prime number” has sharp boundaries as a matter of principle, and that it is mechanically computable and unambiguous whether a given number falls under the predicate’s extension or not. In contrast, while predicates like “red” or “young” have paradigmatic and consensual cases of application and of nonapplication, there appears to be a region of intermediate cases for which the use of these predicates is shifty, more or less arbitrary, and left at the speaker’s discretion. In general, moreover, this region of intermediate cases itself appears to have no clear-cut boundary: the transition from clear cases of a predicate to clear cases of its negation is only gradual. Without stipulation, for example, it seems hardly possible to specify a unique age at which one would cease to be young. Thus, a second distinguishing feature of vague predicates is

that the transition from clear cases of application to clear cases of exclusion is “seamless.” The boundaries of vague concepts are “blurry,” as it is often put (see Smith 2008).

The lack of sharp boundary for vague predicates is related to a further feature, which concerns what we may call the relative plasticity in the use of these predicates. Vague predicates are often described as *tolerant* (Wright 1976): if someone is counted as young, for instance, it seems that anyone only slightly older (by a small increment of just a few days, or months) should count as young too. Little by little, however, the principle of tolerance implies the problematic conclusion that anyone older than a 16-year-old man should count as young. Thus the (apparent) tolerance of vague predicates gives rise to *sorites paradoxes*. Sorites paradoxes are easily constructed from most gradable expressions, namely expressions which can be mapped onto a degree scale and which thus “typify qualitative linguistic expression of quantitative information” (DeVault and Stone 2004). The abstract formulation of a sorites paradox is given by the mathematical principle of induction, which states that if a predicate P holds for 0, and if it holds for $n + 1$ whenever it holds for n , then the predicate holds for every n . The induction principle correctly applies to a sharp predicate like “uniquely decomposable into prime numbers” (Bonini et al. 1999), but yields counter-intuitive consequences in the case of many ordinary predicates, as illustrated with “young” above, and by the paradox of the heap: 0 grains of sand do not make a heap, and it seems that if n grains of sand do not make a heap, then adding only 1 grain will not make a heap either. By induction, it should follow that for any n , n grains of sand do not make a heap, which is clearly paradoxical.

Thus there are three central and distinctive features of the phenomenon of vagueness: the existence of borderline cases, the lack of a sharp boundary along the transition from clear cases to clear counter-instances, and finally, susceptibility to sorites arguments (see for instance Keefe 2000, Kennedy 2007 and Smith 2008 for detailed overviews).

1.1.2 Vagueness in language

It is worthwhile to point out two further aspects of vagueness in order to clarify the scope and focus of this volume.

First, from a logical point of view, vagueness (in the above sense) is primarily an attribute of concepts, as opposed to sentences in particular. Often, “vagueness” is informally attributed to the information conveyed by particular sentences, typically to stress the fact that this information is imprecise or overly general (think of the answer “someone left” to the question “who left?,” which we would typically call a vague answer).

This sense of vagueness, which we may call *underspecificity* (following Keefe 2000: 10), is not without relation to the one we are concerned with (it is opposed to the idea of an answer that would be as *complete* or *precise* as possible), but *prima facie* it involves the consideration neither of borderline cases, nor of boundaries or soriticality, which fundamentally pertain to concepts and their scalar structure.

Second, from a linguistic point of view, vagueness thus understood is obviously not restricted to adjectives (such as “young”, “bald” or “red”), but is found in many other lexical categories, and probably in all categories for which some notion of grading can be relevant (see Sapir 1944, van Rooij, forthcoming). Thus nouns can be vague (not only *heap*, but most common nouns such as *chair*, *apple*, etc., and even proper names such as *London*), as well as verbs (*walk*, *run*), determiners (*many*, *few*, *much*, *little*), adverbs (*quickly*, *surprisingly*, *clearly*) and modifiers (*very*, *somewhat*, *completely*).

Whether all or only part of ordinary language is vague in the sense we have isolated is a moot question. According to Lewis (1986: 212), we must consider that “not all of language is vague”: logical words such as “not,” “and” and “or,” as well as quantifiers such as “none,” “all” and “some,” appear to have a perfectly sharp meaning. According to Russell (1923: 65), however, logical words “when used by human beings, share the vagueness of all other words.” In particular Russell seems to have considered that the vagueness of the concepts “true” and “false” as applied to nonlogical words was likely to impinge on our understanding of the truth conditions of conjunctive or negated sentences. Nevertheless, Russell admits that logical words are at least *less vague* than nonlogical words, and even that their meaning can be made fully precise when applied to nonvague vocabulary. While more would need to be said about the extent of vagueness in language (see below in section 1.3), we take Lewis’s view to be compelling, in particular because if the language of mathematics can be made entirely precise and exempt of vagueness, this suggests that at least part of our logical vocabulary must be exempt of vagueness from the start.

1.2 Main issues and overview of the literature

1.2.1 Research on vagueness since the 1970s

The phenomenon of vagueness has given rise to a large body of research since the early 1970s, in the areas of natural language semantics, philosophy, cognitive psychology and artificial intelligence. From a psychological point of view, one of the main problems raised by the

phenomenon of vagueness concerns the norms of similarity that are involved in our ability to categorize, both conceptually and at the lexical level (see Kamp and Partee 1995, Hampton 2007). From a semantic point of view, vagueness raises the question of which framework might be the most suitable to represent the absence of sharp boundaries for most of our predicates. Classical logic is well suited to deal with sharp predicates, but it offers no direct handle on vagueness: it was developed precisely to deal with the language of exact sciences and with determinate concepts (see Frege 1879).

From the 1960s onward, several alternative frameworks have been developed and employed to deal with vagueness and the sorites paradox, in particular many-valued logics (such as fuzzy logics, see Goguen 1969, Lakoff 1973, Zadeh 1975) and partial logics (in particular supervaluationism, see Lewis 1970, Fine 1975, Kamp 1975). During the 1990s, under the influence of the philosophical work of R. Sorensen (1988, 2001) and T. Williamson (1994), growing attention has been paid to the epistemological and psychological sources of vagueness, in relation to the idea that vagueness might be more fundamentally an epistemic phenomenon (of ignorance, or inexactness) rather than a semantic primitive (such as partiality, or underdetermination of meaning, as in the supervaluationist tradition). Concurrently, various proposals have been made to stress the role and importance of context in the mechanisms of comparison involved in language and perception, as well as for the resolution of the sorites paradox (Sapir 1944, Klein 1980, Kamp 1981, Bosch 1983 are pioneering references on vagueness and context-dependence well before the 1990s; Raffman 1994, 1996, van Deemter 1996, Soames 1999, Fara 2000, Shapiro 2006, Kennedy 2007 represent various more recent proposals at articulating contextualist intuitions about vagueness).

1.2.2 Recent trends in philosophy and linguistics

Several developments in philosophy and linguistics have particularly driven an expansion of interest and debate in the past decade, not only within their respective disciplines, but across them. On the philosophical side, the epistemic theory of vagueness, which makes the forceful claim that vagueness does not call for a revision of classical logic, has generated much debate, often in relation to the sorites paradox (see Wright 1994, Schiffer 2003, Field 2008). On the epistemic conception of vagueness, the tolerance principle is considered false for vague predicates. By classical reasoning, the rejection of tolerance implies the counter-intuitive consequence that vague predicates do have a sharp cutoff as a matter of fact. To make this consequence palatable, however, the epistemic theory

of vagueness proposes to derive our inability to locate this boundary from principles of imperfect discrimination that are logically weaker than the tolerance principle, but still close enough to its original formulation (so-called *margin for error principles*, see Williamson 1994, and Kennedy this volume). On the epistemic conception, more generally, the facts about our use of language that settle the boundary of our vague predicates are too complex for us to know with sufficient accuracy, and can be used to explain why we are uncertain when faced with borderline cases.

While the epistemic conception of vagueness is often met with skepticism in philosophical discussions surrounding the sorites paradox, the simplicity of its architecture has also gradually drawn the attention of linguists. Concurrently, linguistic analysis of the syntax and semantics of comparison and gradable adjectives (see Kennedy 1999, Fara 2000, Kennedy and McNally 2005, Kennedy 2007) has provided new data and further insight into theories of vagueness, in particular by casting light on the structure of scales associated with vague expressions, the compositional semantics of a range of constructions involving gradable predicates, and the connection between vagueness and context-dependence. On the influential accounts of Fara and Kennedy, the model theory of vague predicates is based on classical, bivalent logic, and implies that these predicates have a sharp cutoff. However, the location of the cutoff in question is highly context-dependent, and our ignorance of the cutoff can be described as ignorance regarding the exact value of these contextual parameters (see in particular Fara 2000, or Krifka 2007 for a semantics of antonyms based on the epistemic conception of vagueness; see Kennedy, this volume, for discussion of relevant differences between Williamson's epistemicism and Fara's interest-relative account).

To give an example of this connection between vagueness and context-dependence, let us consider the semantics of gradable adjectives like "tall" or "young." "Tall" is context-sensitive in the sense that whether an individual is judged tall or not depends on an underlying comparison class (whether a snowman counts as "tall" varies depending on whether it is built by a child or by a fraternity; Klein 1980, Kamp and Partee 1995). Relative to a given comparison class moreover, "x is tall" might be analyzed as meaning that "x's height is taller than the average height within the class" (Bartsch and Vennemann 1972), or rather that "x's height is (significantly) taller than a contextually salient standard of comparison" (Fara 2000, Kennedy 2007). Saliency is itself typically a vague and elusive notion, as is the notion of "significant" difference (see Fara 2000 on the relation of this notion to practical interests, and Kennedy this volume, Sweeney and Zardini, this volume). But even when

the theoretical position of the threshold is semantically determined and the discriminatory unit specified, some vagueness can remain. Suppose that “taller” were to mean “taller than the average” (see Kennedy 2007 for arguments against this simplification). In many cases, it may be practically impossible to compute where the average lies in the class. Even supposing the mean value to be known, it may still be hard in some cases to appreciate whether the individual lies above or below the standard. The applicability of a predicate like “tall” is thereby multiply relative: to a comparison class, to the position of a threshold in the class, to the notion of critical departure from the threshold, and ultimately to the discriminatory resources of the subject. At each of these levels, the dependence on context may therefore be seen as a further source of complexity and uncertainty regarding the applicability of the predicate.

Another recent strand of research on vagueness in linguistics deals with the semantics and pragmatics of imprecision and approximation (see Pinkal 1995, Lasersohn 1999, Kennedy 2007, Sauerland and Stateva 2007 and this volume). On a naive view, imprecision might be seen as the essence of vagueness. For instance, one may consider that geometrical concepts such as “square” or “circle” have a perfectly precise denotation, but that in practice, it is vague whether an object is a circle or not because we never deal with perfect squares or perfect circles, but only with approximations to perfect squares or perfect circles. A view of vagueness that would reduce it entirely to imprecision along these lines may be inadequate, however, simply because not all concepts are so easily amenable to mathematical precision. As Kennedy puts it, it is probably more accurate to consider that vagueness directly pertains to *meaning*, while imprecision concerns the *use* of expressions with a precise meaning.

For instance, “12 o’clock” may denote a unique point on the clock, or a sharp interval (say the minute between 12:00 included and 12:01 excluded), but if we consider an adjective like “tall,” there is no unique candidate interval to serve as a fixed denotation, and the meaning of “tall” can be described as vague, unlike that of “12 o’clock.” On the other hand, we can say that John arrived at 12 o’clock when in fact John really arrived at 7 minutes past 12. This suggests that expressions with a precise meaning can be used with slack, depending on how much people are willing to coarsen precise meanings in conversational contexts (see Lasersohn 1999, Sauerland and Stateva this volume).

The same contrast can be observed among quantifiers in natural language. A determiner such as “all” appears to have a precise meaning: “all students left” means that the set of students is included in the set of

people who left. Obviously, the sentence could be uttered in a situation in which all but one student left, if one wants to express disappointment at such a large number of defections. This imprecision, however, does not appear to call into question the precise semantics of “all.” In contrast, a determiner like “many” is more radically vague and context-dependent. For a sentence like “many As are Bs,” there is no fixed number or ratio that can be used across contexts to fix the denotation of the term (see Partee 1989, Lappin 2000). In that sense, the imprecision that attaches to some uses of a precise quantifier like “all” is distinct from the vagueness that attaches to the meaning of “many.”

1.2.3 Vagueness and related notions

The idea that vagueness is not wholly reducible to *imprecision* is relatively uncontroversial today, even though the study of imprecision remains one important aspect of the understanding of vagueness (see in particular Sauerland and Stateva this volume). However, there remains some debate concerning the identification of vagueness with other semantic notions in its vicinity. Let us consider a few of those here.

In semantic theories of vagueness such as supervaluationism, vagueness is essentially tied to the notion of *underdetermination* and *openness* of meaning (see Fine 1975, Keefe 2000, Shapiro 2006, Cobreros this volume). On the supervaluationist account of the meaning of a word like “tall,” the existence of borderline cases for “tall” is related to the possibility of precisifying its meaning in different ways. This idea is crucially distinct from the contextualist position, on which vagueness would derive from the *contextual variability* of a determinate threshold. Although Williamson’s epistemic theory of vagueness also posits a determinate threshold, it contrasts with contextualist positions (such as those of Fara and Raffman) in its basic explanation of vagueness. Williamson’s theory ties vagueness to the notion of *inexactness*, and in particular takes it to fundamentally issue from nontransitivity in discrimination. The link between vagueness and nontransitivity plays a central role in several theories of vagueness and can be elaborated in various ways, however (see in particular Goodman 1951 and Luce 1956 for groundbreaking accounts, and the recent theories of Halpern 2008 and van Rooij forthcoming, and this volume). It should be stressed that the view of vagueness as nontransitivity is relatively neutral regarding the commitments of a theory of vagueness, in particular concerning the location and shape of boundaries for vague concepts.

The neighboring concepts of imprecision, underdetermination, contextual variability, and inexactness are by no means exhaustive of the

ways in which vagueness is approached or conceptualized in the literature. In particular, vagueness is sometimes seen also as the expression of *overdetermination*, *overlap* or *ambivalence* between categories, rather than of underdetermination (as in so-called “glut” theories and paraconsistent approaches more generally: see for instance Hyde 1997 on subvaluationism, an approach dual to supervaluationism, and Beall and van Fraassen 2003 for an overview of many-valued logics involving paraconsistency). Overall, however, we take the connections we listed above to be main avenues along which the phenomenon of vagueness is declined or specified in the literature. It is useful to bear in mind the distinctions between these concepts to appreciate which aspect or aspects of the phenomenon of vagueness are particularly emphasized by a given theory.

1.3 Contents of the volume

1.3.1 Overview and general organization

As the title of this volume indicates, the essays assembled herein are mostly concerned with the manifestations of vagueness in language. Upon reflection, one may wonder how it is possible that vagueness does not radically compromise our ability to successfully communicate. Various answers to this question are conceivable. One hypothesis is that vagueness turns out to be much more limited than what would seem at first (see the quote by Lewis above); another possibility is that we can rely on a number of semantic and pragmatic mechanisms that compensate for the vagueness of most of our expressions; yet a third view, finally, could be that lexical vagueness is itself in fact a key feature to any successful act of linguistic communication (see for instance Parikh 1994, van Deemter 2010).

The contributions in this volume can be broadly related to the second of these hypotheses, namely they explore some of the mechanisms by which vagueness in language is regulated. In the first part of the book, in particular, the chapters deal with the link between vagueness and comparison in language. The leading thread between the essays concerns the issue of inexactness that was raised previously. In particular, all chapters in that section (by Fults, van Rooij, Kennedy and Sassoon) examine the link between exact and inexact measurement and their counterparts in language.

The chapters in Part II deal with the phenomena of approximation and intensification in language. The link between the chapters in this section (by Sauerland and Stateva, Nouwen, Wolf and Cohen, and Barker) concerns the role played by vagueness-related adverbs and modifiers, namely

adverbial expressions whose semantic role is to qualify the degree to which a property holds, or more generally that serve to modulate the standard of precision or comparison relevant for its application.

Part III finally collects essays dealing with the sorites paradox, and which concern specific aspects of the supervaluationist and contextualist approaches to vagueness. The first two chapters (by Cobreros and Fara) deal with the problem of higher-order vagueness and with the treatment of this problem in supervaluationist semantics. The two chapters that follow (by Sweeney and Zardini, and by Pagin) discuss the prospects for contextualist and pragmatic treatments of the sorites paradox.

In each of the three parts, finally, there are duos or trios of chapters that address one another, and which the reader is advised to read jointly: in Part I, Fults, van Rooij and Kennedy's chapters on the semantics of comparatives; in Part II, Wolf and Cohen's chapter on the assertion of clarity, and the response by Barker; in Part III, Cobreros's defense of supervaluationism about higher-order vagueness, and Fara's criticisms of that strategy.

1.3.2 Measurement and comparison

Comparative constructions in language provide an example of a general mechanism whereby vagueness can be reduced and precision increased. Quine (1987: 109), for instance, writes:

there is no place in science for bigness, because of this lack of boundary; but there is a place for the relation of biggerness. Here we see the familiar and widely applicable rectification of vagueness: disclaim the vague positive and cleave to the precise comparative.

This semantic contrast between the positive and the comparative form of gradable adjectives can be evidenced in a number of ways. For example, if we consider a gradable adjective such as "young," and consider only age as the relevant dimension of comparison, we saw earlier that "*x* is young" is typically hard to adjudicate for some cases. In contrast, the comparative sentence "*x* is younger than *y*" is usually unambiguously true or false, including for the same cases that would count as borderline for the positive form, provided we know the ages of *x* and *y*.

The idea that comparative constructions reduce or even eliminate vagueness is widely accepted, but it deserves further examination. Keefe (2000: 13) considers that comparatives can still be vague and have borderline cases, for instance, in the case of "young," if we are uncertain about how to order moments of birth. The same, arguably, could happen

for sizes, if some individuals are so close to each other that we cannot say with enough accuracy whether one is bigger than the other, or of identical size. Another source of problems for comparatives concerns multidimensional adjectives (such as “clever” or “qualified”), which may give rise to uncertainty concerning which dimension is most relevant (see Sassoon this volume), and even to quandaries and intransitivities of a kind familiar in social choice theory (a may be judged more clever than b along one dimension, b more clever than c along another, and c more clever than a along a third; see e.g. van Deemter 2010: 47–51). Arguably, however, these complexities do not call into question the semantic contrast between the positive and the comparative form of gradable adjectives. For the point is that if we consider only one dimension of comparison, and suppose perfect accuracy in measurement, vagueness does affect the extension of the positive form of the adjective, but not that of the comparative.

Two specific issues are addressed in Part I of this book regarding the semantics of comparison. The first concerns the nature of the relation between the positive and the comparative forms. Morphologically, the comparative “younger” is more complex than the positive “young,” and appears to be derived from it. Semantically, however, semantics of gradable adjectives based on degree scales (see in particular Bartsch and Vennemann 1972, Kennedy 2007, and further references in Kennedy, this volume) go in the reverse direction, treating the positive form as effectively a covert comparative. On this view, “young” basically means “younger than a context-dependent threshold,” and a silent morpheme (*pos*) is postulated to derive this meaning in a compositional way. A different approach, found for instance in Wheeler (1972), Kamp (1975) or Klein (1980), is to treat the positive form as semantically primitive, and to try to derive the comparative from it.

A second and related issue concerns the link between (explicit) comparatives like “John is taller than Mary,” and *implicit comparatives* (following Kennedy’s terminology, itself based on Sapir 1944) – sentences like “John is tall in comparison to Mary,” which use the positive form rather than comparative morphology. If John is only 1 cm taller than Mary, one can hardly say “John is tall in comparison to Mary.” However, it remains fine to say simply that “John is taller than Mary.” In the former case, one expects the sentence to be true only if John is *significantly* taller than Mary, suggesting that Mary’s height sets a first standard relative to which John’s height is to be evaluated to count as significantly taller.

Both issues, the relation between the positive form and the comparative form, and the related contrast between explicit and implicit

comparatives, are the topic of the first three chapters of this volume. Scott Fults argues that the semantics of implicit comparatives relies on a system of comparison essentially distinct from the one on which explicit comparatives operate. On Fults' account, the positive form of gradable adjectives is evaluated relative to an analog scale, akin to the approximate number system used for the evaluation of mental magnitudes. By contrast, explicit comparatives are taken to rely on precise scales of measurement.

Van Rooij's proposal goes in a similar direction, but using a different set of premises. On van Rooij's approach, implicit comparatives involve *semi-orders*, while explicit comparatives involve *weak orders*. Semi-order relations, originally defined by Luce to account for intransitive indifferences, are relations intended to express that a quantity is "significantly greater" (or smaller) than another. In that sense, semi-orders provide a qualitative notion of comparison that purports to account for inexact measurement, since two quantities can differ objectively and yet be such that neither is felt or judged significantly greater or smaller than the other. The gist of van Rooij's proposal is moreover to derive both semi-orders and weak orders from a semantics of the positive form of gradable adjectives based on comparison classes, rather than degrees, by means of appropriate postulates on how comparison classes work.

Van Rooij's theory allows him both to account for the contrast between implicit and explicit comparatives, and to restore the intuition that the positive form of gradable adjectives, which incorporates vagueness, is semantically primitive over the explicit comparative form. Some potential limitations of van Rooij's account are discussed in Kennedy's contribution to this volume. Unlike van Rooij, Kennedy does not assume that comparatives, whether explicit or implicit, should necessarily be compositionally derived from a root semantics for the positive. On Kennedy's view, the positive "tall" is covertly comparative, but the silent morpheme *pos* which it incorporates involves the notion of a *significant* difference, in contrast to the explicitly comparative morpheme *-er/more* of English. Kennedy furthermore gives a typological outlook on the difference between explicit and implicit comparatives, and examines the incidence of the distinction for epistemicist as well as supervaluationist accounts of vagueness.

Galit Sassoon's chapter, which closes this section, focuses on a different aspect of the contrast between positive and comparative forms of gradable adjectives. Sassoon uses the idea that comparatives are precise and positives vague to make a distinction between semantic vagueness and vagueness as ignorance. On her account, we may fail to know whether

“Sam is tall,” even when we know Sam’s height, because the cutoff for “tall” is indeterminate. On the other hand, we can also fail to know whether “Sam is taller than Dan,” but this time it should be because we lack information about the heights of Sam and Dan. Sassoon argues that both phenomena can be integrated within a supervaluationist approach on which individuals are identified purely by their *property values*, namely by the measurable extents to which they satisfy properties. More generally, she suggests that our ignorance about the property values of individuals supports a non-Kripkean theory of singular terms, whereby those can denote distinct individuals in different precisifications.

1.3.3 Adverbial modification, approximators and intensifiers

Comparative constructions are not the only mechanisms available to reduce or regulate vagueness in language. Part II of the volume is devoted to the discussion and exploration of adverbs and modifiers in relation to vagueness. In the literature, these expressions are often referred to under the name of “hedges,” following Lakoff (1973). Lakoff called hedges “words whose meaning implicitly involves fuzziness – words whose job is to make things fuzzier or less fuzzy” (1973: 479). Hedges in the sense of Lakoff include a variety of modifying expressions, such as *very*, *roughly*, *completely*, but also expressions like *sort of*, *strictly speaking*, *loosely speaking*, and so on. Sauerland and Stateva (this volume) refer to these expressions under the generic name of “approximators.” Some of these expressions, like “very,” “definitely,” “completely” and so on, are more commonly or readily referred to under the name of “intensifiers” (see e.g. Soames 1999).

Whichever cover term one may prefer, the function of these expressions is thus to increase or to relax precision, and more generally to qualify the degree to which a property holds. Importantly, as the original definition of Lakoff suggests, these expressions themselves are usually not exempt of vagueness. For instance, an adverb like “clearly” inherits the vagueness of the adjective “clear.” Just like the vague adjective “tall,” “clearly tall” still selects a vague region. Nevertheless, the modifier corresponds to an instruction that implies selecting a particular subregion of the region selected by “tall” (namely individuals tall to a high degree). Similarly to “clearly,” a modifier like “very” is also vague. The vagueness of “very” can be evidenced when applied to what Kennedy and McNally (2005), following Unger (1975), have called “absolute gradable adjectives,” namely adjectives that select an endpoint on a degree scale. On this account, “full” is taken to select the maximum degree to which a container can be filled (and is not considered vague for that matter

by Unger or Kennedy). However, “very full” typically implies “less than completely full,” to a measure that is unspecified. By contrast, as applied to a relative gradable adjective such as “tall,” which would correspond to an open scale (see Kennedy 2007), “very tall” instructs to look for a range of individuals that are tall to a significantly higher degree than that relevant for the unmodified “tall” (note in particular that the prepositional phrase “in comparison to Mary” in “John is tall in comparison to Mary” is an approximator in the sense intended by Sauerland and Stateva). Unlike comparative morphemes more generally, approximators do not semantically eliminate or neutralize vagueness, but they serve a related purpose, namely they serve to select subsets or supersets of the default extensions selected by the predicates they modify.

The chapters in Part II deal with several issues concerning the typology of approximators, and the varieties of adverbial modification to which they correspond. Sauerland and Stateva in their essay adduce empirical evidence in favor of the distinction outlined above between imprecision and vagueness proper. They examine the distribution of a large sample of approximators and show that one class (such as “exactly”), which they call *scalar approximators*, can only apply to expressions with a perfectly precise meaning, or which select a completely precise point on a scale (viz. “exactly 50 persons”). Whereas a second class, which they call *epistemic approximators* (“definitely”), modify expressions that are vague (“definitely a house”). Sauerland and Stateva’s data can be seen as giving further evidence for a general distinction congruent to the one between relative and absolute gradable adjectives investigated by Kennedy and McNally (see Kennedy this volume). In the case of scalar vagueness, moreover, they compare several ways in which precise extensions can be semantically tightened or coarsened. Importantly, the notion of epistemic approximator on their account is meant to refer to adverbs whose semantics refers to uncertainty (“maybe, definitely, clearly”). By extension, and in relation to the epistemic account of vagueness, they call epistemically vague predicates about whose cutoff point we are uncertain.

Nouwen’s chapter in turn proposes some typological generalizations based on the behavior of evaluative adverbs such as “surprisingly,” “amazingly,” or “unusually.” Nouwen focuses on the contrast between the role of these adverbs as sentence modifiers (as in ‘surprisingly, John is tall’), and as predicate modifiers (as in “John is surprisingly tall”). Qua predicate modifiers, these adverbs act as degree intensifiers, namely they express that the predicate holds to a high degree, but not so qua sentence modifiers. Nouwen gives a compositional derivation of the relation

between the meaning of these adverbs as sentential and as predicate modifiers. In particular, Nouwen makes the assumption that all gradable predicates are monotone in order to derive several asymmetries between the positive and negative form of such adverbs (as in “usually tall” vs “unusually tall”).

Nouwen’s essay closes with some remarks about the adverb “clearly,” which occupies a central role in several theories of vagueness. In particular, as we saw, the adverb “clearly,” just as the adverb “definitely,” is used to characterize borderline cases, as cases that are neither clearly/definitely P nor clearly/definitely not P . In the philosophical literature on vagueness, the behavior of these adverbs is generally approached from a logical angle, in particular to deal with the problem of higher-order vagueness (see in particular the Appendix to Williamson 1994, Barker 2002, and below in Part III, the contributions by Cobreros and Fara; see also the remarks by Sassoon, and Sauerland and Stateva, this volume). Another aspect of the behavior of the adverb “clearly” concerns its intensifying role in relation to assertion. The issue has been particularly highlighted in recent work by Chris Barker (see Barker and Taranto 2003, Barker 2009) and concerns the use the adverb “clearly” as a sentential operator, and the relation of the adverb to other modalities (such as knowledge, belief, and epistemic modals like “must”).

Barker’s theory is that in asserting clarity we express that the degree of public justification of a proposition is above a vague standard. In their contribution, Wolf and Cohen defend a view on which the point of asserting clarity is primarily to express belief, rather than justification, in the proposition in question. However, they maintain that the belief in question makes normative reference to what ideal believers would believe. Unlike Barker, in particular, they argue that having a high personal degree of belief in the proposition is at least a necessary condition to assert clarity, though not a sufficient one. They furthermore contrast two particular sentential constructions based on the adjective “clear”: the adjectival construction “it is clear that p ” and the adverbial construction “clearly, p .” On their theory, based on a probabilistic semantics, the adverb “clearly” only modifies the strength of the assertion, while “it is clear that” furthermore expresses that the proposition is believed to a high degree by ideal reasoners.

In his response, Barker considers several reasons why an account based on justification and public evidence might not need to appeal to belief in the first place. Barker argues, in particular, that a belief-based account might be too liberal in cases in which one’s degree of belief could be made very high, but in which the actual evidence is yet missing or

not available, typically lottery cases involving assertions about winning tickets (see Kyburg 1961, Hawthorne 2004).

1.3.4 The sorites paradox

The status of the adverb “clearly” makes the connection between Parts II and III of the volume, which concerns more foundational issues, in particular ways of dealing with the sorites paradox.

As already emphasized, the debate between Wolf–Cohen and Barker concerns the status of clarity as predicated of full propositions, namely it concerns the occurrences of the adverb “clearly” in front of full sentences, as in “clearly, your ticket will win.” As mentioned by Nouwen, however, even when “clearly” behaves syntactically as a predicate modifier, it seems that it can be equally given scope over the whole sentence (compare “your ticket will clearly win,” where “clearly” appears in a low position, to “clearly, your ticket will win”). A point worth stressing, however, is that in nearly all the examples discussed by Wolf–Cohen and Barker, the problems would remain essentially the same if the propositions modified by the adverb “clearly” did not involve any vague predicates at all (contrast with Barker 2002 who examines “clearly” and “definitely” as modifying vague predicates). For instance, “your ticket is a winning ticket” typically does not express a vague proposition, for every ticket is taken to be definitely a loser or a winner (compare with the remarks of Soames (1999: 221) on “definitely” and “clearly”: “Although *clearly* and *definitely* are often used with vague or partially defined predicates, there is no requirement that the predicate be such.”) This helps to emphasize why Barker questions the contrast between plain assertion and the assertion of clarity: when propositions have a perfectly determinate truth value in principle, for instance mathematical truths, the point of asserting their clarity cannot be to intensify the truth status of the proposition, as it were, but presumably concerns another dimension (such as justification, or belief, or strength of assertion).

In contrast to that, the debate between Cobreros and Fara concerns the specific contribution of the adverb “definitely” when it modifies a vague predicate (as in “John is definitely tall”). In this case, the function of “definitely” as applied to “tall” is to denote a subset of the set of tall people (namely people with a sufficiently high degree of tallness). More specifically, the topic of Cobreros’ chapter and of Fara’s response is the problem of higher-order vagueness. In the literature on vagueness, the problem is stated equivalently in terms of the adverbs “definitely,” “clearly” or “determinately” (compare Williamson 1994, Soames 1999, 2003, Fara 2003). Consider a vague predicate like “tall.”

First-order vagueness corresponds to the idea that some cases, namely borderline cases of tallness, are cases that are neither clearly or definitely tall, nor clearly or definitely not tall. Second-order vagueness is the idea that the boundary between clear cases and borderline cases is in turn vague. Thus some cases must be neither definitely tall, nor definitely not tall. Higher-order vagueness is the generalization of this idea to further embeddings of the adverb “definitely” (or “clearly” or “determinately”). Higher-order vagueness is generally assumed to be the necessary counterpart of the notion of “seamless” transition between clear instances and clear counter-instances of a vague predicate.

The problem discussed by Cobreros is a version of the sorites paradox put forward by Fara (2003) concerning higher-order vagueness, in relation to the supervaluationist treatment of vague predicates and of the operator “definitely.” Fara’s argument, elaborating on a reductio argument against higher-order vagueness originally due to C. Wright (see Wright 1992, 2010), shows that certain *gap principles* used to capture higher-order vagueness in terms of the “definitely” operator are contradictory with a particular rule of inference validated in standard supervaluationist semantics, the rule of D-introduction. This rule, which constitutes a direct analog of the rule of necessitation in modal logic, implies that if “ x is tall” is true, then “ x is definitely tall” will be true. Fara (2003) originally used the paradox as an argument against the adequacy of supervaluationism. In his chapter, Cobreros proposes a way out of this paradox based on a modification of the standard supervaluationist notion of truth, called *regional truth*. Relative to this concept, the inference from “ x is P ” to “ x is definitely P ” is no longer valid, but still all gap principles can be made compatible with the consideration that the ends of a sorites series are absolutely definite instances of a predicate and of its negation.

In her response to Cobreros, Fara maintains her initial criticisms of the supervaluationist framework (in line with other recent objections she addresses to supervaluationism, see in particular Fara 2010). She first reviews and compares the various notions of truth at stake in standard supervaluationism and in Cobreros’ revision of it, namely bivalent truth, super-truth and regional truth. She formulates several criticisms regarding the adequacy of the notion of regional truth. Her main objection, eventually, is an argument purporting to show that the region-valuationist, just as standard supervaluationists more generally, is committed to the inference from “ x is tall” to “ x is not borderline tall.” Based on this inference, and on the definability of definiteness in terms of borderlineness, she concludes that the region-valuationist’s invalidation of the rule of D-introduction cannot be used as a satisfactory

solution to the paradox. Fara's conclusion to her essay is that a theory of borderline cases based on truth-value gaps is not an adequate theory of vagueness.

While Fara does not state her own account of vagueness in her chapter, Fara (2000) contains an elaborate version of her positive theory, and argues in favor of a contextualist view of vagueness. In that paper, in particular, Fara defends a classical and bivalent semantics for vagueness, which implies the rejection of the tolerance principle. She argues for several principles intended to explain both why we are psychologically inclined to take the tolerance principle to be valid, and also why we are unable to know the cutoff in a sorites series, even though one is assumed to exist.

The critical examination of Fara's theory is in turn the focus of Sweeney and Zardini's contribution to this volume, in which they formulate a series of objections against the adequacy of contextualist treatments of vagueness more generally. In their chapter, Sweeney and Zardini closely discuss the status of one central constraint that Fara puts forward in her theory to account for the sorites paradox, called the constraint of Salient Similarity. This constraint states a principle that is supposed to be weaker than the tolerance principle, namely that if two things are *saliently* similar in a given context, then they are both *P* together or both not *P* together. Fara uses this principle to argue for the fact that cutoffs can exist where things are similar but not saliently so. She also relates it to the tendency for particular instances of the tolerance principle to be compelling. One of Sweeney and Zardini's main criticisms, on the other hand, is that it is not clear that the constraint of Salient Similarity can defuse soritical contradictions. When it appears to do so, for instance when subjects have to deal with soritical series involving a small number of elements, Sweeney and Zardini argue that saliency does in fact defeat the idea that we are necessarily ignorant of cutoffs. More generally, Sweeney and Zardini argue that the constraint of Salient Similarity is too weak or too strong in regard to some of the features of vagueness that Fara's theory purports to account for. While the contribution of their chapter remains mostly critical, Sweeney and Zardini's conclusion is that an adequate theory of vagueness and the sorites should either go on to revise classical logic, or at any rate, should cast doubt on the claim that cutoffs are unknowable.

The last essay in this volume, by Peter Pagin, examines a different kind of strategy for dealing with the sorites. Pagin considers that the right approach to the sorites paradox is in substance pragmatic, and that the paradox can be dealt with without modifying the underlying logic. Unlike Fara, however, Pagin does not exactly rely on the idea that our

use of vague predicates implies the move of a “thin” context-dependent boundary, nor on a constraint akin to Fara’s Similarity Constraint. Rather, Pagin considers that the tolerance principle is distinctive of vague predicates, but that our use of those predicates is systematically regulated by two parameters, which concern standards of precision relevant for the application of the predicate on the one hand, and the assignment of a central gap or “thick” separation between clear positive and clear negative instances of the predicate on the other.

Thus Pagin observes that on at least one way of formulating the tolerance principle, no paradox will result if tolerance is evaluated relative to a sequence of individuals in which at least two consecutive individuals in the series are more distant than the standard relevant for the application of the predicate. In the case of “tall,” a central gap corresponds to a difference in height that exceeds the size of the step used to formulate the tolerance principle. Based on this observation, Pagin’s suggestion is that ordinary uses of gradable adjectives and vague predicates can satisfy the tolerance principle without contradiction, provided our use of these predicates pragmatically presupposes central gaps. Pagin suggests moreover that the mechanism by which central gaps are used can be compared to familiar mechanisms of domain restriction for quantifiers. Interestingly, Pagin’s approach also communicates with the chapters included in the first part of this volume: like Fults and van Rooij, in particular, Pagin starts out from considerations about measurement theory to discuss the notion of tolerance; his strategy is furthermore echoed in van Rooij’s approach based on comparison classes.

Whether any of the approaches discussed in this volume – the supervaluationist, the contextualist, the pragmatic – ultimately succeeds in dealing with the sorites, or whether more radical departures from classical logic are necessary, is a question we shall not try to adjudicate in this introduction. At any rate, we are pleased to stress the fact that all essays here assembled enable us to see more clearly into the nature of vagueness, as well as into some of the mechanisms by which vagueness can be tamed in ordinary language.

References

- Barker, C. (2002). The dynamics of vagueness. *Linguistics and Philosophy*, 25(1): 1–36.
- Barker, C. (2009). Clarity and the grammar of skepticism. *Mind and Language*, 24(3): 253–73.

- Barker, C. and Taranto, G. (2003). The paradox of asserting clarity. In Koskinen, P. (ed.), *Proceedings of the Western Conference on Linguistics (WECOL) 2002*, pp. 10, 21.
- Bartsch, R. and Vennemann, T. (1972). The grammar of relative adjectives and comparison. *Linguistische Berichte*, 20: 19–32.
- Beall, J. and van Fraassen, B. (2003). *Possibilities and Paradox: an Introduction to Modal and Many-Valued Logic*. Oxford University Press.
- Bonini, N., Osherson, D., Viale, R., and Williamson, T. (1999). On the psychology of vague predicates. *Mind & Language*, 14(4): 377–93.
- Bosch, P. (1983). Vagueness is context-dependence. In Ballmer, T. T. and Pinkal, M. (eds), *Approaching Vagueness*. North-Holland Publishing Company.
- DeVault, D. and Stone, M. (2004). Interpreting vague utterances in context. In *The 20th International Conference on Computational Linguistics*. University of Geneva.
- Fara, D. (2000). Shifting sands: an interest-relative theory of vagueness. *Philosophical Topics*, 28(1): 45–81. (Originally published under the name Delia Graff.)
- Fara, D. (2003). Gap principles, penumbral consequence, and infinitely higher-order vagueness. In Beall, J. (ed.), *Liars and Heaps: New Essays on Paradox*, pp. 195–221. Oxford University Press.
- Fara, D. (2010). Scope confusions and unsatisfiable disjuncts: two problems for supervaluationism. In Dietz, R. and Moruzzi, S. (eds), *Cuts and Clouts. Vagueness, its Nature and its Logic*, pp. 128–50. Oxford University Press.
- Field, H. (2008). *Saving Truth from Paradox*. Oxford University Press, Oxford.
- Fine, K. (1975). Vagueness, truth, and logic. *Synthese*, 39 (265–300).
- Frege, G. (1879). *Begriffsschrift. Eine der arithmetischen nachgebildete Formelsprache des reinen Denkens*. Halle.
- Goguen, J. (1969). The logic of inexact concepts. *Synthese*, 19: 325–73.
- Goodman, N. (1951). *The Structure of Appearance*. Synthese Library.
- Halpern, J. (2008). Intransitivity and vagueness. *The Review of Symbolic Logic*, 1(4): 530–47.
- Hampton, J. (2007). Typicality, graded membership and vagueness. *Cognitive Science*, 31: 355–83.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford University Press, Oxford.
- Hyde, D. (1997). From heaps and gaps to heaps of gluts. *Mind*, 106: 641–60.
- Kamp, H. (1975). Two theories about adjectives. In Keenan, E. (ed.), *Formal Semantics of Natural Language*. Cambridge University Press.
- Kamp, H. (1981). The paradox of the heap. In Mönnich, U. (ed.), *Aspects of Philosophical Logic*, pp. 225–77. Reidel.
- Kamp, H. and Partee, B. (1995). Prototype theory and compositionality. *Cognition*, 57: 129–91.
- Keefe, R. (2000). *Theories of Vagueness*. Cambridge University Press.
- Kennedy, C. (1999). *Projecting the Adjective: the Syntax and Semantics of Gradability and Comparison*. Garland, New York.
- Kennedy, C. (2007). Vagueness and grammar: the semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy*, 30(1): 1–45.
- Kennedy, C. and McNally, L. (2005). Scale structure, degree modification, and the semantics of gradable predicates. *Language*, 81(2): 345–81.
- Klein, E. (1980). A semantics for positive and comparative adjectives. *Linguistics and Philosophy*, 4: 1–45.

- Krifka, M. (2007). Negated antonyms: creating and filling the gap. In Sauerland, U. and Stateva, P. (eds), *Presupposition and Implicature in Compositional Semantics*, pp. 163–77. Palgrave Macmillan.
- Kyburg, H. (1961). *Probability and the Logic of Rational Belief*. Wesleyan University Press, Middleton.
- Lakoff, G. (1973). Hedges: a study in meaning criteria and the logic of fuzzy concepts. *Journal of Philosophical Logic*, 2 (458–508).
- Lappin, S. (2000). An intensional parametric semantics for vague quantifiers. *Linguistics and Philosophy*, 23: 599–620.
- Lasersohn, P. (1999). Pragmatic halos. *Language*, 75(3): 522–51.
- Lewis, D. (1970). General semantics. *Synthese*, 22 (18–67).
- Lewis, D. (1986). *On the Plurality of Worlds*. Blackwell.
- Luce, R. (1956). Semiorders and a theory of utility discrimination. *Econometrica*, 24: 178–91.
- Parikh, R. (1994). Vagueness and utility: the semantics of common nouns. *Linguistics and Philosophy*, 17(6): 521–35.
- Partee, B. (1989). Many quantifiers. In *Proceedings of the 5th Eastern States Conference on Linguistics*, Powers, J. and de Jong, K. (eds), pp. 383–402. Ohio State University, Columbus. Reissued in *Compositionality in Formal Semantics, Selected Papers by Barbara H. Partee*. Blackwell, 2004, ch. 12.
- Pinkal, M. (1995). *Logic and Lexicon*. Kluwer, Dordrecht.
- Quine, W. V. (1987). *Quiddities*. The Belknap Press of Harvard University Press.
- Raffman, D. (1994). Vagueness without paradox. *Philosophical Review*, 103: 41–74.
- Raffman, D. (1996). Vagueness and context-relativity. *Philosophical Studies*, 81: 175–92.
- Russell, B. (1923). Vagueness. *Australasian Journal of Philosophy and Psychology*. 1. Reissued in R. Keefe and P. Smith, *Vagueness: a Reader*, MIT Press, 1997, pp. 61–8.
- Sapir, E. (1944). Grading: a study in semantics. *Philosophy of Science*, 11: 93–116.
- Sauerland, U. and Stateva, P. (2007). Scalar vs. epistemic vagueness. In *Proceedings of SALT 17*, Ithaca, NY. CLC Publications, Cornell University.
- Schiffer, S. (2003). *The Things We Mean*. Oxford University Press.
- Shapiro, S. (2006). *Vagueness in Context*. Oxford University Press.
- Smith, N. (2008). *Vagueness and Degrees of Truth*. Oxford University Press.
- Soames, S. (1999). *Vagueness, Partiality and the Sorites Paradox*, pp. 203–27. Oxford University Press, Oxford.
- Soames, S. (2003). Higher-order vagueness for partially defined predicates. In Beall, J. (ed.), *Liars and Heaps: New Essays on Paradox*, pp. 128–50. Oxford University Press.
- Sorensen, R. (1988). *Blindspots*. Oxford University Press, Oxford.
- Sorensen, R. (2001). *Vagueness and Contradiction*. Oxford University Press, Oxford.
- Unger, P. (1975). *Ignorance*. Oxford University Press, New York.
- van Deemter, K. (1996). The sorites fallacy and the context-dependence of vague predicates. In Kanazawa, M., Piñon, C. and de Swart, H. (eds), *Quantifiers, Deduction, and Context*, pp. 59–86. CSLI.
- van Deemter, K. (2010). *Not Exactly – In Praise of Vagueness*. Oxford University Press.
- van Rooij, R. (forthcoming). Vagueness and language. In Ronzitti, G. (ed.), *The Vagueness Handbook*. Springer.

- Wheeler, S. (1972). Attributives and their modifiers. *Noûs*, 6(4): 310–34.
- Williamson, T. (1994). *Vagueness*. Routledge.
- Wright, C. (1976). Language mastery and the sorites paradox. In Evans, G. and McDowell, J. (eds), *Truth and Meaning: Essays in Semantics*, pp. 223–47. Clarendon Press.
- Wright, C. (1992). Is higher order vagueness coherent? *Analysis*, 52(3): 129–39.
- Wright, C. (1994). The epistemic conception of vagueness. *The Southern Journal of Philosophy* (Supplement), 33: 133–59.
- Wright, C. (2010). The illusion of higher-order vagueness. In Dietz, R. and Moruzzi, S. (eds), *Cuts and Clouds. Vagueness, its Nature and its Logic*, pp. 223–49. Oxford University Press.
- Zadeh, L. (1975). The concept of a linguistic variable and its applications to approximate reasoning (Parts i, ii and iii). *Information Sciences*, 8–9.

This page intentionally left blank

Part I

Measurement and Comparison

This page intentionally left blank

2

Vagueness and Scales

Scott Fults

2.1 Introduction

Scales have become so ubiquitous in linguistic analyses of gradable adjectives, that when the little sentence in (1) was introduced into the literature by Beck et al. (2004), not many took notice of its importance.¹ Beck et al. (2004) provided the truth conditions in (2) as a semantic analysis of (1):²

(1) John is tall compared to Bill.

(2) $\exists d$ [John is d -tall and $d > d_S$]

where d_S is the size standard made salient by the utterance context
 $d_S :=$ Bill's height

(2) says that John's height is greater than some standard height, and that standard height is Bill's height. Degree theories assume *precise* scales as primitives of the metalanguage used to write truth conditions.³ These scales are mappings from the domain onto the real numbers, whereby each number, called a *degree*, represents an extent of some property; in the case of *tall* entities are mapped to their heights on a tallness scale.⁴ The relation between the entity predicated of the gradable adjective (henceforth, the *referent*) and the standard of comparison (given in the *compared to* phrase of (1) or a *than* phrase of a comparative) is stated in terms of the *greater than* relation, $>$, which is defined exactly as it is over the real numbers.

The sentence in (1) is, of course, vague. Borderline cases are easy to find: if Bill is six feet tall and John is six feet and two inches, it is difficult to say whether John is *tall compared to Bill*. Maybe if he were six foot three inches? Or six foot four? Also, the predicate lacks sharp boundaries: if we

were to set up a series of individuals, each only a fraction of an inch taller than the previous member of the series, it is impossible to find an exact boundary between those that are *tall compared to Bill* and those that are not. The predicate is also susceptible to sorites reasoning. Anyone who is only 1/100th of an inch shorter than someone who is tall compared to Bill, is also tall compared to Bill; hence, someone who is only an inch tall is tall compared to Bill. Scales that map entities onto the real numbers and order them with $>$, are too precise to capture any of these properties.

The importance of (1), however, does not just lie in the failure of degree theories to account for its vagueness. The sentence allows us to control for a confound present in the cases of vagueness that have typically been the focus of attention; a confound that has, I believe, led us down the wrong path. Sentences like those in (3) are often used to point out that vague predicates are context-dependent (Kamp 1981; Kamp and Partee 1995):

- (3) a. John is tall.
 b. John is tall for a jockey.
 c. My two-year-old son built a really tall snowman yesterday.
 d. The D.U. fraternity brothers built a really tall snowman last weekend.

In (3a), we have to determine what the comparison class is. In (3b), the comparison class is given, but what is the typical height of a jockey? In (3c) and (3d), what are the typical heights of snowmen built by two-year-olds and fraternity brothers? The contextualist story is that vagueness arises from the process of determining the value of a precise standard based on the context and knowledge of how the world works (Fara 2000, Kamp 1981, Kamp and Partee 1995, Kennedy 1999, 2007, Lewis 1970, Soames 1999). It could be a matter of narrowing down the possible values of the standard degree, and for some purposes a precise amount is not needed, so it is left partially but not fully determined. Or perhaps our interests and purposes can shift quickly and without notice, changing the standard as we might otherwise be approaching it. Or maybe comparison classes are inadequate for determining the value of the standard either because we do not have enough experience or because people use vague predicates so haphazardly that we could never determine the normal use of a standard. Or, we are simply incapable of knowing exactly what a prototypical instance of a comparison class is no matter how much experience we have (Williamson 1996).

Beck et al.'s analysis in (2) is a version of the contextualist story; the *compared to* phrase simply sets the value of the standard degree. One

problem with this, though, is that with (1), the value of the standard can be precisely known. It can be given explicitly by the context, say, if Bill were exactly six feet tall, and it is still vague. It can even be provided within the sentence itself:

(4) John is tall compared to someone who measures exactly six feet in height.

(4) is vague in the same way that (1) is: it is impossible to determine the exact minimal height that John would have to be for (4) to be true. In these cases, the context has provided precise values for the standard, and yet vagueness remains.

Furthermore, comparatives are not vague:⁵

(5) John is taller than Bill.

We know how tall John has to be to be *taller than Bill*, e.g. John can be 1/100th of an inch taller, or 1/1000th of an inch, etc. The point is that the same information can be provided about the standard in both (1) and (5), and yet one is vague and the other is not. Accounts which treat vagueness as an effect of determining the standard from contextual and world knowledge will be hard pressed to find a reason why positives are different from comparatives. This is not to say that one cannot provide a semantics for positives that relies on the context in a way that comparatives do not. Rather, one would have to do so in a way that the information provided by *Bill* or *someone who measures exactly six feet in height* is not enough to determine a precise standard with positives, but is so with comparatives. And if this were not difficult enough, one would have to do so in a way that respects compositionality, for everyone's intuition is that *taller* is made up of *tall* and *-er*.⁶ So, part of the importance of (1) is that it allows us to dispense with accounts of vagueness that place the blame on using contextual and world knowledge to determine the standard. This is not to say these are not factors in determining the standard; sentences like those in (3) clearly show us they do. But once we keep those things fixed, vagueness remains.

It seems obvious, from a naive perspective anyway, that the truth conditions of constructions containing gradable adjectives involve something like scales. Let us assume that this is so; what type of scale would it be? The hypothesis I want to explore here is that positive gradable adjectives do not involve scales like the ones assumed in degree theories, while comparatives do. By reviewing the new data supplied by positives with *compared to* phrases, a number of properties will emerge

that will constrain the types of scales appropriate for encoding vague predicates like positives. Specifically, the borderline cases of vague predicates populate what I will call a *buffer zone* between the definitely false and the definitely true that can be characterized in terms of its size, which is scalable, and its shape, which encodes levels of uncertainty across borderline cases. I will present this data in section 2.2.

In section 2.3, I will review two scale types, *ratio scales* and *analog magnitude scales*. Ratio scales are the precise scales we are familiar with; they map entities onto the real numbers. Analog magnitude scales map entities onto the *approximate number system* (ANS), which is a continuous, noisy ordering of *mental magnitudes*. It is an evolutionarily ancient cognitive resource, common to humans and non-verbal animals, and it is used to represent and reason about the world. According to one line of thought, the ANS underlies our species-specific arithmetic knowledge of the integers and the real numbers, though the details of how this might be so are unknown (Gallistel and Gibbon 2000). In fact, a complete formal definition of the ANS has not yet been developed, only a partial list of its properties.

In section 2.4, I will show how the properties of analog magnitude scales correlate quite well with those properties of vague predicates outlined in section 2.2. But, because the ANS and analog magnitude scales are so poorly understood, I will not be able to provide a detailed compositional semantics for gradable adjective constructions. Instead, I want to focus on how analog magnitude scales seem to be implicated in the truth conditions of positives, while precise ratio scales are appropriate for comparatives, as discussed in section 2.5.

The idea of associating analog magnitude scales and the ANS with vagueness is appealing beyond the facts though, since it ties vagueness to non-verbal numerical cognition in humans and animals. Likewise, if it is correct that precision in natural language such as that found in comparatives is associated with ratio scales, the real numbers, and the integers, then precision can be tied to the learned, verbal, numerical cognitive capacities that are unique to humans. I hope at least to convince the reader that the ultimate explanation of vagueness lies in answering two difficult and seemingly unrelated questions: (i) how is *taller* compositionally related to *tall*? And, (ii) how are the real numbers and integers related to the approximate number system?

2.2 Properties of vagueness

Even when positive and comparative adjectives appear with phrases containing the same information about the standard, they nonetheless

behave differently with respect to the traditional diagnostics for vagueness: borderline cases, fuzzy boundaries and ease of forming sorites paradoxes. In this section, the positive form with *compared to* phrases will be used to recast the notion of fuzzy boundaries and borderline cases: it will be shown that borderline cases do not exactly coincide with the boundary, be it fuzzy or not. Then several other properties of vagueness will be revealed, properties that give us insight into the type of scales that would be appropriate in describing their meanings. Specifically, vagueness has a *shape* and it is *scalable*.

2.2.1 Buffer zones and maximum uncertainty

There is not a clear, definite line between the extension and nonextension of a vague predicate. As an often used metaphor states, there is only a continuous change such as with a blurred shadow fading into the background (Wright 1976). The idea behind the metaphor is that blurriness in the boundary symbolizes our uncertainty when making judgments about borderline cases: the less certain we are, the more blurry it is. This metaphor, though, does not quite work in the way it is usually intended. Our uncertainty is not represented in the boundary line, as can be seen by looking closely at sentences like *John is tall compared to Bill*, since these are vague but also tell us exactly where the boundary line is (or should be).

Let us first look at a situation where there is a minimal difference between the standard and the referent. Comparatives allow what Kennedy (2007) calls *crisp* judgments, while positives do not. His example is in (6):

- (6) Context: a 100-page novel and a 99-page novel
- a. #This novel is long compared to that one.
 - b. This novel is longer than that one.

While (6b) is definitely true, (6a) is false. What would it take for (6a) to become true? The judgment remains roughly the same if a page is added to the longer novel. But if we build a series of books, each one a page longer than the previous, and we check our judgments after each additional page, then we go through a gradual change of certainty – from definitely false, to uncertain, and finally to definitely true.

In order to judge a book to be *long compared to that one*, or a person to be *tall compared to Bill*, it is required that there be some amount of space between the referent and the standard on the relevant scale (Kennedy 2007). I will call this space the *buffer zone*. The buffer zone marks our uncertainty about whether members of the series belong to

the extension; it is where borderline cases exist. It is similar to the gap in supervaluationist accounts of vagueness, except that the buffer zone does not have precise boundaries. It just continues in both directions as long as there is uncertainty. The buffer zone, however, should not be mixed up with the actual boundary.

The actual boundary can be known precisely. In this case, it is the 99-page book. If this boundary were blurry and blurriness represented our uncertainty, we should be uncertain at the 99-page position in the series. In fact, the uncertainty should continue below that, since blurriness does not have a sharp cutoff point. However, there is no uncertainty in judging a 99- or 98-page book: a 99-page book is certainly not long compared to another 99-page book, nor is any shorter book. Thus, the standard marks off the *lower end* of the buffer zone, where certainty returns.

The higher end of the buffer zone is roughly where objects are definitely in the extension of the predicate, for example, where books become *long compared to that 99-page book*. Let us assume for ease of exposition that it is roughly around 115 pages. A 115-page book might not be long, but it can still be long compared to that 99-page book.

There has been relative silence about the intuition that there is a *maximum level of uncertainty*. Where uncertainty reaches a maximum is an empirical question; one that could and should be tested under experimental conditions. My own intuitions suggest that it is at a point roughly around the middle of the buffer zone. On both sides of this peak point (or points), uncertainty fades away giving rise to more certain judgments. Using the example above, maximum uncertainty would be at about 107 or so pages. Below the peak, books are closer to the nonextension of the predicate, and above it, they are closer to the extension. Thus we can say that the buffer zone has some kind of shape to it.

2.2.2 Shape

That the buffer zone has a shape is not controversial, because that is to say simply that there are borderline cases, i.e. there is not a sharp cutoff point between the extension and nonextension. What shape best characterizes the buffer zone is more interesting and has been, at least implicitly, argued about for some time. Supervaluationism (Fine 1975, Kamp 1975, Klein 1980) makes no claims about the levels of uncertainty inside the extension gap, but does claim that there are sharp cutoff points at its upper and lower limits. This, as many have noted, leads to the problem of higher-order vagueness. Intuitions suggest gradual changes of uncertainty even at the purported gap's boundaries.⁷ Given these intuitions along with the conclusion that there is a maximum level of uncertainty,

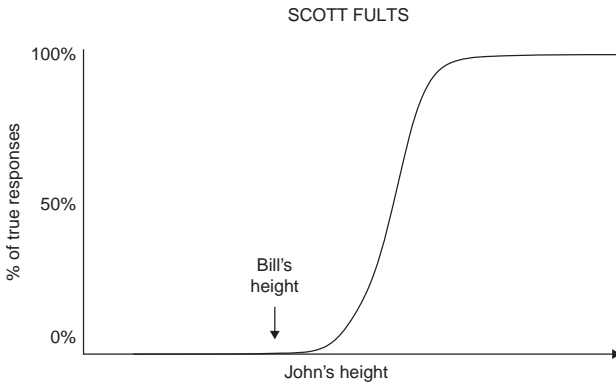


Figure 2.1 Hypothesized results of experiment

a better representation of the buffer zone would be captured by something more like what Black (1937) suggests.

Imagine an experiment in which subjects are asked to judge the truth or falsity of a sentence like *John is tall compared to Bill* within a context that explicitly states the heights of John and Bill. The value of the standard remains constant over each trial, say if Bill is exactly six feet tall, but the height of the referent varies across the buffer zone. Sometimes there would be unanimity in responses, both within and across subjects: if John was six feet tall, just like Bill, everyone would agree that the sentence was false; and if John were seven feet tall, everyone would agree that it was true. But for some trials, there would be variation: maybe only 25 or 50 percent of responses are “true.” In these cases, the assumption would be that variation, both within and across subjects, indicates uncertainty. Maximum uncertainty is reached in the situation where 50 percent of responses are “true” and 50 percent are “false.” If there are only 25 percent “true” responses, then there is half as much uncertainty.

The data collected from such an experiment (with a suitably large number of trials and participants) would give us insight into how uncertainty varies over the buffer zone. Black (1937) suggests that the results would be similar to the S-curve in Figure 2.1.⁸

The percentage of true responses is plotted as a function of John’s height. The hypothesis is that the buffer zone has a certain shape to it which can be measured by the percentage of true responses. Certainty is attained at 0 and 100 percent, and maximum uncertainty is where the line is at 50 percent “true” responses and 50 percent “false.” So, one way of stating the hypothesized shape of the buffer zone is to say it is a classic

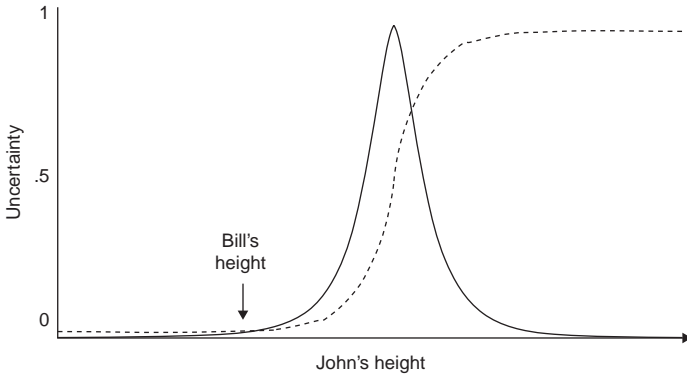


Figure 2.2 Shape of the buffer zone

S-curve. If uncertainty is plotted as a function of John's height, it has a hypothesized shape as in Figure 2.2.

These representations of the buffer zone do not fall into the trap of higher-order vagueness because there are no sharp points of change from uncertainty to certainty (except of course at the maximum point). I am going to assume that the buffer zone has a shape something like this, though the actual facts may prove that assumption wrong.

2.2.3 Scalability

The final property of the buffer zone that emerges with *compared to* constructions is that it is *scalable*: the size of the buffer zone is directly proportional to the value of the standard. If Bill is six feet tall, then for John to be tall compared to Bill, he might have to be about five or so inches taller than Bill. That difference, about five inches, does not carry over to other comparisons:

- (7) That building is tall compared to this one.
- (8) This toy soldier is tall compared to that one.

If we are comparing a 100-foot tall building with another that is only 5 inches taller, then (7) would be false. A 15- or 20-foot difference might make one building tall compared to the other, but not a 5-inch difference. If we are comparing toy soldiers, and one is 2 inches tall, the other does not have to be 5 inches taller in order to be tall compared to it.

This scalable effect indicates that the buffer zone is proportional to the size of the standard, meaning that the curves in Figures 2.1 and 2.2

stretch or compress depending on the standard. In order for the referent to be outside the buffer zone, the ratio of the standard to the referent must be below some, perhaps constant, fraction. For instance, if the ratio required were about 6:7, then the referent would have to be about 15 percent greater than the standard. Of course, it is an empirical question what the value of the ratio is (for tallness, book length, etc.), but the fact that vagueness is scalable suggests that there is such a requirement.

In summary, *compared to* phrases allow us to control for context and world knowledge. This in turn allows us to notice that: borderline cases occur not at the boundary, but some amount beyond the boundary; they have a kind of “shape” to them in that they vary in uncertainty in a nonlinear fashion; and that the span of borderline cases is proportional to the size of the standard. We therefore want a theory of vagueness to predict these facts. And, if gradable adjectives produce truth conditions that include scales, we want a scale that has these properties.

2.3 Scale types

The purpose of this section is to review the properties of analog magnitude scales and the approximate number system (ANS) that they are built on top of. These properties correlate exactly with those of vagueness outlined in section 2.2. I will start with a brief review of how precise scales are defined in order to facilitate the discussion of analog magnitude scales, but also because they appear to be relevant for stating the truth conditions of comparatives.

2.3.1 Precise scales

Measurement theory is concerned with the rules by which we use numbers to represent properties of objects and events.⁹ A scale is generally defined as a measuring function that maps entities or events in the domain onto the real numbers. Scales are then categorized by the kinds of transformations that can be performed on them without a loss of meaning.

The discussion here will be made easier if we adopt the example in Klein (1991).¹⁰ Assume that the domain includes a set of rods, D , with differing lengths. Rods can be compared with respect to their lengths using the relation \geq , such that $a \geq b$ if and only if a is longer than or equal to b . This results in an *empirical relational structure*, $\langle D, \geq \rangle$. A measuring function, φ , maps rods onto the structure, $\langle \mathbb{R}, \geq \rangle$, where \mathbb{R} is the real numbers and \geq is the greater than or equal to relation.

Ordinal scales are measuring functions that map the domain onto the real numbers in such a way that rank order remains constant:

$$(9) \quad \varphi(a) \geq \varphi(b) \quad \text{iff } a \succeq b$$

If the domain consists of four rods, $\{a, b, c, d\}$ and $\varphi(a) = 1$, $\varphi(b) = 2$, $\varphi(c) = 3$, and $\varphi(d) = 4$, then d is the longest rod and a is the shortest. All that ordinal scales encode is relative ordering, because the structure $\langle \mathbb{R}, \geq \rangle$ is only defined in terms of \geq . Thus they entail nothing about the differences between objects. In order to express differences, this structure must also include addition, and the empirical relational structure must have a correlate to addition.

Ratio scales are created from domains that are extensive, meaning any two objects with a gradable property can be *concatenated* to form a new object with the same property. Concatenation of a with b is represented as: $a \circ b$. Concatenations can be compared: $a \circ b \succeq c \circ d$. The empirical relational structure on a domain that includes concatenation can be defined as: $\langle D, \succeq, \circ \rangle$. Concatenation corresponds to addition in the real numbers: $\langle \mathbb{R}, \geq, + \rangle$.

Also assume that we can create perfect copies of anything in the domain and concatenate them: if c_1 is in the domain, so is its copy c_2 , and their concatenation is $c_1 \circ c_2$. Ratio scales map some entity c in the domain to 1; that entity is then the *unit of measurement*. Then, $c_1 \circ c_2$ is mapped to 2, $c_1 \circ c_2 \circ c_3$ is mapped to 3, etc. Let us symbolize concatenations of n copies of the same object c as nc .

The value of $\varphi(b)$ depends on the unit of measurement. If the unit is c , then b can be assigned a number m by concatenating copies of c , until the number of copies, n , is such that: $(n+1)c \succeq b \succeq nc$. Then, $n+1 \geq m \geq n$, and this has an error of as much as one unit. This process can be repeated using a different unit that is half the size of the first, producing an error that is also half the size of the first unit. Measurements can be made more and more precise by repeating the process and selecting smaller and smaller units. Taking the limit of such a process yields an entirely precise measurement.

Ratio scales start at zero, encoding both order *and* magnitude. Thus, magnitudes on ratio scales have meaning. For instance, height scales are ratio scales; if something is 100 meters tall, and something else is 200 meters tall, the second is twice as tall as the first. Also, differences between numbers represent differences in magnitude. If Al is six feet tall, Bob is five feet tall, and Carrie is four feet tall, then Al and Bob differ in height by the same amount that Bob and Carrie differ in height, i.e.

one foot, and the ratio of the differences is 1 : 1. These ratios are preserved when the units change; any ratio scale can be transformed by multiplication of a positive real number and the ratios between magnitudes are preserved. Ratio scales, then, can be used to measure heights, cost, baldness, etc., as long as a constant unit can be defined (an inch, a dollar, a hair) and the domain is extensive.

2.3.2 Analog magnitude scales

It is a bit ironic that precise scales are understood so well and analog magnitude scales so poorly. The former rely on our species-specific ability to represent precise numbers, while the latter employ a cognitive resource that does not require explicit training to use and is shared by most animals. What is known has come from studying how humans and animals measure quantities and magnitudes with their perceptual systems and then use those measurements to reason about and plan in the world. They are used to represent a wide range of physical stimuli such as brightness, loudness, intensity of smell, weight, etc., but also less sensory-specific properties like time durations and quantities of discrete items.¹¹

Analog magnitude scales order *mental magnitudes* along a continuous spectrum in much the same way that the real numbers are ordered on a dense scale: between every mental magnitude is another mental magnitude. However, mental magnitudes are not the exact, discrete items that the real numbers represent. Instead, they are something like noisy approximations.

In a classic experiment, subjects are blindfolded before they compare two weights. Subjects have difficulty discriminating similar weights, indicated by a decrease in accuracy and consistency and an increase in response time. Discrimination becomes easier as the difference between the weights increases. This is explained as an effect of how noisy the mental magnitudes are and how much their noise overlaps: greater overlap means more difficulty in discrimination.

Typically, a Gaussian distribution curve (or normal distribution) is used to represent a mental magnitude, as in Figure 2.3 with the x -axis representing magnitude values (like a mental “number line”), and the y -axis representing mental activation. The width of the distribution indicates its variance or noise, while the height encodes likelihood or certainty. In experiments where subjects estimate magnitudes of an external stimulus, their responses cluster around the mean at the center of the curve. It is crucial, however, that mental magnitudes are not precise values with an error range. The shape of the curve *is* the mental magnitude.

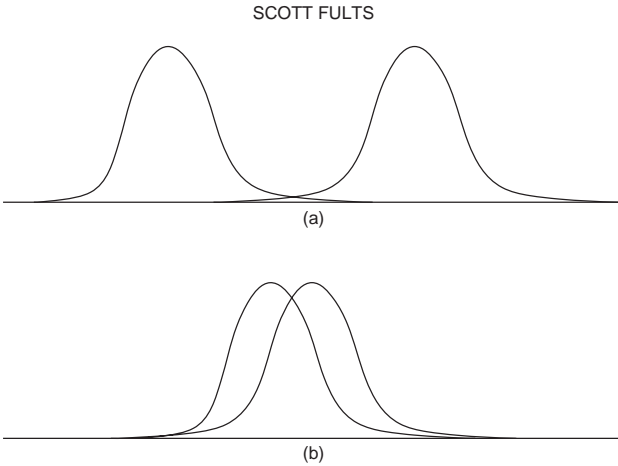


Figure 2.3 Gaussian distributions representing mental magnitudes: discrimination is determined by overlap

The change in ability to discriminate two mental magnitudes as their values get closer is not linear. If subjects in a discrimination task are asked whether the value/magnitude of some stimulus is greater than another and we plot the number of “yes” responses as a function of how close the two magnitudes are to each other, then an S-curve emerges. Maximum uncertainty of discrimination obtains when the two mental magnitudes overlap completely. This is what we would expect if uncertainty were a function of how much overlap there is between two Gaussian distribution curves.

In addition, discrimination becomes more difficult as mental magnitudes get larger, indicating that the width of their distribution increases in proportion to their value. This is called *scalar variability*, and it is stated as *Weber’s law*: the smallest amount of difference Δ in a stimulus magnitude ϕ necessary to be noticed by an observer is determined by a constant fraction, c , called the Weber fraction:

(10) Weber’s law:

$$\Delta\phi = c\phi$$

or, alternatively,

$$\frac{\Delta\phi}{\phi} = c$$




Figure 2.4 Analog magnitudes with scalar variability

The *difference threshold* is the least amount of change in the intensity of a stimulus necessary for an observer to report a difference, i.e. the amount of difference necessary to be just noticeable. Weber's law states that the size of the difference threshold $\Delta\phi$ is proportional to the size of the original magnitude ϕ . To predict the difference threshold for ϕ , simply multiply its value by the fraction c . Because c is a constant, as ϕ increases, so does $\Delta\phi$.

This means that the ability to discriminate two stimuli is determined by the ratio of their magnitudes. For example, when adults compare quantities of discrete objects (without counting verbally), the two quantities must stand in at most about a 7 : 8 ratio, i.e. the Weber fraction would be about $1/7$.¹² Anything greater than that is too close to discriminate accurately (Barth et al. 2003). So, discriminating the quantities 7 and 8 is about as easy as discriminating 70 and 80, 700 and 800, etc. Weber's law holds across modalities, though the actual value of the Weber fraction might differ greatly.

If we think of mental magnitudes in terms of the distribution curves discussed above, the Weber fraction can be thought of as encoding how much overlap between two magnitudes that can be tolerated by an observer before she loses the ability to discriminate them, while Weber's law encodes scalar variability.

A picture of an analog magnitude scale starts to emerge like that in Figure 2.4. This is a linear model presented with discrete quanta rather than continuous values.

2.3.3 Approximate number system

A growing number of researchers believe that there is a single cognitive resource, the approximate number system, used to represent estimations

of numerosity and time duration (Dehaene 1997, Feigenson et al. 2004, Gallistel and Gelman 2000, 2005). Numerosity and time duration are of a less sensory-specific nature than those studied in classical psychophysics, though the ANS may also be capable of ordering these as well. There are several reasons for this conclusion.

First, humans and animals not only generate mental magnitudes and order them, but also perform mathematical operations on them, including addition, subtraction, multiplication and division (Barth et al. 2003, Boysen and Bernston 1989, Brannon and Terrace 1998, Gallistel and Gelman 2000, 2005, Gibbon and Fairhurst 1994, Hauser et al. 1996, Olthof et al. 1998, Spelke and Dehaene 1999, Wynn 1992). While it is known that these operations are carried out, there are still questions remaining about how they work. For instance, it is unclear how the noise in the operands is propagated to the result.

Second, mental magnitudes are used to reason, make decisions and devise plans of action. For example, in Balci et al. (2009), humans and mice were put in experimental situations where, in order to act with optimal efficiency, they had to take into account the uncertainty of their mental magnitude representations of time duration (as well as perform division with them). Both the humans and the mice computed the optimal behavior rapidly and accurately, indicating that mental magnitudes are part of our mental models of the world.

Third, humans appear to have a learned bidirectional mapping from number symbols and words to analog magnitudes. This mapping preserves the relationship between mental magnitudes and the world: the mental magnitude associated with a symbol is the same mental magnitude that is generated when estimating the size of a set containing the number of objects represented by that symbol¹³ (Gallistel and Gelman 2005). In Moyer and Landauer (1967), subjects were required to judge the relative order of two Arabic numerals as rapidly as possible. Reaction times were longer when the numbers were numerically close together, and shorter when they were further apart. Parkman (1971) showed that reaction times were long when the smaller digit was large, indicating scalar variability – even when subjects were not estimating numerosities, but interpreting symbols. Similar results have been shown for verbal numbers (Cordes et al. 2001, Whalen et al. 1999).

All of this suggests that the ANS, and scales built on top of it, are integral to our mental models of the world; we generate them, we make computations over them, reason with them, and we use language to talk about them.

I will assume, then, that analog magnitude scales could in principle be given a formal description like that given for ratio scales above, but

to the best of my knowledge, no one has determined how. Recall that to define a ratio scale, one needs concatenation in the domain and addition in the number system. Without a meaningful understanding of the arithmetic operations that can be performed on the number system, the mapping from the world to that number system is meaningless.¹⁴ While it is known that humans and animals perform addition, subtraction, multiplication and division with mental magnitudes, not much is known about the properties of these operations (Gallistel and Gelman 2005). Hence, for now, it is impossible to give a precise and accurate measure-theoretic description of the mapping from the world to analog magnitude scales, though I will attempt to give a sketch of what one would look like provided the necessary pieces were given.

First though, it is worth briefly discussing proposals which use the notion of a *semi-order* to describe this mapping onto analog magnitude scales. For instance, van Rooij (this volume), based on work by Luce (1956) and Scott and Suppes (1958), describes a semi-order in this way, where X is a finite set, ϵ a positive real number, \succ_P the ordering relation based on the predicate P , and f_P a real-valued function:

$$(11) \quad \langle X, \succ_P \rangle \text{ is a semi-order iff } \exists f_P \text{ s.t.} \\ \forall x, y \in X : x \succ_P y \text{ iff } f_P(x) > f_P(y) + \epsilon$$

With a semi-order, objects x, y are ordered with respect to one another unless they are too close, where “too close” is determined by the value of ϵ . Call this a fixed-margin-of-error ordering, because ϵ is a fixed value. If \succ_P is an ordering based on tallness, then $x \succ_P y$ if and only if the height of x is greater than the height of y plus some fixed amount ϵ . If the difference between the heights of x and y is smaller than ϵ , they are indistinguishable.

At first glance, semi-orders look like an appropriate way to start building an analog magnitude scale since they capture the property of indistinguishability between items that are too similar to one another, as well as the intransitivity of indistinguishability.¹⁵ If x, y are indistinguishable in terms of \succ_P , i.e. $\neg(x \succ_P y)$ and $\neg(y \succ_P x)$, and y, z are also indistinguishable, i.e. $\neg(y \succ_P z)$ and $\neg(z \succ_P y)$, we cannot conclude that x, z are also indistinguishable. But there are several reasons why semi-orders are incapable of describing analog magnitudes and vagueness. First, ϵ is too sharp of a cutoff point: analog magnitudes and vagueness do not have a sharp points of discrimination (the problem of higher-order vagueness). Relatedly, semi-orders are incapable of providing an account of variable uncertainty, since again, two values $f(x), f(y)$ are either within ϵ to each other or not; there are no intermediate levels. And lastly, ϵ is

fixed. As we have seen, analog magnitudes and vagueness show scalar variability. If semi-orders are to be used for either, ϵ must change based on the value of the magnitudes being compared.¹⁶

All this being said, perhaps we can give a sketch of what a theory of analog magnitude scales would look like. It will necessarily be incomplete, therefore the following should be taken with a (perhaps big) grain of salt. But until analog magnitudes are better understood, it is the best we can do. Nonetheless, the following description will be helpful in describing the relation between analog magnitudes and vague predicates given in the next section.

Let us assume that there is an analog measuring function μ , which maps entities onto the approximate numbers \mathbb{A} . The ANS has its own relation, \succsim , that orders members of \mathbb{A} , though it does so only approximately. Also assume that there is an addition operation on \mathbb{A} , \oplus , that can add two approximate numbers to form another approximate number. Thus, the empirical relational structure is $(\mathbb{A}, \succsim, \oplus)$. This is incomplete because we cannot define what \oplus means.¹⁷ The problem is that analog magnitude scales and the ANS do not employ consistent units like ratio scales and the real numbers.

The relation \succsim orders the members of \mathbb{A} approximately, meaning that the noise of each mental magnitude in \mathbb{A} is taken into account. We need two things: (i) a limit of overlap (a difference threshold) where discrimination becomes impossible, and (ii) scalar variability. Recall that the Weber fraction accounts for the difference threshold of external stimuli, and Weber's law accounts for scalar variability. But Weber's law cannot be simply inserted into the definition of \succsim , since it is defined in terms of precisely measured external stimuli, i.e. in (10), $\phi \in \mathbb{R}$ (the real numbers), and \succsim is defined over the ANS.

So instead of using Weber's law, let us say that two mental magnitudes can be ordered with respect to each other if and only if they are sufficiently distant from one another, where *sufficiently distant* is determined by their ratio. The difference between two mental magnitudes necessary to discriminate them is a constant fraction k of the smaller magnitude, as in (12), where ψ is a mental magnitude:¹⁸

$$(12) \quad \Delta\psi = k\psi$$

(12) encodes both a discrimination limit defined by k and scalar variability in the same way the Weber's law does, but applies to mental magnitudes rather than the real numbers. What (12) fails to do is encode variable certainty, the key concept here, since k , like c in Weber's

law, is defined as an acceptable amount of certainty (or accuracy in discrimination). Instead, certainty should be determined by how much two mental magnitudes overlap: the more two magnitudes overlap, the more uncertain one would be in discriminating them. Let us assume that the mental domain has something like (12) to governing \succsim to help determine a threshold of certainty, whereby observance of (12) results in discrimination.

Note that while we might think of k as determining a boundary (or discrimination limit), it is not a real number, but an approximate number as well. So even that boundary is fuzzy. Hence, while we cannot define \succsim , we can say that it orders any $a, b \in \mathbb{A}$ as long as $b - a$ is at least ka .

2.4 Positive gradable adjectives and analog magnitude scales

Recall these properties of vague positive adjectives from section 2.2:

- (13) Buffer zones:
- a. are where borderline cases exist
 - b. have a shape with a maximum uncertainty
 - c. are scalable

The purpose of this section is first to show how well these properties correlate to the properties of analog magnitudes and the ANS. The second purpose is to make an attempt at describing how the properties in (13) could be captured with analog magnitude scales.¹⁹

The insight provided by using analog magnitude scales is the notion of fuzzy values. Mental magnitudes are fuzzy, and therefore, if the boundaries of vague predicates are mental magnitudes, we can explain why they seem to be fuzzy as well. But the idea of a fuzzy boundary has been around for a while, and, as discussed in section 2.2, assuming a fuzzy boundary does not result in the properties given in (13). Analog magnitude scales, however, allow us to treat both the standard and referent as fuzzy. That is, in a sentence like *John is tall compared to Bill*, both Bill's height *and* John's height can be represented by mental magnitudes. Once we make that assumption, the properties in (13) fall out easily.

A first attempt at the truth conditions for positive gradable adjectives appearing in sentences like (14) can be written as in (15), where v is a vague positive gradable adjective and a, b are elements from the domain, μ is the analog measuring function that maps entities onto \mathbb{A} :

(14) John is tall compared to Bill.

(15) $\llbracket v(a, b) \rrbracket = 1$ iff $\exists a \exists b [\mu(a) = a \ \& \ \mu(b) = b \ \& \ a \succsim b]$

With truth conditions like (15), we can explain why the buffer zone contains borderline cases. Members of \mathbb{A} are only approximate, but with an uncertainty that tapers off. If $a, b \in \mathbb{A}$ and they are sufficiently distant from each other as defined by (12), \succsim orders them with certainty. If a, b are too close, however, \succsim has an inherent uncertainty determined by how much a and b overlap. Recall that \succsim orders any two approximate numbers as long as there is a difference between them of at least k , as stated in (12).

We can also explain why the maximum level of uncertainty is above the value of the standard. Recall that it is natural to assume that the maximum level of uncertainty coincides with the boundary (i.e. the standard). The minimal difference necessary to discriminate two approximate numbers defines a threshold, stated in terms of proportionality, which in a sense forms a new boundary. If (12) defines a mental threshold of certainty, then above the threshold we are more certain of truth, and below it, we are more certain of falsity. Thus, the boundary might seem to be above the standard, but that is because the threshold is above the standard as defined by (12). Borderline cases, then, are those situations in which certainty is less than optimal because mental magnitudes mapped onto the ANS overlap, and maximum uncertainty occurs around the threshold. The shape of the buffer zone is determined by two approximate numbers approaching one another, beginning to overlap, and then reaching and passing the threshold.

Scalability is accounted for because \succsim is dependent on (12), which states that the size of the difference necessary to discriminate two mental magnitudes is proportional to the size of the standard. Note also that if k is stated precisely, higher-order vagueness could be a worry. But k only determines the threshold, which is not a precise value, because in the ANS there are no definite values of proportionality: $k \in \mathbb{A}$.

(15), however, cannot be correct, since it encodes uncertainty only indirectly through the (as yet imprecisely defined) relation \succsim . The problem also lies in stating the truth conditions with only $\{0, 1\}$. In order to encode uncertainty, it seems like truth must also be placed on a scale, perhaps even an analog magnitude scale. The details of what such a drastic change would look like are beyond the scope of this chapter, so just the informal description will have to do for now.

One benefit of using analog magnitude scales in describing the meanings of vague positive adjectives is that they make a prediction about

what shape the buffer zone should have. Many-valued logics, such as fuzzy logic and probabilistic logic, have been used (more or less successfully) to describe levels of uncertainty across borderline cases. But while these theories may be capable of *describing* whatever shape the levels of uncertainty a vague predicate may have, none make *predictions* about shape. That is, many-valued logics are capable of describing a plethora of shapes, and offer no reason why vague predicates might have the shape they do have. Analog magnitude scales make such a prediction, namely that vague predicates should have uncertainty levels that look like Gaussian curves with distribution widths that are proportional to their mean.²⁰

It should be noted that the proposal here is not that sentences with vague predicates are *verified* using analog magnitude scales, though sometimes they may be. The examples thus far have not involved estimating someone's height or a book's length; the exact heights and lengths have been given precisely by the context. The proposal is that the representation of a vague predicate's meaning, as seen through truth conditions or models of the world, actually invoke analog magnitude scales and the ANS. That is, analog magnitude scales are part of the ontology of truth conditional models. If humans and animals build models of the world with analog magnitudes, and then use those models to make plans of action and carry them out, I see no reason to believe that we are incapable of using them to build truth conditional models for natural language sentences in general, and in fact, maybe we must use them for vague predicates. Once the truth conditional model is built in such a way, using a precise measurement to verify it does not make the truth conditions more precise. The way to make a comparison precise with natural language, apparently, is to use a comparative form.

2.5 Comparative gradable adjectives and ratio scales

This section is short, since there are abundant analyses of comparatives that use precise scales.²¹ I will just briefly describe how typical truth conditions of comparatives are stated.

Recall that precise scales assume an empirical extensive structure $\langle D, \succeq, \circ \rangle$ with a structure on the real numbers $\langle \mathbb{R}, \geq, + \rangle$. The relations \succeq and \geq correspond to one another, as do concatenation and addition:

- (16) a. $a \succeq b$ iff $\varphi(a) \geq \varphi(b)$
 b. $\varphi(a \circ b) = \varphi(a) + \varphi(b)$

Thus, comparative constructions such as (17) can be given truth conditional values as in (18), where v is a comparative adjective, and φ is a measuring function that maps its argument onto \mathbb{R} :

(17) John is taller than Bill.

(18) $\llbracket v(x, y) \rrbracket = 1$ iff $\exists d \exists d' [\varphi(x) = d \ \& \ \varphi(y) = d' \ \& \ d > d']$

Degree analyses like this one assume that the real numbers, measure functions like φ , and the relation $>$ are part of the semantic ontology. (18) simply says that (if we are talking about tallness) in order for x to be *taller than* y , x 's height must be greater than y 's height, where $>$ is defined over the real numbers as it typically is in mathematics, that is, *precisely*. Thus, degree analyses predict that there should not be any borderline cases, fuzzy boundaries, scalability, etc. This differs from (15) in several respects. First, in (15), objects are mapped by a measure function not onto the real numbers, but rather onto the approximate numbers. Second, the approximate numbers are ordered with respect to \sim .²²

There is a fair amount of work done by the measuring function φ in (18). Most importantly, it encodes the choosing of a sufficiently small unit of measurement. This is where what Pinkal (1995) calls *precisification* and Sauerland and Stateva (2007, and this volume) call *scalar vagueness* is found. It concerns the granularity of our measurements, and is therefore subjective: it *shifts* along with our interests and purposes at any given moment, to use the terminology in Fara (2000). What counts as *two meters* depends on what are purposes are, how precise our measurements need to be, or what measurement techniques we have available, etc. If we are just using our naked eye, what counts as *two meters* will not be very precise. (19) is vague in this way, since we might not care to make very precise measurements about what counts as three inches:

(19) John is three inches taller than Bill.

Likewise, comparatives without a measure phrase can display this granularity effect, depending on how fine or coarse we care to make our measurements:

(20) John is taller than Bill.

(20) will be true or false depending on whether we can see a difference between John and Bill, but also depending on what our purpose for making the comparison is. If to be taller than Bill means he gets a different

size suit coat, (20) could be true or false depending on how much taller than Bill John is. However, it must be admitted that these examples do not seem like typical cases of vagueness. When Russell (1923) famously claimed that most of natural language was hopelessly vague, he used measure phrases as examples, but added that they were vague "... to a lesser degree" than words like *tall* and *red*. The difficulty with accepting that this is vagueness lies in the fact that in principle we can always be more precise, whereas with vague predicates, precision is impossible. This may be an argument in favor of using precise scales for comparatives, since they provide a natural way to precisify by choosing smaller and smaller units of measurement.

2.6 Conclusion

The meanings of vague positive adjectives appear to involve analog magnitude scales which map entities onto the ANS. This explains why positive adjectives lack precise boundaries, have borderline cases, demand a buffer zone with a unique shape between items being compared, and are scalable. Comparatives lack all of these properties, and therefore appear to involve precise scales which map entities onto an exact number system such as the real numbers.

It is still unclear whether this provides an easy answer to the sorites paradox. The paradox is centered around a precise scale, i.e. a series of individuals whose differences are precisely known. Once a relation is established between two individuals on the scale using a vague predicate like *tall compared to Bill*, we inductively apply the predicate to the rest of the scale. We seem to readily accept that the vague predicate is hereditary on the series, even though doing so results in a false conclusion. Why does it seem that *tall compared to Bill* should be hereditary on a precise scale? Perhaps it is because we are using a predicate that does not belong on a precise scale, and reasoning about it as such results in errors. Another possibility is that mental magnitudes have distributions that are in some sense infinite; their noise trails off towards zero, but never actually gets there. Whatever the answer to the sorites paradox, this chapter tries to provide a framework for finding it.

The answer lies in part in understanding how analog magnitude scales and precise scales are related to one another, if indeed they are. This is a difficult question to answer since analog magnitude scales and the ANS are not fully understood. According to one line of thought, the ANS provides the cognitive foundation for our knowledge of the integers and reals, and indeed all of our species-specific arithmetic abilities (Gallistel

and Gibbon 2000). But this too is only a vague hypothesis. How the ANS might have given rise to precise concepts of numbers is a mystery. Still, I do not believe we have to wait for an answer to this question, in order to arrive at an answer to ours.

There is one final point to be made about sentences like *John is tall compared to Bill* which makes them worth considering: we now have a sentence which forms a minimal pair with its comparative counterpart, *John is taller than Bill*. There is a meaning difference that seems to be triggered by the morphosyntactic difference. This is only to make an obvious point which others have made: *taller* is made up of *tall* and *-er*. One wants to know how a vague gradable adjective becomes precise simply by adding *-er*. If it is true that analog magnitude scales are somehow implicated in the meanings of vague adjectives and precise scales in comparatives, then understanding vagueness might just be a matter of understanding how the meaning of *-er* and its morphosyntactic structure compositionally turn analog magnitude scales into precise ones.

Notes

1. But see Fults (2006), Kennedy (2007, this volume) and van Rooij (this volume).
2. As pointed out by Beck et al. (2004), this sentence has many variations. For example, *compared to* can be replaced with *in comparison to*, *with respect to*, *with regards to*, *relative to*, etc. These phrases are also quite mobile:
 - (i) a. Compared to Bill, John is tall.
 - b. John, in comparison to Bill, is tall.
 - c. John is, with respect to Bill, tall.

However, their movement is restricted in that they obey island constraints, indicating that they are base generated in or around the adjective phrase and do not act like context setters as Beck et al. (2004) propose. See Fults (2006), especially Chapter 2, for details and further arguments that these sentences are not hidden conditionals. I will use *compared to* phrases as in (1) as examples throughout the chapter, but all conclusions apply to these variations as well.

3. By “degree theories” I mean work begun in Cresswell (1977), and expanded by Bierwisch (1989), Heim (1985, 2000), Kennedy (1999, 2007), Kennedy and McNally (2005), Schwarzschild and Wilkinson (2002) and Stechow (1984), among many others. I do not mean to refer to theories that invoke degrees of truth to describe vagueness, such as in Edgington (1996), Machina (1976) and Tye (1994), although there are similarities.
4. Some theories map entities onto an exact degree (Kennedy 1999), some onto sets of degrees (Stechow 1984), and some onto intervals (Schwarzschild and Wilkinson 2002), but these details are irrelevant here.
5. At least not in the same way that positives are. They may possess what Sauerland and Staleva (2007, and this volume) call *scalar vagueness*, though this

may be what Kennedy (2007) (following Pinkal 1995) calls imprecision. See section 2.5 for a short discussion.

6. See Kennedy (2007) for an attempt, but one in which the meaning of *tall* is not included in the meaning of *taller*.
7. Even fuzzy set theory is subject to this criticism if membership in a fuzzy set is dependent on precise positions on a scale (Keefe and Smith 1996).
8. I have modified Black's graph a bit, but this is basically what he predicts. (He does not actually run the experiment either.)
9. See Gescheider (1997: 186–91) for a detailed, but informal description of measurement theory, and Krantz et al. (1971) and Luce and Narens (1987) for more mathematically detailed descriptions.
10. Klein (1991) is following Krantz et al. (1971). I will leave out certain details, so the reader is directed to both for more detail, especially Klein (1991) who presents his discussion in the context of the natural language semantics of comparatives.
11. See Feigenson et al. (2004), Gallistel and Colman (2005), Gelman and Gallistel (2004) and Gescheider (1997) for detailed reviews.
12. Success in discriminating two stimuli is usually defined as about 85 percent accuracy (Gescheider, 1997).
13. Of course, humans may also have a mapping from number symbols and words onto whatever concepts are involved in precise arithmetic.
14. See Gallistel and Gelman (2005) for some discussion of this, but also see Krantz et al. (1971) and Luce (1990).
15. The purpose of introducing semi-orders in Luce (1956) was, in fact, to capture these properties.
16. My suspicion is that the problem of higher-order vagueness will creep up whenever real numbers are used to define borderlines. (See Pagin, this volume, for another similar attempt.)
17. Furthermore, it may not be a good idea to assume that concatenation \circ is the analog to \oplus .
18. This may just be a restatement of Ekman's law, which states that the change in mental magnitude resulting from a just noticeable difference in a stimulus is a constant fraction of that mental magnitude. See Gescheider (1997: 351–6).
19. Again I must admit that without a precise definition of analog magnitude scales, this goal will remain elusive.
20. Égré and Bonnay (to appear) provide a logical model of vagueness and certainty based on signal detection theory, a probabilistic account of subject responses in typical psychophysics tasks. Their account captures variable certainty without running into the problem of higher-order vagueness, though there may be problems with scalability and deriving the precise meanings of comparatives from the meanings of positives. The reader is directed there for an analysis that is much in line with the spirit of the current chapter.
21. Klein (1991) and Kennedy (1999) give particularly nice reviews, and both contain many references.
22. It is worth noting here that introducing *both* approximate numbers and real numbers into the semantic ontology is rather cumbersome, and I think this speaks against either analysis. An alternative account of comparative constructions is presented in Pietroski (2007) which does not refer to real numbers.

References

- Balci, F., Freestone, D. and Gallistel, C. R. (2009). Risk assessment in man and mouse. In *Proceedings of the National Academy of Sciences*, 106: 2459–63.
- Barth, H., Kanwisher, N. and Spelke, E. (2003). The construction of large number representations in adults. *Cognition*, 86: 201–21.
- Beck, S., Oda, T. and Sugisaki, K. (2004). Parametric variation in the semantics of comparison: Japanese vs. English. *Journal of East Asian Linguistics*, 13: 289–343.
- Bierwisch, M. (1989). The semantics of gradation. In Bierwisch, M. and Lang, E. (eds), *Dimensional Adjectives: Grammatical Structure and Conceptual Interpretation*. Springer-Verlag, Berlin.
- Black, M. (1937). Vagueness: an exercise in logical analysis. *Philosophy of Science*, 4: 427–55.
- Boysen, S. T. and Berntson, G. G. (1989). Numerical competence in a chimpanzee. *Journal of Comparative Psychology*, 103: 23–31.
- Brannon, E. M. and Terrace, H. S. (1998). Ordering of the numerosities 1 to 9 by monkeys. *Science*, 282: 746–9.
- Cordes, S., Gelman, R., Gallistel, C. R. and Whalen, J. (2001). Variability signatures distinguish verbal from nonverbal counting for both large and small numbers. *Psychonomic Bulletin & Review*, 8(4): 698–707.
- Cresswell, M. J. (1977). The semantics of degree. In Partee, B. (ed.), *Montague Grammar*, pp. 261–92. Academic Press.
- Dehaene, S. (1997). *The Number Sense: How the Mind Creates Mathematics*. Oxford University Press, New York.
- Edgington, D. (1996). Vagueness by degrees. In Keefe, R. and Smith, P. (eds), *Vagueness: a Reader*. MIT Press, Cambridge, Mass.
- Égré, P. and Bonnay, D. (2010). Vagueness, uncertainty and degrees of clarity. *Synthese*, 174(1): 47–78.
- Fara, D. G. (2000). Shifting sands: an interest-relative theory of vagueness. *Philosophical Topics*, 28: 45–81.
- Feigenson, L., Dehaene, S. and Spelke, E. S. (2004). Core systems of number. *Trends in Cognitive Sciences*, 8: 307–14.
- Fine, K. (1975). Vagueness, truth and logic. *Synthese*, 30: 265–300.
- Fulst, S. (2006). The structure of comparison: an investigation of gradable adjectives. PhD thesis, University of Maryland, College Park, Md.
- Gallistel, C. and Gelman, R. (2000). Non-verbal numerical cognition: from reals to integers. *Trends in Cognitive Sciences*, 4(2): 59–65.
- Gallistel, C. R. and Gelman, R. (2005). Mathematical cognition. In Holyoak, K. and Morrison, R. (eds), *The Cambridge Handbook of Thinking and Reasoning*, pp. 559–88. Cambridge University Press, New York.
- Gallistel, C. R. and Gibbon, J. (2000). Time, rate, and conditioning. *Psychological Review*, 107(2): 289–344.
- Gelman, R. and Gallistel, C. R. (2004). Language and the origin of numerical concepts. *Science*, 306(5695): 441–3.
- Gescheider, G. A. (1997). *Psychophysics: the Fundamentals*. Lawrence Erlbaum Associates, Inc., Mahwah, NJ.
- Gibbon, J. and Fairhurst, S. (1994). Ratio versus difference comparators in choice. *Journal of Experimental Analysis of Behavior*, 62: 409–34.

- Hauser, M. D., MacNeilage, P. and Ware, M. (1996). Numerical representations in primates. *Proceedings of the National Academy of Sciences of the United States of America*, 93: 1514–17.
- Heim, I. (1985). Notes on comparatives and related matters. University of Texas, Austin.
- Heim, I. (2000). Degree operators and scope. In Jackson, B. and Matthews, T. (eds), *Semantics and Linguistic Theory 10*. CLC Publications, Ithaca, NY.
- Kamp, H. (1975). Two theories of adjectives. In Keenan, E. (ed.), *Formal Semantics of Natural Language*, pp. 123–55. Cambridge University Press, New York.
- Kamp, H. (1981). The paradox of the heap. In Mönnich, U. (ed.), *Aspects of Philosophical Logic*, pp. 225–77. D. Reidel Publishing Company, Dordrecht.
- Kamp, H. and Partee, B. (1995). Prototype theory and compositionality. *Cognition*, 57: 129–91.
- Keefe, R. and Smith, P. (1996). Introduction: theories of vagueness. In Keefe, R. and Smith, P. (eds), *Vagueness: a Reader*. MIT Press, Cambridge, Mass.
- Kennedy, C. (1999). *Projecting the Adjective: the Syntax and Semantics of Gradability and Comparison*. Garland, New York.
- Kennedy, C. (2007). Vagueness and grammar: the semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy*, 30(1): 1–45.
- Kennedy, C. and McNally, L. (2005). Scale structure and the semantic typology of gradable predicates. *Language*, 81(2): 345–81.
- Klein, E. (1980). A semantics for positive and comparative adjectives. *Linguistics and Philosophy*, 4: 1–45.
- Klein, E. (1991). Comparatives. In Stechow, A. V. and Wunderlich, D. (eds), *Semantics: an International Handbook of Contemporary Research*, pp. 673–91. Walter de Gruyter, New York.
- Krantz, D. H., Luce, R. D., Suppes, P. and Tversky, A. (1971). *Foundations of Measurement*, Vol. 1. Academic Press, New York.
- Lewis, D. (1970). General semantics. *Synthese*, 22: 18–67.
- Luce, R. D. (1956). Semiorders and a theory of utility discrimination. *Econometrica*, 24: 178–91.
- Luce, R. D. (1990). What is a ratio in ratio scaling. In Bolanowski, S. J. and Gescheider, G. A. (eds), *Ratio Scaling of Psychological Magnitude: In Honor of the Memory of S. S. Stevens*. Lawrence Erlbaum Associates, Inc., Hillsdale, NJ.
- Luce, R. D. and Narens, L. (1987). Theory of measurement. In Eatwell, J., Milgate, M. and Newman, P. (eds), *The New Palgrave: a Dictionary of Economic Theory and Doctrine*, pp. 428–32. The Macmillan Press Ltd, New York.
- Machina, K. F. (1976). Truth, belief and vagueness. *Journal of Philosophical Logic*, 5: 47–78.
- Moyer, R. S. and Landauer, T. K. (1967). Time required for judgments of numerical inequality. *Nature*, 215: 1519–20.
- Olthof, A., Iden, C. M. and Roberts, W. A. (1998). Judgments of ordinality and summation of number symbols by squirrel monkeys. *Journal of Experimental Psychology: Animal Behavior Processes*, 23(3): 325–39.
- Parkman, J. M. (1971). Temporal aspects of digit and letter inequality judgments. *Journal of Mathematical Psychology*, 91: 191–205.
- Pietroski, P. (2007). Induction and comparison. In Conroy, A., Jing, C., Nakao, C. and Takahashi, E. (eds), pp. 154–88. University of Maryland Working Papers in Linguistics 15, UMWPiL, College Park, Md.

- Pinkal, M. (1995). *Logic and Lexicon*. Kluwer, Dordrecht.
- Russell, B. (1923). Vagueness. *Australasian Journal of Philosophy and Psychology*, 1: 84–92.
- Sauerland, U. and Stateva, P. (2007, and this volume). Scalar vs. epistemic vagueness. In Gibson, M. and Friedman, T. (eds), *Proceedings of SALT 17*. CLC Publications, Ithaca, NY.
- Schwarzschild, R. and Wilkinson, K. (2002). Quantifiers in comparatives: a semantics of degree based on intervals. *Natural Language Semantics*, 10: 1–41.
- Scott, D. and Suppes, P. (1958). Foundational aspects of theories of measurement. *Journal of Symbolic Logic*, 23: 113–28.
- Soames, S. (1999). *Understanding Truth*. Oxford University Press, New York.
- Spelke, E. S. and Dehaene, S. (1999). Biological foundations of numerical thinking. *Trends in Cognitive Sciences*, 3: 365–6.
- Stechow, A. V. (1984). Comparing semantic theories of comparison. *Journal of Semantics*, 3: 1–77.
- Tye, M. (1994). Sorites paradoxes and the semantics of vagueness. In Tomberlin, J. E. (ed.), pp. 189–206. *Philosophical Perspectives 8: Logic and Language*. Ridgeview Publishing Co., Atascadero, Calif.
- Whalen, J., Gelman, R. and Gallistel, C. R. (1999). Nonverbal counting in humans: the psychophysics of number representation. *Psychological Science*, 10: 130–7.
- Williamson, T. (1996). Vagueness and ignorance. In Keefe, R. and Smith, P. (eds), *Vagueness: a Reader*. MIT Press, Cambridge, Mass.
- Wright, C. (1976). Language mastery and the sorites paradox. In Evans, G. and McDowell, J. (eds), pp. 223–47. *Truth and Meaning*. Oxford University Press, Oxford.
- Wynn, K. (1992). Addition and subtraction by human infants. *Nature*, 358: 749–50.


3

Implicit versus Explicit Comparatives

Robert van Rooij*
ILLC, University of Amsterdam

3.1 Introduction

Consider the following figure, picturing the lengths of John and Mary:



Chris Kennedy (this volume) observed that according to most people's intuition, this picture allows us to say (1a). At the same time, most people would say that (1b) is false.

- (1) a. John is taller than Mary.
b. Compared to Mary, John is tall, but compared to John, Mary is not tall.

This observation is stated as a challenge to followers of Klein (1980). According to Klein (1980), (1a) is true if and only if *there is* a comparison class according to which John is tall and Mary is not, i.e. (Klein). It is standardly assumed (e.g. von Stechow 1984) that Klein's actual analysis of (1a) is *equivalent* with the simpler analysis of van Benthem (Benthem), which will be discussed in section 3.2:

- (Klein) (1a) is true iff *there is* a comparison class such that John is tall here and Mary is not.
(Benthem) (1a) is true iff John is tall in comparison class {John, Mary}, but Mary is not.

But according to (Bentham), (1a) and (1b) have the same truth conditions, and it is impossible to account for the contrast between them in truth conditional terms. We will meet Kennedy's challenge by making a distinction between analyses (Klein) and (Bentham).

3.2 Comparatives and comparison classes

Although expressions of many lexical categories are vague, most research on vagueness concentrates on adjectives like "tall" and "bald." In linguistics these adjectives are known as *gradable* adjectives and should be distinguished from *nongradable* adjectives like "pregnant" and "even." The latter adjectives do not give rise to (much) vagueness. There exist two major types of approaches to the analysis of gradable adjectives: degree-based approaches and delineation approaches. Degree-based approaches (e.g. Seuren 1973, von Stechow 1984, Kennedy 1999) analyze gradable adjectives as relations between individuals and degrees, where these degrees are thought of as scales associated with the dimension referred to by the adjective. Individuals can possess a property to a certain measurable degree. The truth conditions of sentences involving these adjectives are stated in terms of degrees. According to the most straightforward degree-based approach, a sentence like *John is tall* is true iff the degree to which John is tall is (significantly) greater than a (contextually given) standard of height. The comparative *John is taller than Mary* is true iff the (maximal) degree to which John is tall is greater than the (maximal) degree to which Mary is tall.¹

Delineation approaches (Lewis 1970, Kamp 1975, Klein 1980, 1991) analyze gradable adjectives like "tall" as simple predicates, but assume that the extension of these terms is crucially context-dependent. If one accounts for comparatives in terms of supervaluation structures as Lewis and Kamp do, the most obvious way to account for this context-dependence is very similar to the one used in degree-based approaches: context just determines the cutoff point. But, of course, context-dependence is more fine-grained than that. For Jumbo to be *a small elephant*, for instance, means that Jumbo is small for an elephant, but that does not mean that Jumbo is small. For instance, Jumbo will be much bigger than any object that counts as *a big mouse*. One way to make this explicit is to assume with Klein (1980) that every adjective should be interpreted with respect to a *comparison class*, i.e. a set of individuals. The truth of a sentence like *John is tall* depends on the contextually given comparison class: it is true in context (or comparison class) *c* iff John is counted as tall in this class. Klein (1980) further proposes that

the meaning of the comparative *John is taller than Mary* is context-independent and the sentence is true iff there is a context (henceforth we will use “context” instead of the more cumbersome “comparison class”) according to which John counts as tall, while Mary does not. If there is any context in which this is the case, it will also be the case in the context containing only John and Mary.

Klein (1980) favors the delineation approach towards comparatives for a number of reasons. First, a degree-based approach only makes sense in case the comparative gives rise to a *total* ordering. But for at least some cases (e.g. *more clever than*) this does not seem to be true, because *clever* is a multidimensional adjective. Second, the delineation account assumes that the meaning of the comparative “taller than” is a function of the meaning of “tall.” This is much in line with Frege’s principle of compositionality,² and also accounts for the fact that in a wide variety of languages the positive is formally unmarked in relation to the comparative.

An analysis of comparatives in terms of comparison classes is sometimes stated as if it *presupposes* that the domain of all individuals of any gradable adjective has an inherent ordering imposed upon it, and that the ordering on a comparison class must preserve the initial ordering on the domain of the adjective in order to avoid undesirable entailments. Kennedy (1997) states this in terms of the following *consistency postulate*:

For any context in which *a is φ* is true and $b \geq a$ with respect to the original ordering on the domain of φ , then *b is φ* is also true, and for any context in which *a is φ* is false, and $a \geq b$, then *b is φ* is also false.

But if this were so, the delineation analysis of comparatives would be reduced to an initial comparison ordering, and thus the delineation approach would not take the positive use of the adjective as basic. Fortunately, van Benthem (1982) has shown that what Kennedy calls the initial ordering can be *derived* from how we positively use the adjective in certain contexts, plus some additional constraints on how the meaning of the adjective can change from context to context. This is done in terms of the notion of a *context structure*, M , being a triple $\langle X, C, V \rangle$, where X is a nonempty set of individuals, the set of contexts, C , consists of all finite subsets of X , and the valuation V assigns to each context $c \in C$ and each property T those individuals in c which are to count as “being T in c .”

This definition leaves room for the most diverse behavior of individuals across contexts. Based on intuition (for instance by visualizing sticks of various lengths), however, the following plausible cross-contextual

principles make sense, which constrain the possible variation. Take two individuals x and y in context c such that $M, c \models T(x) \wedge \neg T(y)$. We now constrain the set of contexts C by the following three principles: No Reversal (NR), which forbids x and y to change roles in other contexts:

$$(NR) \quad \neg \exists c' \in C : M, c' \models T(y) \wedge \neg T(x)$$

This constraint does not prevent x and y both from being tall in larger contexts than c . However, once we look at such larger contexts, the Upward Difference (UD) constraint demands that there should be at least one difference pair:

$$(UD) \quad \forall c' \in C [c \subseteq c' \rightarrow \exists z_1, z_2 : M, c' \models T(z_1) \wedge \neg T(z_2)]$$

The final Downward Difference (DD) principle constrains in a very similar way what is allowed if we look at subsets of c : if x and y are elements of this subset, there should still be a difference pair:

$$(DD) \quad \forall c' \in C [(c' \subseteq c \ \& \ x, y \in c') \rightarrow \exists z_1, z_2 : M, c' \models T(z_1) \wedge \neg T(z_2)]$$

If we say that “John is taller than Mary” is true if and only if there is a context c such that $M, c \models T(j) \wedge \neg T(m)$, van Benthem shows that the comparative (given the above constraints on context structures) as defined above has exactly those properties which we intuitively want for most comparatives (see below). Thus we have seen that on the basis of the initial idea of the delineation approach we can *derive* the ordering relation that Kennedy (1997) claims delineation approaches must already take for granted to begin with.

In the definition of a context structure we used above, context structures give rise to orderings for any context-dependent adjective. For convenience, we will just limit ourselves to one adjective: P . If we do so, we can think of a context structure as a triple $\langle X, C, P \rangle$, where P can be thought of as a *choice function*, rather than a general valuation function.

Definition 1

A *context structure* M is a triple $\langle X, C, P \rangle$, where X is a nonempty set of individuals, the set of contexts, C , consists of all finite subsets of X , and the choice function P assigns to each context $c \in C$ one of its subsets.

Notice that $P(c)$ (with respect to context structure M) corresponds to the set $\{x \in X : M, c \models P(x)\}$ in our earlier formulation. To state the cross-contextual constraints somewhat more compactly than we did

above, we define the notion of a difference pair: $\langle x, y \rangle \in D_P(c)$ iff_{def} $x \in P(c)$ and $y \in (c - P(c))$. Now we can define the constraints as follows (where c^2 abbreviates $c \times c$, and $D_P^{-1}(c) =_{\text{def}} \{y, x\} : \langle x, y \rangle \in D_P(c)\}$):

- (NR) $\forall c, c' \in C : D_P(c) \cap D_P^{-1}(c') = \emptyset$
- (UD) $c \subseteq c'$ and $D_P(c) \neq \emptyset$, then $D_P(c') \neq \emptyset$
- (DD) $c \subseteq c'$ and $D_P(c') \cap c^2 \neq \emptyset$, then $D_P(c) \neq \emptyset$

If we say that $x >_P y$, iff_{def} $x \in P(\{x, y\})$ and $y \notin P(\{x, y\})$, van Benthem (1982) shows that the ordering as defined above gives rise to a *strict weak order*. A structure $\langle X, R \rangle$, with R a binary relation on X , is a strict weak order just in case R is irreflexive (IR), transitive (TR), and almost connected (AC).

Definition 2

A (*strict*) *weak order* is a structure $\langle X, R \rangle$, with R a binary relation on X that satisfies the following conditions:

- (IR) $\forall x : \neg R(x, x)$
- (TR) $\forall x, y, z : (R(x, y) \wedge R(y, z)) \rightarrow R(x, z)$
- (AC) $\forall x, y, z : R(x, y) \rightarrow (R(x, z) \vee R(z, y))$


The constraint that R should be almost connected is in some circles better known under its contrapositive guise as co-transitivity:

$$\forall x, y, z : (\neg R(x, z) \wedge \neg R(z, y)) \rightarrow \neg R(x, y)$$

If we now define the indifference relation, “ I ,” or in our case “ \approx_P ,” as follows: $x \approx_P y$ iff_{def} neither $x >_P y$ nor $y >_P x$, it is clear that “ \approx_P ” is an equivalence relation. It is well-known from measurement theory (e.g. Krantz et al. 1971) that in case “ $>_P$ ” gives rise to a (strict) weak order, it can be represented numerically by a real valued function f_P such that for all $x, y \in X$: $x >_P y$ iff $f_P(x) > f_P(y)$, and $x \approx_P y$ iff $f_P(x) = f_P(y)$.

3.3 Explicit versus implicit comparison

Consider again the following figure, picturing the lengths of John and Mary:



How can we account for the fact that the *explicit* comparative (1a) is intuitively true, while the *implicit* comparative (1b) is false?³

(1a) John is taller than Mary.

(1b) Compared to Mary, John is tall, but compared to John, Mary is not tall.

It is clear that according to (Benthem) (1a) and (1b) have the same truth conditions, and it is impossible to account for the contrast between them in truth conditional terms.

(Benthem) (1a) is true iff John is tall in comparison class {John, Mary}, but Mary is not.

Of course, Klein's (1980) original analysis (Klein) was a bit different, so it seems possible to account for the difference between (1a) and (1b) in terms of the difference between (Klein) and (Benthem).

(Klein) (1a) is true iff *there is* a comparison class s.t. John is tall here and Mary is not.

It can be easily shown, however, that *in case* Klein had adopted van Benthem's (1982) cross-contextual constraints, (Klein) is equivalent to (Benthem). It is immediately clear that by van Benthem's definition of a context structure and by adopting his constraints, analysis (Klein) *follows from* analysis (Benthem). But it is important to see why the reverse also holds. So suppose that (1a) is true according to (Klein). That means that there exists a comparison class $c \in C$ containing John and Mary such that $j \in P(c) \wedge m \notin P(c)$. But this means that $\langle j, m \rangle$ is a difference pair in c , and by (DD) it follows that $\langle j, m \rangle$ will be a difference pair in all $c' \in C$ that are subsets of c containing both John and Mary. By the assumption that C contains all finite subsets of X , it follows that $\langle j, m \rangle$ is also a difference pair for comparison class {John, Mary}, which means that (1a) is also predicted to be true by (Benthem).

I would like to point out that the derivation of the truth of the comparative according to (Benthem) from its truth according to (Klein) is based on (at least) *two assumptions*. The first is that we do not really make a distinction between it not being the case that Mary is tall, and Mary's being *not-tall*, or, perhaps equivalently, Mary's being *short*. A second crucial presupposition on which the above derivation is based is the assumption that C contains *all* finite subsets of X , in particular that {John, Mary} is

an element of C . The *first assumption*, however, was explicitly rejected by Klein (1980). Klein (1980) explicitly proposed that an adjective gives rise to a *three-way* partition of the comparison class c : some individuals in c are (definitely) tall, some are (definitely) not-tall, and some are neither. Klein used a three-valued logic, but the same intuition can be captured by inducing the three-way partition by a set of contrary predicates: e.g. the adjective “tall” and its antonym “short.” Although no individual in c is tall *and* short, it is possible that some individuals are neither. Klein (1980) assumed that an adjective gives rise to a three-way partition to account for *vagueness*. I will argue in the next section that doing so is indeed natural to account for the sorites paradox. In the next section I will also argue that *for the very same reason* it is also natural to give up van Benthem’s (1981) assumption that *all* finite subsets of X are appropriate comparison classes. In section 3.5 I will then show that if one gives up either of these assumptions one can generate an ordering relation that properly represents vagueness, and, or so I will argue, in terms of which one can give a natural account of the distinction between (1a) and (1b).

3.4 The sorites and semi-orders

3.4.1 Vagueness and semi-orders

Consider a long series of people ordered in terms of their height. Of each of them you are asked whether they are tall or not. We assume that the variance between two subsequent persons is always *indistinguishable*. Now, if you decide that the first individual presented to you, the tallest, is tall, it seems only reasonable to judge the second individual to be tall as well, since you cannot distinguish their heights. But, then, by the same token, the third person must be tall as well, and so on indefinitely. In particular, this makes also the last person tall, which is a counterintuitive conclusion, given that it is in contradiction with our intuition that this last, and shortest individual, is short, and thus not tall.

This so-called sorites reasoning is elementary, based only on our intuition that the first individual is tall, the last short, and the following inductive premise, which seems unobjectionable:

[P] If you call one individual tall, and this individual is not visibly taller than another individual, you have to call the other one tall too.

Our above sorites reasoning involved the predicate “tall,” but that was obviously not essential. Take any predicate P that gives rise to a complete ordering “as P than.” Let us assume that “ \sim_p ” is the indistinguishability,

or indifference, relation between individuals with respect to predicate P . Now we can state the inductive premise somewhat more formally as follows:

$$[P] \text{ For any } x, y \in X : (P(x) \wedge x \sim_P y) \rightarrow P(y)$$

If we assume that it is possible that

$$\exists x_1, \dots, x_n : x_1 \sim_P x_2 \wedge \dots \wedge x_{n-1} \sim_P x_n,$$

but $P(x_1)$ and $\neg P(x_n)$, the paradox will arise. Recall that if $P(x_1)$ and $\neg P(x_n)$, it is required that x_1 must be *visibly* or *significantly* P -er than x_n , denoted by $x_1 \succ_P x_n$. In section 3.2 of this chapter we have defined a relation “ \succ_P ” in terms of the behavior of predicate P . The constraints discussed there, however, did not allow for the possibility that

$$\exists x_1, \dots, x_n : x_1 \sim_P x_2 \wedge \dots \wedge x_{n-1} \sim_P x_n,$$

but $P(x_1)$ and $\neg P(x_n)$, and the defined comparison relation could not really be interpreted as meaning “being *visibly/significantly* P -er than.” Fortunately, there is a well-known ordering that should be interpreted this way: what Luce (1956) calls a *semi-order*. Following Scott and Suppes’ (1958) (equivalent, but still) simpler definition, a structure $\langle X, R \rangle$, with R a binary relation on X , is a semi-order just in case R is irreflexive (IR), satisfies the interval-order (IO) condition, and is semi-transitive (STr).⁴

Definition 3

A *semi-order* is a structure $\langle X, R \rangle$, with R a binary relation on X that satisfies the following conditions:

$$(IR) \quad \forall x : \neg R(x, x)$$

$$(IO) \quad \forall x, y, v, w : (R(x, y) \wedge R(v, w)) \rightarrow (R(x, w) \vee R(v, y))$$

$$(STr) \quad \forall x, y, z, v : (R(x, y) \wedge R(y, z)) \rightarrow (R(x, v) \vee R(v, z))$$

It is important to see that if we interpret the relation “ \succ_P ” as a semi-order, it is irreflexive and transitive, but need not be almost connected. Perhaps the easiest way to grasp what it means to be a semi-order is to look at its (intended) measure-theoretical interpretation. On the intended interpretation, “ $x \succ_P y$ ” is true iff the height of x is higher than the height of y plus some fixed (small) real number ϵ , which can be thought of as a *margin of error*.⁵ Indeed, as already suggested by Luce (1956) and

rigorously proved by Scott and Suppes (1958), if X is a finite set and ϵ a positive number, $\langle X, \succ_P \rangle$ is a semi-order if and only if there is a real-valued function f_P such that for all x ,

$$y \in X : x \succ_P y \text{ iff } f_P(x) > f_P(y) + \epsilon$$

This fact helps to explain the constraints. That the order should be irreflexive is trivial, because there can be no x such that $f_P(x) > f_P(x) + \epsilon$. As for (IO), consider two cases: $f_P(x) \geq f_P(v)$ or $f_P(v) \geq f(x)$. In the first case we have $f_P(x) \geq f_P(v) > f_P(w) + \epsilon$, and thus $x \succ w$. In the second case we have $f_P(v) \geq f_P(x) > f_P(y) + \epsilon$, and thus $v \succ y$. To see that (STr) holds, suppose that $x \succ_P y$ and $y \succ_P z$. Then $f_P(x) >_\epsilon f_P(y) >_\epsilon f_P(z)$ (with $a >_\epsilon b$ iff $a > b + \epsilon$). But then $f_P(v) \geq f_P(y)$ implies $v \succ z$, and $f_P(v) \leq f_P(y)$ implies $x \succ_P v$.

In terms of “ \succ_P ” we can define a similarity relation “ \sim_P ” as follows: $x \sim_P y$ iff neither $x \succ_P y$ nor $y \succ_P x$. The relation “ \sim_P ” is reflexive and symmetric, but need not be transitive. Thus, “ \sim_P ” does not give rise to an equivalence relation. One should think of this similarity relation as one of *indifference* or *indistinguishability*. Measure theoretically “ $x \sim_P y$ ” is true iff the difference in height between x and y is less than ϵ .⁶ In case $\epsilon = 0$, the semi-order is a weak order.

3.4.2 Some proposed solutions to the sorites

The standard reaction to the sorites paradox taken by proponents of fuzzy logic and/or supervaluation theory is to say that the argument is valid, but that the inductive premise [P] (or one of its instantiations) is *false*. But why, then, does it at least *seem* to us that the inductive premise is true? According to the standard accounts of vagueness making use of fuzzy logic and supervaluation theory, this is so because the instantiations of the inductive premise are *almost* true (in fuzzy logic), or almost all instantiations are true in the complete valuations (in supervaluation theory).

Linguists (e.g. Kamp 1975, Klein 1980, Pinkal 1995) typically do not like the fuzzy logic approach to vagueness, because that cannot account for what Fine (1975) called “penumbral” connections. The treatment of vagueness and the sorites paradox in supervaluation theory is not unproblematic either, however. The use of complete refinements in supervaluation theory assumes that we *can* always make sharp cutoff points: vagueness exists only because in daily life we are *too lazy* to make them. But this assumption seems to be wrong: vagueness exists, according to Dummett (1975), because we *cannot* make such sharp cutoff points

even if we wanted to. In terms of what we discussed above, this suggests that the relation “ \sim_P ” of indifference or indistinguishability should be intransitive, just as it is for a semi-order. But because this relation is still symmetric, it is very natural to claim that something like [P] is true.

For a while, the so-called “contextualist” solution to the sorites paradox was quite popular (e.g. Kamp 1981, Pinkal 1984, Raffman 1994, 1996). Most proponents of the contextualist solution follow Kamp (1981) in trying to preserve (most of) [P] by giving up some standard logical assumptions, and/or by making use of a mechanism of *context change*. But with Keefe (2007) we do not believe that context change is essential to save natural language from the sorites paradox. We rather believe that any solution involves some notion of *partiality*. We will briefly discuss two such proposed “solutions” in this section (without pretending to be complete or assuming that they are undoubtedly successful), and use the motivations behind those “solutions” in the following sections to propose some new cross-contextual constraints on the behavior of predicates which generate semi-orders.

A first solution is closely related with recent work of Raffman (2005) and Shapiro (2006) and makes use of partiality in a rather direct way:⁷ in terms of three-valued logic (Shapiro), or in terms of pairs of contrary antonymms (Raffman). The idea – just as what Klein (1980) proposed earlier – is that with respect to a comparison class c , predicate P and its antonym \bar{P} do not necessarily partition c , and there might be elements in c that neither (clearly) have property P nor property \bar{P} , but are somewhere “in the middle.” Once one makes such a move it is very natural to assume that the inductive principle [P] is not valid, but a weakened version of it, [P_w], is:

[P_w] If you call one individual tall in a particular context, and this individual is not visibly/relevantly taller than another individual, you will/should not call the other one short/not tall.

Thus, for any $x, y \in I, c \in C : (P(x, c) \wedge x \sim_P y) \rightarrow \neg \bar{P}(y, c)$.

Of course, principle [P_w] can only be different from the original [P] if $\neg \bar{P}(y, c)$ does not come down to the same as $P(y, c)$.⁸ Thus, a gap between the sets of P - and \bar{P} -individuals (with respect to c) is required. Notice that the sorites paradox can now be “solved” in a familiar way: $P(x_1, c)$ and $\bar{P}(x_n, c)$ are true in context c , and modus ponens is valid, but the inductive hypothesis, or (all) its instantiations, are not. However, because we adopt [P_w] as a valid principle of language use, we can

explain why inductive hypothesis [P] *seems* so natural. To illustrate, if $c = \{x, y, z\}$, it might be that with respect to a particular context structure $P(c) = \{x\}$, $\bar{P}(c) = \{z\}$, and $x \sim_P y \sim_P z$. Notice that such a context structure satisfies [P_w] but not [P].

A second “solution” is more radically *pragmatic* in nature and seems very much in line with Wittgenstein’s *Philosophische Untersuchungen*.⁹ In normal discourse, we talk about relatively few objects, all of which are easily discernible from the others. In those circumstances, [P] will not give rise to inconsistency, but serves its purpose quite well. Only in exceptional situations, i.e. when we are confronted with long sequences of pairwise indistinguishable objects, do things go wrong. But in such situations, we should not be using vague predicates like “tall” but precisely measurable predicates instead. A weak version of this reaction can be formalized naturally in terms of comparison classes. The idea is that it only makes sense to use a predicate P in a context – i.e. with respect to a comparison class – if it helps to clearly demarcate the set of individuals that have property P from those that do not. Following Gaifman (2002),¹⁰ we will implement this idea by assuming that any subset of X can only be an element of the set of *pragmatically appropriate* comparison classes C just in case the gap between the last individual(s) that have property P and the first that do(es) not must be between individuals x and y such that x is clearly, or significantly, P -er than y . This is not the case if the graph of the relation “ \sim_P ” is closed in $c \times c$.¹¹ Indeed, it is exactly in those cases that the sorites paradox arises. Notice also that this analysis makes use of partiality, but this now consists in the idea that certain comparison classes are not appropriate for the use of a particular predicate P .

How does such a proposal deal with the sorites paradox? Well, it claims that *in all contexts in which P can be used appropriately*, [P] is true. If we assume in addition that the first element x_1 of a sorites series is the absolute most P -individual, and the last element x_n the absolute least P -individual, it also claims that *in all contexts c in which it is appropriate to use predicate P in combination with x_1 and x_n* , “ $P(x_1, c)$ ” is true and “ $P(x_n, c)$ ” is false. Thus, in all *appropriate* contexts, the premises of the sorites argument are considered to be true. Still, no contradiction can be derived, because using predicate P when *explicitly* confronted with a set of objects that form a sorites series is *inappropriate*. Thus, in contrast to the original contextualist approaches of Kamp (1981), Pinkal (1984), and others, the sorites paradox is not avoided by assuming that the meaning (or extension) of the predicate changes as the discourse proceeds. Rather,

the sorites paradox is avoided by claiming that the use of predicate P is inappropriate when confronted with a sorites series of objects.¹²

The above sketch of some proposed solutions was rather unsophisticated and I do not want to claim in this chapter that they, or their more sophisticated variants, are completely successful. I also do not want to go into their relative pros and cons, or argue that they are (clearly) preferred to other proposals (though I sympathize with them). I only sketched them because the motivations behind those proposals clearly suggest some new cross-contextual constraints on the behavior of predicates which can be shown to generate (or even characterize) semi-orders.

3.4.3 Semi-orders and semantic gaps

The first of the above “solutions” to the sorites paradox is in essence *three-valued*. Either because a three-valued logic was used, or because we made use of pairs of antonyms. Recall again that Klein (1980) already used a three-valued logic: not all individuals in a particular comparison class need to be either tall or not-tall (in fact, Klein used supervaluations to make up for this “deficiency”). In this section we will indicate that once we follow this line of thought, it becomes easy to generate semi-orders, instead of weak orders. Our derivation makes use of *two choice functions*. Let us say that $P(c)$ selects the elements of c that (clearly) have property P (e.g. “tall”), while $\bar{P}(c)$ selects the elements that (clearly) have property \bar{P} (e.g. “short”). Now we can give the following four constraints:¹³

- (P*) $\forall c \in C : P(c) \cap \bar{P}(c) = \emptyset$
- (NR*) $\forall c, c' : D_P(c) \cap D_{\bar{P}}^{-1}(c') = \emptyset$ and $D_{\bar{P}}(c) \cap D_P^{-1}(c') = \emptyset$
- (UD*) If $c \subseteq c'$ and $D_{P\bar{P}}(c) \neq \emptyset$, then $D_{P\bar{P}}(c') \neq \emptyset$
- (DD*) If $c \subseteq c'$ and $D_{P\bar{P}}(c') \cap c^2 \neq \emptyset$, then $D_{P\bar{P}}(c) \neq \emptyset$

Constraint (P*) assures that P and \bar{P} behave as contraries, while (NR*) is the obvious generalization of van Benthem’s (1982) No Reversal constraint. Constraints (UD*) and (DD*) are very similar to the earlier Upward and Downward Difference constraints of van Benthem (1982), but still crucially different. The difference is that in this case we look at *contrary* pairs, and not merely at *contradictory* pairs. We define the ordering relation as follows: $x \succ_P y$ iff_{def} $x \in P(\{x, y\})$ and $y \in \bar{P}(\{x, y\})$. Then we can prove that this relation is irreflexive and transitive, but it need *not* satisfy *almost connectedness*: if $x \succ_P y$, it is possible that neither $x \succ_P z$ nor $z \succ_P y$, because (DD*) does not require either of them to hold if $P(\{x, y, z\}) = \{x\}$ and $\bar{P}(\{x, y, z\}) = \{y\}$. Now we can prove the following theorem (see van Rooij, 2011):

Theorem 1. Any context structure $\langle X, C, P, \bar{P} \rangle$ with X and C as defined above such that P and \bar{P} obey axioms (P^*) , (NR^*) , (UD^*) , and (DD^*) , gives rise to a semi-order $\langle X, \succ_P \rangle$, if we define $x \succ_P y$ as $x \in P(\{x, y\})$ and $y \in \bar{P}(\{x, y\})$.

3.4.4 Semi-orders and pragmatic gaps

Recall that according to the “pragmatic solution” of the sorites paradox not all subsets of X are assumed to be *appropriate* comparison classes. Whether c is an appropriate comparison class/context set was defined in terms of the relations “ \succ_P ” and “ \sim_P ”: the relation “ \sim_P ” should not connect all elements in c . In this section we want to turn that idea around: find some principles to generate all and only all appropriate context sets and then derive the relations “ \succ_P ” and “ \sim_P ” from that. The idea is that we just start with subsets of C that consist of two distinguishable elements and close this set of subsets of C under some closure conditions such that they will generate *all and only all appropriate contexts*. That is, the idea is to find some closure conditions such that we will generate just those subsets of X for which also vague predicates can clearly partition the context without giving rise to the sorites paradox. In conjunction with this, we will assume the same cross-contextual constraints on the behavior of P as van Benthem (1982) did, and define also $x \succ_P y$ as he did: $x \succ_P y$ iff_{def} $x \in P(\{x, y\})$ and $y \notin P(\{x, y\})$.

The closure conditions that jointly do this job are the following:

- (P1) $\forall c \in C : \forall x \in \bigcup C : c \cup \{x\} \in GAP \rightarrow c \cup \{x\} \in C$
- (OR) $\forall c \in C, \{x, y\} \in C : c \cup \{x\} \in C$ or $c \cup \{y\} \in C$
- (P2) $\forall c \in C, x \in X : c \in GAP_2 \rightarrow c \cup \{x\} \in C$

In terms of these constraints we want to generate all appropriate comparison classes starting with a set of appropriate comparison classes of just two elements. These constraints mention “ GAP ” and “ GAP_2 ,” which intuitively stand for gaps. Still, it is important to realize that the formalization does not make use of any predefined notion of a gap. The two notions “ GAP ” and “ GAP_2 ” will be defined in terms of such sets of appropriate comparison classes.

Before we discuss these constraints, it is important to see that the closure conditions do not guarantee that C necessarily contains *all* finite subsets of X (or better, not all subsets of X with cardinality 2 or 3). This is essential, because otherwise we could conclude with van Benthem (1982) that the resulting ordering relation would be a weak order and satisfies (AC) $\forall x, y, z : x \succ_P y \rightarrow (x \succ_P z \vee z \succ_P y)$. It suffices to observe that because neither GAP nor GAP_2 (both notions are defined below) is

always satisfied, no constraint formulated above forces us to assume that $\{x, y, z\} \in C$ if $x \succ_P y$, which is all that we need.

Now we will discuss these constraints. Constraint (P1) says that to any element c of C one can add any element $x \in X$ (thus, also $c \cup \{x\} \in C$) that is in an ordering relation with respect to at least one other element, on the condition that $c \cup \{x\}$ satisfies constraint *GAP*. To state this constraint, suppose that c contains n elements (written by c^n). Then the constraint says that there must be at least $n - 1$ subsets c' of c with cardinality $n - 1$ such that all these c' are also elements of C .

$$c^n \in \text{GAP} \quad \text{iff}_{\text{def}} \quad \exists_{n-1} c' \subset c : \text{card}(c') = n - 1 \wedge c' \in C$$

The *intuition* behind this condition is that only those subsets of X satisfy *GAP* if there is at least one gap in this subset with respect to the relevant property. It is easy to show that this closure condition guarantees that the resulting ordering relation will satisfy transitivity and will thus be a strict partial order.¹⁴ Constraint (OR) guarantees that if $\{x, y\}$ and $\{v, w\}$ are in C , then either $\{x, y, v\}$ or $\{x, y, w\}$ belongs to C as well. We will see below that by adopting this constraint the resulting ordering relation will satisfy the interval ordering condition. Constraint (P2) implements the intuition that if c gives rise to *two* gaps (again, only intuitively speaking), one can always add at least one arbitrary element of the domain to it, without closing all gaps. Constraint (P2) is defined in terms of predicate GAP_2 which is defined as follows:

$$c^n \in \text{GAP}_2 \quad \text{iff}_{\text{def}} \quad \exists_n c' \subset c : \text{card}(c') = n - 1 \wedge c' \in C$$

Notice the subtle difference between GAP_2 and *GAP*: whereas the former requires that there are at least n subsets of c^n in C with cardinality $n - 1$, the latter requires this only for $n - 1$ subsets. The intuition between this formal difference is that whereas c satisfies *GAP* already if it contains at least *one* gap, c can only satisfy GAP_2 if it has at least *two* gaps. Consider subsets of the natural numbers and assume that such a subset has a gap if at least one number in the order is missing. Thus, $\{1, 2, 3, 4\}$ has no gap, $\{1, 2, 3, 5\}$ has 1, but $\{1, 3, 4, 6\}$ and $\{1, 3, 5, 7\}$ have two or more gaps. The set $\{1, 2, 3, 5\}$ has the following subsets of three numbers with a gap: $\{1, 2, 5\}$, $\{1, 3, 5\}$, and $\{2, 3, 5\}$. Thus it has three such subsets, which means that $\{1, 2, 3, 5\}$ satisfies *GAP*, but not GAP_2 . The set $\{1, 3, 4, 6\}$, on the other hand, has four subsets of three elements with a gap: $\{1, 3, 4\}$, $\{1, 3, 6\}$, $\{1, 4, 6\}$, and $\{3, 4, 6\}$ which means that it satisfies both *GAP* and GAP_2 . The same holds for the set $\{1, 3, 5, 7\}$. The idea

behind constraint (P2) is that to the set $\{1, 3, 4, 6\}$ we can always add an arbitrary natural number and still have a gap (and thus be an appropriate context), but that this does not hold for $\{1, 2, 3, 5\}$: adding 4 would result in an inappropriate context. Now we can state the desired theorem (see van Rooij 2011):

Theorem 2. *Any context structure $\langle X, C, P \rangle$ with X a set of individuals, where P obeys axioms (NR), (DD), (UD) of section 3.2, and where C is closed under (P1), (OR), and (P2) gives rise to a semi-order $\langle X, >_P \rangle$, if we define $x >_P y$ as $x \in P(\{x, y\})$ and $y \notin P(\{x, y\})$.*

3.5 Comparisons revisited

Consider once more the following figure, picturing the lengths of John and Mary:



How can we account for the intuition that this picture allows us to say that the *explicit* comparative (1a) is true while the *implicit* comparative (1b) is false?

- (1a) John is taller than Mary.
 (1b) Compared to Mary, John is tall, but compared to John, Mary is not tall.

I would like to suggest that the difference between explicit and implicit comparatives is closely related with the difference between weak orders and semi-orders. As already suggested in section 3.1, weak orders are very natural representations of standard explicit comparatives like (1a). I propose that the semi-order relation *significantly taller than*, i.e. “ $>_T$,” is what is relevant to evaluate the truth of implicit comparatives like (1b). Thus, (1b) is true just in the case when John is significantly taller than Mary. This immediately explains why (1a) can be inferred from (1b), but not the other way around.

A weak order “ $>_P$ ” can be thought of as *at least as informative*, or *refined*, as a corresponding semi-order “ $>_P$ ” in the sense that for all $\langle x, y \rangle \in X \times X$, if $x >_P y$, then $x > y$ as well. There is, however, another sense in which it is natural to think of the semi-order as the basic one, and *derive* a

corresponding weak order. Note, though, that for an arbitrary semi-order there might always be several weak orders that are compatible with it. The following two weak orders can, for instance, be derived from the semi-order “ $>_P$ ”: “ $>_P^1$ ” defined as

$$x >_P^1 y \text{ iff}_{def} \exists Z : (x \sim_P Z \wedge Z >_P y)$$

and “ $>_P^2$ ” defined as

$$x >_P^2 y \text{ iff}_{def} \exists Z : (x >_P Z \wedge Z \sim_P y)$$

Fortunately, for each semi-order there is also a *unique most refined* weak order that can be derived from it. As already shown by Luce (1956), this *unique* strict weak order “ $>_P$ ” can be defined as follows:


$$x >_P y \text{ iff}_{def} \exists Z : (x \sim_P Z \wedge Z >_P y) \vee (x >_P Z \wedge Z \sim_P y)$$

The corresponding relation “ \approx_P ” defined as $x \approx_P y \text{ iff}_{def} x \not>_P y$ and $y \not>_P x$ is an equivalence relation, which could also be defined directly as

$$x \approx_P y \text{ iff}_{def} \forall Z \in I : x \sim_P Z \text{ iff } y \sim_P Z$$

What I would like to suggest is that if we start with the semi-order “ $>_P$ ” in terms of which we interpret implicit comparatives, it is the strongest derived weak order “ $>_P$ ” that is relevant to interpret explicit comparatives.

Recall from section 3.1 that if “ P ” is “tall,” measure theoretically “ $x \sim_P y$ ” is true iff the difference in length between x and y is less than a fixed *margin of error* ϵ . Suppose that John’s length is $\delta < \epsilon$ higher than Mary’s length, which is $\delta < \epsilon$ higher than Sue’s length, but that the difference between John’s and Sue’s length is higher than ϵ . This situation can be pictured as follows:



In this situation, $j \sim_P m$, $m \sim_P s$, but $j >_P s$. Given our claim that implicit comparatives should be interpreted as “significantly taller than”

using semi-orders, we correctly predict that (1b) is false: $j \not>_p m$. But given that we interpret implicit comparatives in terms of the weak order “ $>_p$ ” defined in terms of the semi-order, we also correctly predict that (1a) comes out as being true. The reason is that $j >_p m$, because there is an s , i.e. Sue, such that $j >_p s$ and $s \sim_p m$. This is an encouraging result, and enough to “explain” the difference between (1a) and (1b). But we wanted more: we wanted to explain the difference between explicit and implicit comparatives in terms of the behavior of adjectives in comparison classes. In the remainder of this section I want to discuss (i) how we can do that, and (ii) whether the way we suggested to interpret explicit comparatives is really strong enough.

An obvious way to account for the distinction between explicit and implicit comparisons in terms of comparison classes is just to look at the differences between the constraints on context structures in section 3.1 versus those in sections 3.4.3 and 3.4.4. If we adopt the semantic model of section 3.4.3 we might say that the distinction between explicit versus implicit comparatives corresponds to whether we assume the equivalence of “It is not the case that Mary is tall in c ” and “Mary is small in c ” or not. If we adopt the pragmatic model of section 3.4.4, on the other hand, the distinction between explicit versus implicit comparatives can be said to correspond with what we take to be appropriate comparison classes: *all* subsets of X , or only those in which there is a significant gap.

But what if this freedom is not allowed? What if we cannot play with what is an appropriate comparison class and assume the equivalence of “It is not the case that Mary is tall in c ” and “Mary is small in c ”? Even in that case, I claim, we can make a distinction between explicit and implicit comparatives, because adopting the analysis of section 3.4.3, we can make a difference between (Klein) and (Benthem), as introduced in the beginning of this chapter. Thus, I propose the following interpretation rules:


- (Klein) (1a) is true iff *there is* a comparison class such that John is tall here and Mary is not: $\exists c : M, c \models T(j) \wedge \neg T(m)$.
- (Benthem) (1b) is true iff John is tall in comparison class $\{\text{John, Mary}\}$, and Mary is not: $M, \{j, m\} \models T(j) \wedge \neg T(m)$.

Let us first note that in the semantic model of section 3.4.3, it is now at least possible that (1a) is true, but (1b) is false. It is very natural to assume that in the comparison class $\{\text{John, Mary, Sue}\}$, we count John as tall, Sue as short, and Mary as neither tall nor short. But this is enough for the explicit comparative (1a) to be true in the above situation. Moreover, it

is natural to assume that in the comparison class {John, Mary}, we count John as tall if and only if Mary is counted as tall, which means that the implicit comparative (1b) is correctly predicted to be false. Unfortunately, the constraints given in section 3.4.3 do not rule out the possibility that given the above situation, John is counted as tall in comparison class {John, Mary} but Mary is not. But in that situation (1b) is falsely predicted to be true, just as (1a). A situation like this is ruled out if we adopt the natural constraint that for all comparison classes c , $P(c) = \emptyset$ iff $\bar{P}(c) = \emptyset$. Adopting this constraint, it can only be the case that John is counted as tall in comparison class {John, Mary} but Mary not if Mary is counted as *short* in this comparison class, and thus that John is *significantly taller* than Mary. But this is in contradiction with what we assumed.

How does this proposal relate with our earlier suggestion to interpret explicit comparatives in terms of the unique strongest weak order derived from a semi-order? We can prove that it comes down to the same thing. Suppose an explicit comparative “ x is P -er than y ” is true according to (Klein) because there is a comparison class c such that $x \in P(c)$ and $y \in c - P(c)$. If $c = \{x, y\}$ then $x >_P y$ holds. Because $>_P \subseteq \succeq_P$ it follows that $x \succeq_P y$, so that is OK. But now suppose that it is not the case that $x \in P(\{x, y\})$ and $y \notin P(\{x, y\})$. Then there must be a superset c of $\{x, y\}$ for which $x \in P(c)$ and $y \notin P(c)$ holds. Adopting constraint (DD*) we have already ruled out the possibility that $y \in \bar{P}(c)$. So it must be that $y \in c$ but $y \notin P(c)$ and $y \notin \bar{P}(c)$. Let c for instance be $\{x, y, z\}$. By our constraint that for all comparison classes c , $P(c) = \emptyset$ iff $\bar{P}(c) = \emptyset$ it follows that $z \in \bar{P}(c)$. By (DD*) it follows that $x \in P(\{x, z\})$ and $z \in \bar{P}(\{x, z\})$, and thus that $x >_P z$. What we have to show is that $z \sim_P y$. Because $z \in \bar{P}(c)$, $y \notin P(c)$ and $y \notin \bar{P}(c)$, it follows by (NR*) that it is not the case that $z \in P(\{y, z\})$ and $y \in \bar{P}(\{y, z\})$, which means that $z \not>_P y$. But that means that $y \succeq_P z$. If $y \not>_P z$ we are done, so suppose $y >_P z$. In that case it is *natural to assume* that *there is another* (i.e. taller) z' such that $x \in P(\{x, y, z'\})$, $z' \in \bar{P}(\{x, y, z'\})$, $y \notin P(\{x, y, z'\})$ and $y \notin \bar{P}(\{x, y, z'\})$ and for which $y \not>_P z'$. But this means that $y \sim_P z'$ which is what we wanted.

The last but one sentence indicated something important that, in fact, already appeared much earlier: if we want to guarantee that John is counted as taller than Mary if their relative lengths are pictured as below,



we have to assume that we have enough other individuals John and Mary can be compared with. There either has to be somebody like Sue who is significantly shorter than John but similar to Mary, or another individual who is significantly taller than Mary but similar to John. A very natural way to guarantee these kind of *witnesses* to exist is to adopt the following constraints on models: for all individuals y with at most two exceptions (the tallest and shortest individuals, if they exist), $\exists x, z : x \succ_P z \wedge x \sim_P y \sim_P z$. For the other two individuals v , if they exist, it just holds that $\exists w : w \neq v \wedge v \sim_P w$. Thus, we demand that with at most two exceptions, any object is “indistinguishable” from at least two others. If we take semi-orders to be primitive, this constraint has a direct effect. Otherwise, the constraint should be reformulated in terms of context structures. In whatever way we do this, it is clear that it has the desired effect: any small difference in length between John and Mary is enough to make the explicit comparative true.

Adopting the above type of witness constraint is costly, but how costly is it really? Degree-based theories of comparatives make use of witness constraints as well. If we look at the algebraic structures that are faithfully represented by means of the measure functions (see Krantz et al. 1971), we see that in case the numbers are really crucial (as is the case in so-called “interval scales” and “ratio scales”) it has to be assumed that there *exists a witness for every possible degree* required for the homomorphic function. It is clear that our witness constraint is much less involved, i.e. we did not secretly presuppose degrees after all.

3.6 Conclusion

I claimed in this chapter that the distinction between explicit and implicit comparatives corresponds to the difference between (strict) weak orders and semi-orders. Moreover, I showed that both can be characterized naturally in terms of constraints on the behavior of predicates among different comparison classes, and thereby meeting the challenge Kennedy and others have posed upon comparison class-based approaches of comparative statements. How can degree-based approaches account for the difference between explicit and implicit comparatives? The most natural way for them to make a distinction between (1a) and (1b) would be to claim that for the latter there must be a specific number $\epsilon > 0$ such that the length of John minus ϵ is more than the length of Mary. But if we think of ϵ as the threshold, degree-based approaches must make a distinction between explicit and implicit comparatives very much like we did involving weak and semi-orders.

An interesting question that arises is whether we really want the threshold to be the same for any pair of individuals. This was assumed in this chapter (by making use of semi-orders), but should perhaps be rejected in general. In De Jaegher and van Rooij (to appear) it is shown that prospect theory can be used to account for the intuition that the threshold depends on the individuals involved. It would take us too far to investigate this issue here.

Notes

* I would like to thank an anonymous reviewer and the editors (Paul Égré and Nathan Klinedinst) for their useful comments on an earlier version of this chapter. I would like Chris Kennedy for stating the challenge and for pointing out the reference to Sapir.

1. More complex sentences suggest that this simple picture is naive, and there has been a lot of discussion of how to improve on it. I will ignore this discussion in this chapter.
2. But see von Stechow (1984) for an argument saying that also the degree-based approach is in line with Frege's principle.
3. According to Chris Kennedy (p.c.) the names "explicit" versus "implicit" comparatives go back to Sapir (1944).
4. Any relation that is irreflexive and satisfies the interval-order condition is called an *interval order*. All interval orders are also transitive, meaning that they are stronger than strict partial orders.
5. One can think of Williamson's (1990) epistemic analysis of vagueness based on semi-orders as well.
6. The fact that " \sim_p " is intransitive has the consequence that semi-orders cannot be given a full measure-theoretic interpretation f in the sense that there is no set of transformations such that f is *unique* up to this set of transformations. This fate it shares with, among others, partial orders.
7. Shapiro (2006) argues that his solution is closely related to Waismann's notion of "open texture." For what it is worth, I believe that Waismann's notion is more related to what I call the second "solution" of the sorites discussed on p. 61.
8. In this chapter I do not really distinguish thinking of the comparison class as part of the context (as I did until now), or of thinking of it as an argument of an adjective. For present purposes, this distinction is irrelevant.
9. See in particular section 85–87: "A rule stands like a signpost ... The signpost is in order – if, under normal circumstances, it fulfills its purpose." The observation that our pragmatic solution is very much in line with Wittgenstein's later philosophy, I owe to Frank Veltman.
10. One might argue that Gaifman's solution was already anticipated – though in a rather different way – by Kamp (1981). A theory much more similar to Gaifman's was proposed by Pagin (this volume). The editors of this volume pointed out to me that Gomez-Torrente (2010) argues for much the same idea.
11. Notice also that in discrete cases the relation " \sim_p " can be closed in $c \times c$. It just depends on how " \sim_p " is defined.

12. Williamson (p.c.) and a reviewer of this chapter ask what are the *semantic* consequences of using a pragmatically inappropriate comparison class. The main answer is that if pushed one can still choose between, for instance, an epistemic approach or a three-valued approach. Adopting this approach, the answer to this question should, I think, be of little theoretical importance: I do not think we have very strong semantic intuitions about things that go against what we ought to do and normally do.
13. The formulation of the constraints is much simpler, though equivalent, to the formulation I used in van Rooij (2009). I thank Frank Veltman for pointing out that my earlier formulation was needlessly complex.
14. For a proof of this result, and the others mentioned below, see van Rooij (to appear).

References

- Dummett, M. (1975). Wang's paradox. *Synthese*, 30: 301–24.
- Fine, K. (1975). Vagueness, truth and logic. *Synthese*, 30: 265–300.
- Gaifman, H. (2002). Vagueness, tolerance and contextual logic, manuscript, Columbia University.
- Gomez-Torrente, M. (2008). The sorites, linguistic preconceptions and the dual picture of vagueness. In Dietz, R. and Moruzzi, S. (eds), *Cuts and Clouds. Essays in the Nature and Logic of Vagueness*. Oxford University Press, Oxford.
- Graff, D. (2000). Shifting sands: an interest relative theory of vagueness. *Philosophical Topics*, 28: 45–81.
- Jaegher, K. de and van Rooij, R. (to appear). Strategic vagueness and appropriate contexts. In Benz, A. et al. (eds), *Games and Pragmatics*. Texts in Logic and Games, LNCS, Springer.
- Kamp, H. (1975). Two theories of adjectives. In Keenan, E. (ed.), *Formal Semantics of Natural Language*, pp. 123–55. Cambridge University Press, Cambridge.
- Kamp, H. (1981). The paradox of the heap. In Mönnich, U. (ed.), *Aspects of Philosophical Logic*, pp. 225–77. D. Reidel, Dordrecht.
- Keefe, R. (2007). Vagueness without context change. *Mind*, 116: 462.
- Kennedy, C. (1999). *Projecting the Adjective: the Syntax and Semantics of Gradability and Comparison*. Garland Press, New York.
- Klein, E. (1980). The semantics of positive and comparative adjectives. *Linguistics and Philosophy*, 4: 1–45.
- Klein, E. (1991). Comparatives. In von Stechow, A. and Wunderlich, D. (eds), *Semantik: Ein internationales Handbuch der zeitgenössischen Forschung*. Walter de Gruyter.
- Krantz, D., Luce, R. Duncan, Suppes, P. and Tversky, T. (1971). *Foundations of Measurement*, Vol. 1. Academic Press, New York and London.
- Lewis, D. (1970). General semantics. *Synthese*, 22: 18–67.
- Luce, R. D. (1956). Semiorders and a theory of utility discrimination. *Econometrica*, 24: 178–91.
- Pinkal, M. (1984). Consistency and context change: the Sorites paradox. In Landman, F. and Veltman, F. (eds), *Varieties of Formal Semantics, Proceedings of the Fourth Amsterdam Colloquium*. Groningen–Amsterdam Studies in Semantics, Foris Publications, Dordrecht.

- Raffman, D. (1994). Vagueness without paradox. *The Philosophical Review*, 103: 41–74.
- Raffman, D. (1996). Vagueness and context-relativity. *Philosophical Studies*, 81: 175–92.
- Raffman, D. (2005). Borderline cases and bivalence. *Philosophical Review*, 114: 1–31.
- Sapir, E. (1944). Grading: a study in semantics. *Philosophy of Science*, 11: 93–116.
- Scott, D. and Suppes, P. (1958). Foundational aspects of theories of measurement. *Journal of Symbolic Logic*, 23: 113–28.
- Seuren, P. (1973). The comparative. In Kiefer, F. and Ruwet, N. (eds), *Generative Grammar in Europe*, pp. 528–64. Reidel.
- Shapiro, S. (2006). *Vagueness in Context*. Clarendon Press, Oxford.
- Stechow, A. von (1984). Comparing semantic theories of comparison. *Journal of Semantics*, 3: 1–77.
- van Benthem, J. (1982). Later than late: on the logical origin of the temporal order. *Pacific Philosophical Quarterly*, 63: 193–203.
- van Rooij, R. (2011). Revealed preference and satisficing behaviour. *Synthese*, 2011. DOI 10.1007/s11229-010-9826-z.
- van Rooij, R. (to appear). Vagueness and language. In Ronzitti, G. (ed.), *The Vagueness Handbook*. Springer.
- Waismann, F. (1968). Verifiability. In Flew, A. (ed.), *Logic and Language*. Basil Blackwell, Oxford.
- Wittgenstein, L. (1953). *Philosophical Investigations/Philosophische Untersuchungen*. Oxford, Blackwell.

4

Vagueness and Comparison*

Christopher Kennedy
University of Chicago

4.1 Introduction

Vagueness and comparison are linked in a number of different ways. The most obvious is that, at least when they are expressed by adjectives (in English), vague predicates have morphosyntactically comparative forms. These forms are often invoked in formulations of the inductive premise of the sorites paradox:

- (1) If x is {tall, big, heavy, ...}, and y is just a little bit {short-, small-, light-, ...}-**er than** x , then y is {tall, big, heavy, ...}.

The relation between vagueness and comparison is also exploited in a class of important semantic analyses of morphological comparatives that attempt to explain their truth conditions in terms of prior analyses of the vagueness of the morphologically unmarked forms (Wheeler 1972, Kamp 1975, Klein 1980, van Benthem 1982, van Rooij this volume). Finally, constraints on judgments about pairwise comparisons are often invoked in the development and evaluation of theories of vagueness. In particular, a crucial evaluation criterion for a theory of vagueness is that it should derive the descriptive constraint in (2) (Fara 2000).

- (2) *The Similarity Constraint*

When x and y differ to only a very small degree in the property that a vague predicate g is used to express, we are unable or unwilling to judge the proposition that x is g true and that y is g false.

Theories differ both in whether they derive (2) and in how: some derive (2) as a function of language use (e.g. Soames 1999); some as a by-product

of epistemic uncertainty (e.g. Williamson 1994); and some as a feature of the meaning of vague predicates (e.g. Raffman 1994, 1996, Fara 2000).

My goal in this chapter is to take a close look at two ways of expressing comparison, which differ in both their morphosyntactic properties and semantic/pragmatic properties, with the goal of showing how they can help us assess theories of vagueness and explanations of the Similarity Constraint on the one hand, and semantic analyses of the positive and the comparative forms (and the relation between them) on the other. The facts will suggest that the Similarity Constraint (and so features of vagueness more generally) is due to a semantic property of vague predicates, and that this property is a feature of the positive form but not the comparative form. This can be easily accommodated if both forms are derived from a more abstract source, but it is difficult (though perhaps not impossible) to explain if the comparative is derived from the positive.

4.2 Explicit and implicit comparison

4.2.1 Modes of comparison

Consider the asymmetric size relation between the planets Uranus and Venus, as determined by diameter, which is shown to scale in Figure 4.1. (To make differences in diameter easily perceptible, I will represent the sizes of the planets as concentric circles in the figures to follow.) A speaker might describe this relation by uttering one of the following sentences:

- (3) a. Uranus is bigger than Venus.
 b. Venus is smaller than Uranus.
 c. Uranus is the bigger one/of the two.
 d. Venus is the smaller one/of the two.

The main predicates in (3a, b) and the pronominal modifiers in (3c, d) contain morphologically comparative forms of the adjectives *big* and




Figure 4.1 Uranus (51,118 km) vs Venus (12,100 km)

small. Like other kinds of degree morphology, the function and conventional meaning of the comparative is to introduce an ordering entailment, which is relativized to the polarity of the adjective. (3a, b) entail that the target of comparison (the subject) is ordered above the standard (marked by *than*) relative to the appropriate size concept (increasing for *big*; decreasing for *small*). (3c, d) each presuppose (thanks to the contribution of the definite article) that there is a unique object that is ordered above another one relative to the relevant size concept, and entail that the subject is that object.

The size relation shown in Figure 4.1 could also be described using one of the sentences in (4), which differ from those in (3) in that the main predicates and the modifiers do not include overt degree morphology, but instead involve the unmarked, “positive” form of the adjectives.¹

- (4) a. Uranus is big compared to Venus.
 b. Venus is small compared to Uranus.
 c. Uranus is the big one.
 d. Venus is the small one.

The fact that the positive form can be used to express comparison follows from two features of its meaning. First, it is context-dependent: what counts as *big* or *small* can vary according to a number of different contextual factors, such as an implicit or explicit comparison class (Wheeler 1972, Klein 1980), the interests/expectations of a designated discourse participant (Bogusławski 1975, Fara 2000, Richard 2004), or other salient contextual information (Kamp and Partee 1995). Second, no matter how the meaning and context sensitivity of the positive form are ultimately cashed out, the two “consistency constraints” in (5a, b) hold.²

- (5) *Consistency constraints*:
- a. For any positive form gradable predicate g and objects in its domain x, y and for any context c , if $g(x)(c)$ is true and $g(y)(c)$ is false, then x exceeds y relative to the scalar concept encoded by g .
- b. For any positive form gradable predicate g and objects in its domain x, y , if there is a context c such that $g(x)(c)$ is true and $g(y)(c)$ is false, then for any c' such that $g(y)(c')$ is true, $g(x)(c')$ is also true.

These features of positive form meaning interact with other general semantic/pragmatic properties to derive the comparative entailments in the sentences in (4). Given a general “informativity” constraint on

predicate valuations which requires both the positive and negative extension of a vague predicate to be nonempty (such predicates are not useful if there are not things that they are true of and things that they are false of; see Klein 1980), we can account for the use of (4a, b) to make comparisons by hypothesizing that the function of *compared to* is to evaluate the main predication of *big/small* relative to a context consisting only of Uranus and Venus (Kennedy 2007a). Asserting that, for example, Uranus is big relative to such a context entails that Venus is not big (in that context); given (5a, b), we may conclude that Uranus is bigger than Venus in all contexts, including the context of utterance.

In (4c, d), a similar result is achieved thanks to the uniqueness presupposition of the definite article (Syrett et al. 2010). An assertion of, for example, (4c) in a context containing just Uranus and Venus commits the speaker to the proposition that Uranus is big in that context and that Venus is not, since this is the only way to satisfy uniqueness relative to the definite description (in the absence of prior context that could supply anaphoric *one* with a meaning that would do the job). Given (5a, b), it must also be the case that Uranus is bigger than Venus.

Borrowing terminology from Sapir (1944), Kennedy (2007a) refers to constructions like those in (3) as instances of EXPLICIT COMPARISON, and constructions like those in (4) as instances of IMPLICIT COMPARISON, providing the definitions for these terms in (6)–(7).

(6) *Explicit comparison:*

Constructions which establish an ordering between objects x and y with respect to gradable property g by using a morphosyntactic form of g whose conventional meaning has the consequence that the degree to which x is g exceeds the degree to which y is g .

(7) *Implicit comparison:*

Constructions which establish an ordering between objects x and y with respect to gradable property g by using the positive form of g and manipulating the context in such a way that the positive form is true of x and false of y .

The point of making this distinction in Kennedy (2007a) was to explore the possibility that the explicit/implicit distinction is a point of typological variation in the expression of comparison in the world's languages, and in particular, to ask whether some languages have only implicit comparison. This possibility is suggested by the fact that some languages lack overt comparative morphology and instead rely on collocations that appear to involve the positive form in order to

express comparison. A particularly striking example is the class of “conjoined comparatives” (Stassen 1985), illustrated by the Samoan example in (8).³

- (8) Ua tele le Queen Mary, ua la’itiiti le Aquitania.
 is big the Queen Mary, is small the Aquitania
 ‘The *Queen Mary* is bigger than the *Aquitania*.’

4.2.2 Crisp judgments

Although the superficial morphosyntactic facts of languages like Samoan lend some credence to the idea that some languages might lack explicit comparison, given the existence of null morphology, we cannot draw this conclusion simply on the basis of the absence of an overtly marked comparative form. Instead we need to identify tests that differentiate between instances of (6) and instances of (7) on a distributional, semantic or pragmatic basis. Several such tests are provided in Kennedy (2007a) (see also Sawada 2009); of interest to us here is one that involves CRISP JUDGMENTS: whether a particular expression can be used to describe differences of a very small degree.

For an illustration, consider Figure 4.2, which shows the relative size of Uranus and Neptune. As the picture indicates, the two planets differ in size by a relatively small amount. This difference could be felicitously characterized by any of the explicit comparison constructions in (9):

- (9) a. Uranus is bigger than Neptune.
 b. Neptune is smaller than Uranus.
 c. Uranus is the bigger one/of the two.
 d. Neptune is the smaller one/of the two.

In contrast, the implicit comparison constructions in (10) are infelicitous: they do not support crisp judgments.⁴




Figure 4.2 Uranus (51,118 km) vs Neptune (49,500 km)

- (10) a. #Uranus is big compared to Neptune.
 b. #Neptune is small compared to Uranus.
 c. #Uranus is the big one.
 d. #Neptune is the small one.

At first glance, the infelicity of these sentences as descriptions of the scenario in Figure 4.2 appears to follow straightforwardly, given that the kinds of judgments involved in evaluating them are exactly the kind of judgments that the Similarity Constraint makes reference to. If this constraint applies to any context of evaluation of a vague predicate, the similarity in size between Uranus and Neptune means that either both planets must be in the positive extension of the predicate or both must be in the negative extension. If the semantic characterization given above for *compared to* sentences is correct, then (10a, b) are infelicitous because they violate the constraint that in every context of evaluation, both the positive and negative extension of the predicate should be nonempty. Similarly, (10c, d) violate the presuppositions of the definite, since it must be the case (according to (2)) either that both planets count as big/small (violating uniqueness) or that both fail to count as big/small (violating existence).

However, as noted above, the Similarity Constraint is a descriptive generalization that should be derived rather than stipulated, and different theories of vagueness derive it in different ways, typically in the context of explaining judgments about the inductive premise of the sorites. But the inductive premise is a universal claim about objects that are not under active consideration, while the comparisons in (10) involve claims about two objects whose relevant properties are directly observable. It is possible, then, that different theories of vagueness might provide equally adequate accounts of sorites judgments, but still make different predictions about data like (10). At the same time, we need to ask of particular compositional semantic analyses of positive and comparative adjectives whether they make the right distinction between explicit and implicit comparatives in crisp judgment contexts. These are the questions I take up in the next section.

Before moving to this discussion, though, let me quickly point out that the facts we are considering here really are facts about vague positive form gradable predicates in particular, not facts about the positive-comparative distinction more generally. As discussed in Kennedy (2007b), there exist classes of gradable predicates that have both positive and comparative forms, but which are not vague. Borrowing terminology from Unger (1975) (see also Kennedy and McNally 2005), Kennedy refers to

gradable predicates that are vague in the positive form as *RELATIVE* and those that are not vague as *ABSOLUTE*. Unlike relative gradable predicates, absolute gradable predicates do not display crisp judgment effects in implicit comparison constructions.⁵

This can be seen by considering the adjective *old*, which has both an absolute and a relative sense. The relative sense behaves in the same way as *big* in crisp judgment contexts. If we are looking at two young boys, Julian and Sterling, who are similar in appearance, but have birthdays that are two years apart, we might distinguish them by saying either of the sentences in (11):

- (11) a. Julian is the old one.
b. Julian is the older one.

If their birthdays are two days apart, however, only (11b) is acceptable.

The absolute sense of *old* behaves differently. Imagine a context in which there are two occurrences of the file `foo.txt`, one on my laptop computer and one on my desktop computer, which are identical except that the one on my desktop was copied from my laptop, immediately modified by changing one character in the text, and then saved. I could then felicitously describe the situation by saying either (12a) or (12b):

- (12) a. The file on my laptop is the old one.
b. The file on my laptop is the older one.

This sense of *old* is roughly similar to (though not the same as) the meaning of *former*, and so is in some sense inherently comparative. In the terms of Kennedy (2007b), it is probably best classified as a “minimum standard” absolute adjective, since it is true of any object whose age (in the relevant sense) diverges by some positive (but possibly quite small) degree from whatever is “most recent.” This is not a vague concept, and so (12a) is correspondingly acceptable in crisp judgment contexts.⁶

4.3 Implications for analyses of vagueness and comparison

4.3.1 Derived vagueness

As a starting point, let us consider analyses of gradable predicates, comparatives and vagueness that work together to make exactly the distinction we need to make in order to explain the crisp judgment facts. Clearly, any analysis in which the positive form contains an element

of meaning that has the Similarity Constraint as a consequence, and in which the comparative form lacks this element of meaning, will be one that makes the right predictions. Assuming compositionality, such a theory would necessarily be one in which the comparative is not derived from the positive, but instead both forms must be derived from a more basic meaning which does not itself contain the element of meaning responsible for crisp judgment effects. There are several such theories on the market; the one I will present here is based on the hypothesis that gradable predicates like *big* do not denote properties of individuals, but rather denote functions from individuals to scalar values, traditionally called DEGREES (Bartsch and Vennemann 1973, Kennedy 1999, 2007b). Gradable predicates are converted into properties by DEGREE MORPHOLOGY; different degree morphemes introduce different requirements into the truth conditions.⁷

There are a number of truth-conditionally equivalent ways of stating the denotation of the comparative morphology given these initial assumptions, which differ in their specific claims about the semantic contribution of other parts of the sentence, such as the semantic type of the gradable predicate (see note 7) and the compositional interpretation of the standard phrase (the *than*-constituent). Since these distinctions are not relevant to the main point of this chapter, we can assume that the meaning of the comparative morpheme in the kinds of examples we are considering is (13a): MORE combines with a gradable adjective and returns a relation between individuals which is true in a context of utterance iff the predicate maps the target of comparison (the external argument) to a higher value on the relevant scale than the standard.

- (13) a. $\llbracket \text{MORE} \rrbracket^c = \lambda g_{(e,d)} \lambda y \lambda x. g(x) > g(y)$
 b. $\llbracket \text{MORE big} \rrbracket^c = \lambda y \lambda x. \mathbf{big}(x) > \mathbf{big}(y)$

For example, composition of MORE with the adjective *big*, as in (13b), returns a relation between individuals which is true just in case the size of the target exceeds the size of the standard.

Importantly, in this kind of analysis, the “unmarked” positive form must also include a degree morpheme, albeit a phonologically null one, since direct composition of a gradable predicate with an individual returns a degree, not something that is truth evaluable.⁸ Since this morpheme is present in the positive form but not in other forms (and not present in the adjectival root), this kind of analysis provides a straightforward means of accounting for differences between the positive and

other forms: we simply assign to the positive form degree morphology whatever semantic features need to be invoked to explain those differences. Since our interest is in saying why the positive but not the comparative form is subject to the Similarity Constraint, let us hypothesize with Fara (2000) that the denotation of the positive morpheme (henceforth POS) is (14a), where $>_1$ is the interest-relative relation “significantly exceed” and **stnd** is a function that, given an adjective meaning and a context, picks out an appropriate standard of comparison in the context for the kind of measurement encoded by the adjective (cf. Bogusławski 1975, Richard 2004, Kennedy 2007b).

- (14) a. $\llbracket \text{POS} \rrbracket^c = \lambda g_{(e,d)} \lambda x. g(x) >_1 \text{stnd}(g)(c)$
 b. $\llbracket \text{POS big} \rrbracket^c = \lambda x. \text{big}(x) >_1 \text{stnd}(\text{big})(c)$

Composition of POS with e.g. *big* returns the interest-relative property in (14b), which is true of an object iff its size exceeds a contextual standard in size by a degree that is significant given some set of interests. According to Fara, such a property is subject to the Similarity Constraint because our interests (in particular our interests in efficiency) are such that a small difference in size can never be significant. That is, it can never be the case that one object exceeds a standard of comparison for size in a way that is significant given my interests, while another object that is nearly the same size does not, and vice versa.

Whether Fara’s analysis is ultimately the right way to derive the Similarity Constraint is to some extent beside the point (see Stanley 2003 for criticisms, and Fara 2008 for a response); I adopt it here because it provides a way of illustrating the more important compositional point: in this kind of analysis of adjective meaning, the positive form, and the comparative form, it is POS that is the locus of whatever aspect of meaning derives Similarity, not MORE and not the adjectival root. This distinction allows us to explain the difference between explicit and implicit comparison constructions in crisp judgment contexts.

First consider the explicit comparative (15a), in which the predicate denotes the property in (15b):

- (15) a. Uranus is bigger than Neptune.
 b. $\lambda x. \text{big}(x) > \text{big}(\text{neptune})$

This property is true of an object just in case its size exceeds that of Neptune, so an assertion of (15a) just commits the speaker to the claim that there is a size difference between the two planets, which is a perfectly

reasonable thing to assert, and also true, given the facts represented in Figure 4.2.

Now consider the implicit comparative (16a). Let us assume that the semantic function of *compared to* is to modify the contextual parameter relative to which the positive form adjective is evaluated. Specifically, it causes the adjective to be evaluated relative to a context that is just like the actual one, except that its domain includes just the two planets Uranus and Neptune (Kennedy 2007a, b; cf. Wheeler 1972). Representing such a context as $c[\mathbf{uranus}, \mathbf{neptune}]$, the predicate in (16a) denotes the property in (16b):

- (16) a. #Uranus is big compared to Neptune.
 b. $\lambda x.\mathbf{big}(x) \succ_! \mathbf{stnd}(\mathbf{big})(c[\mathbf{uranus}, \mathbf{neptune}])$

Given the informativity constraint on predicate valuations, which requires both the positive and negative extension of a vague predicate to be nonempty, an assertion of (16a) commits a speaker to the position that Uranus has a degree of size that significantly exceeds the standard in the {Uranus, Neptune} context, but Neptune does not. This entails that Uranus is bigger than Neptune, and would therefore seem to predict that (16a) is true as a description of the state of affairs in Figure 4.2. However, if Fara's claims about how the interest-relative semantics of the positive form derives the Similarity Constraint are correct, then it cannot be the case that *big* can be true of Uranus and false of Neptune in any context. We thus have a conflict between the meaning of the predicate in (16a) and the informativity constraint on predicate valuations; this conflict is source of the unacceptability of the example. (The examples with definite descriptions work the same way, except that in these cases the conflict arises from the clash between the interest-relative meaning of the adjective and the uniqueness/existence presuppositions of the definite article.)

It is important to recognize that the analysis of explicit and implicit comparatives in crisp judgment contexts that I just sketched has two parts, which are logically distinct: the interest-relative analysis of vagueness and the compositional analysis of positive and comparative adjectives. The interest-relative analysis of vagueness provides an account of the Similarity Constraint, and could be replaced with a different theory of the semantic source of vagueness that achieves the same result.⁹ The compositional analysis of positive and comparative adjectives provides a way of ensuring that this constraint applies only to the positive form, not to the comparative form, by locating the source of the Similarity

Constraint in the meaning of the POS morpheme and not in the adjective. We may now ask whether different assumptions about vagueness and different assumptions about the compositional analysis of positive and comparative adjectives would equally well account for the crisp judgment data, or whether they make incorrect predictions. There is insufficient space to consider all alternative accounts of vagueness here, so I will instead discuss just a couple of the most central analyses, in an effort to illustrate how the explicit/implicit comparison data can be used to test different hypotheses.

4.3.2 Epistemicism

Let us begin by considering an alternative account of vagueness that maintains the assumptions about the composition of positive and comparative adjectival predicates that we made in the previous section: a gradable adjective denotes a measure function, and must combine with degree morphology to derive a property. At first glance, it appears that these compositional assumptions ensure that any account of vagueness that provides *some* way of deriving the Similarity Constraint will achieve the same results that we achieved with the interest-relative account, because the morphological difference between positives and comparatives (composition of the adjective with POS VS MORE) will always allow us to make the necessary distinction between the two forms. This is not the case, however; this result is ensured only if the analysis of vagueness is linked to a particular semantic feature of the positive form (i.e. to a feature of the meaning of POS). If, instead, the explanation of Similarity depends on some extralinguistic property that interacts with the meaning of the positive form but is not itself part of its meaning, then an explanation of crisp judgment effects in comparison is not guaranteed.

One approach to vagueness that works this way is Williamson's (1992, 1994) epistemic analysis. Williamson starts from the assumption that vague predicates (and in fact all predicates in natural language) sharply define a positive and negative extension. This is also true of the interest-relative account discussed above; the difference between the two analyses is that in Williamson's account, our judgments about how a vague predicate applies to objects that are very similar with respect to the relevant property (true of both or false of both) are not due to a semantic feature of the predicate (like interest relativity), but rather to a more general principle of epistemic uncertainty.

Specifically, Williamson observes that unlike many predicates, whose extensions may be stabilized by natural divisions (cf. Putnam 1975), the extensions of vague predicates cannot be so stabilized: a slight shift in

our disposition to say that Venus is big, for example, would slightly shift the extension of *big*. The boundary is sharp, but not fixed. But this in turn means that an object whose size puts it just below (or above) the threshold for counting as *big* could easily have been (or not been) counted as *big* had the facts (in particular, the linguistic facts) been slightly different – different in ways that are too complex for us to even fully catalogue, let alone compute. Given this instability, there will always be objects whose position on the relevant continuum make it such that we can never really know whether or not they are *big*. This last point leads to the “margin for error” principle in (17) (Williamson 1992: 161):

- (17) For a given way of measuring differences in measurements relevant to the application of property *P*, there will be a small but non zero constant *c* such that if *x* and *y* differ in those measurements by less than *c* and *x* is known to be *P*, then *y* is *P*.

The upshot of this reasoning is that it is impossible to know whether, for example, *big* applies differently to objects *x* and *y* that differ in size by less than *c*; this is the source of the Similarity Constraint on the epistemic analysis of vagueness.

While this analysis presents a plausible account of our reaction to the inductive premise of the sorites paradox (though see Fara 2000 for criticism), it runs into problems when we consider crisp judgment effects in implicit comparison. To be precise, it runs into problems given the assumption that the semantic content of the positive form is itself consistent with crisp judgments, unlike what we saw in the previous section.¹⁰ The simplest way to implement this assumption is to say that POS has the denotation in (18a), which is just like (14a) except that the interest-relative relation \succ_1 is replaced with a regular asymmetric ordering relation.


- (18) a. $\llbracket \text{POS} \rrbracket^c = \lambda g_{(e,d)} \lambda x. g(x) \succ \mathbf{std}(g)(c)$
 b. $\llbracket \text{POS big} \rrbracket^c = \lambda x. \mathbf{big}(x) \succ \mathbf{std}(\mathbf{big})(c)$

The **std** function returns a standard degree for the adjective in the context of utterance, so that the positive form of e.g. *big* denotes the property in (18b), which is true of an object as long as its size exceeds the standard (by any amount). More generally, as long as **std** is defined in the context for the adjective that POS composes with, the result will be a property that gives back true or false for any object in its domain strictly based on

how that object gets projected onto the scale: whether it is mapped to a degree that is ordered above or below the standard.

Let us now see how this analysis fares in crisp judgment cases, focusing on examples with definite descriptions. As shown by the example in (19), even when it is made clear both through preceding context and visual stimulus that there is a difference in size between two objects, if that difference is very slight, it is infelicitous to use a definite description based on an implicit comparison to refer to the larger of two objects, but it is felicitous to use the corresponding definite description based on an explicit comparison.

- (19) The planets Uranus and Neptune differ in size. This difference is represented schematically in the diagram on the right, where the circles represent the sizes of the two planets, drawn to scale. Uranus is the {#big, bigger} one.



The problem for the epistemic analysis of vagueness is that this is a situation in which the linguistic and nonlinguistic contextual factors are such that there should be enough certainty about the cutoff point between the big and the nonbig things to justify the use of *the big one* to uniquely describe Uranus. The argument runs as follows.

First, we know based on both prior discourse and observation of the graphical representations, that the planets under consideration have different sizes. Second, we know based on our knowledge of the meaning of the positive form, that the standard of comparison for *big* can vary in different contexts of utterance. (In fact, children as young as three years of age know this, as shown experimentally in Syrett et al. 2010.) Given our semantic assumptions, this means that **stnd** can return different values in different contexts. Finally, we know, based on our knowledge of the presuppositions introduced by the use of the definite article, that one and only one of the two objects under discussion should have a size that exceeds the standard for *big* in the context. We should therefore accommodate these presuppositions in the discourse in (19) by assuming that we are in a context in which the degree returned by the **stnd** function is one that is ordered below the size of Uranus, making *big one* true of it, and above the size of Neptune, making *big one* false of it. But if this is the case, then it should be perfectly acceptable to refer to Uranus as *the big one* in this context, contrary to fact.

The epistemicist might respond to this argument by saying that the implicit comparison form in (19) is infelicitous because the difference in size between Uranus and Neptune is less than the margin for error constant c . If this is the case, then accommodating a standard of comparison which will allow us to know that *big one* is true of Uranus in the context of utterance will, in virtue of (17), make the description true of Neptune, thereby violating the presuppositions of the definite article. Indeed, Williamson's explanation of the sorites paradox – and our judgments about it – relies on exactly this sort of reasoning; see Williamson (1992: 161).

This response cannot be correct, however. The reason that judgments about predications involving vague predicates are typically sensitive to a margin for error is because their extensions are, in general, unstable. As noted above, changes in our dispositions can result in changes in the extension of a vague predicate in ways that are too complicated to calculate, necessitating (or giving rise to) a principle like (17). However, thanks to the semantic/pragmatic contribution of the definite article, this instability disappears – or is at least significantly reduced – when a vague predicate is used to distinguish between two objects that differ, in an observable way, solely along the dimension of measurement that the predicate encodes, as is the case in (19). The definite article imposes uniqueness and existence requirements as a matter of meaning; the meaning of the positive form of *big* allows for a context-dependent standard of comparison which can (by hypothesis) make fine-grained distinctions in size; therefore the only kind of disposition that could lead to a degree of instability in the extension of *big* in contexts like (19) which would justify a margin for error larger than the size difference between the two planets would be a disposition to behave in a way that is inconsistent with the semantic/pragmatic requirements of the expressions of our language. Assuming that this sort of pathological disposition is not normally at play, the result is that the margin for error in contexts like (19) should be so small as to be irrelevant, leading to the incorrect prediction that implicit comparison should be acceptable.

We could salvage the epistemicist's response, however, by returning to our earlier assumption that the positive form (and in particular, the POS morpheme) has a richer meaning than the simple ordering relation in (18). In particular, if we assume that POS introduces whatever element of meaning is responsible for the kind of instability that gives rise to a margin for error in the first place, and that this instability persists even in the presence of the semantic/pragmatic contributions of the definite article, then the epistemicist's response goes through. For example, it

could be that instability in the extension of the positive form of *big* arises from interest relativity in the meaning of POS, as in the analysis we considered in the previous section: since our interests are constantly shifting, the extensions of the positive form are also shifting. If it can be shown that this instability persists even in the presence of the uniqueness and existence presuppositions of the definite article, then the epistemist's response can be maintained.¹¹ It is important to observe, though, that this version of the response crucially imparts an aspect of meaning to the positive form that is not present in the comparative. This means that crisp judgment effects – and by extension, the Similarity Constraint – are ultimately rooted in a semantic property of this class of vague predicates.

4.3.3 Underlying vagueness

Let us now consider supervaluationist accounts of vagueness. There is a way of implementing this sort of analysis that can in principle account for the difference between implicit vs explicit comparison in crisp judgment contexts: we could maintain the compositional assumptions about gradable adjective meaning and explicit comparative morphology that we adopted in section 4.3.1, but take no stand on whether the semantics of POS is interest relative, and instead just make the general assumption that regardless of its specific contributions to the meaning of the positive form, the result is a vague property that should be given a standard supervaluationist analysis. To the extent that such an analysis can be made to derive the Similarity Constraint, for example by stipulating a global constraint that disallows the sort of “fine-grained” precisifications that are necessary to make the positive form usable in crisp judgment contexts (see van Rooij (this volume); but see also Kamp (1975: 145), who explicitly resists this move, and Fara (2000), who rejects supervaluationism on other grounds), it will account for the facts.

These are not the compositional assumptions usually made in supervaluationist analyses that actually address the relation between positive and comparative gradable predicates, however. This brings us back to the second question about the implicit/explicit comparison distinction: what does it tell us about the compositional relation between these forms? To understand the significance of this question, we need to step back a bit and take note of a robust typological generalization about the world's languages: if the positive and comparative forms of a gradable predicate stand in a morphological markedness relation to each other, it is always the comparative that is the marked form. While there are many languages that do not make a morphological distinction between

positive and comparative predicates (as noted in section 4.2), there are no known languages in which the comparative form is morphologically simple and the positive is morphologically complex.¹²

The typological facts can certainly be accommodated in a degree-based theory of the sort we started this section with (e.g. by hypothesizing that the move from a one-place measure function to a one-place property is achieved by a type-shifting rule while the move to a two-place comparative relation requires extra morphology; see note 8). A major selling point of most supervaluationist analyses, along with their close cousins, the comparison-class based analyses of Wheeler (1972), Klein (1980), van Benthem (1982) and van Rooij (this volume), is that they actually predict the typological facts. This is because such analyses provide a way of compositionally deriving the meaning of comparative predicates from the (vague) meaning of the positive, and as an “added bonus,” they can achieve this result without introducing abstract objects like degrees into the semantic ontology.

For example, Kamp (1975) defines the comparative morpheme roughly as an operator that ranges over different ways of precisifying a vague predicate, and says that there are ways of doing so that make it true of the target of comparison and false of the standard, but not vice versa. For example, on this view, (20) is true relative to a model M just in case there is a model M' , which may differ from M in how it populates the positive and negative extension of *big* (and in whether it does so in a total or partial way) such that M' assigns Uranus to the positive extension of *big* and Neptune to its negative extension.

(20) Uranus is bigger than Neptune.

The Consistency Constraints (however they are derived) will ensure that this meaning gives the right truth conditions, i.e. that the size of Uranus exceeds the size of Neptune. But crucially, in order to account for the fact that (20) can be felicitously used in crisp judgment contexts, as we have seen, it must be the case that the range of models with respect to which the positive root can be evaluated includes ones that make very fine-grained distinctions (see Kamp 1975: 145).

Given this, it is difficult to see what rules out implicit comparison constructions in the same contexts. That is, if there is a way to precisify *big* in just the way that (20) requires, then surely if the linguistic and discourse context tells us that we are restricting our attention to just Uranus and Neptune, as is the case in implicit comparison, it is precisely this precisification that we will need to invoke. We could avoid this result by stipulating that the positive root is restricted to be evaluated only with

respect to coarse models – i.e. by putting a limit on precisification – but then we would incorrectly predict that explicit comparatives should also show crisp judgment effects, since they are compositionally derived from the positive. Alternatively, we could hypothesize that the positive contains an element of meaning that adds the restriction to coarse-grained models, but this would effectively amount to giving up on the hypothesis that the comparative is compositionally derived from the positive.

Turning to comparison class-based analyses of positive and comparative adjectives, it appears at first glance that they run into the same sorts of problems as supervaluation analyses. In this kind of approach, the positive form denotes a property whose extension is determined relative to a comparison class c , which may be either explicit or implicit. The resulting function from individuals to truth values may be partial – distinguishing between those elements in the comparison class that definitely have the property, those that do not, and those that fall in an “extension gap” – leading to a treatment of vagueness that is much the same as in a supervaluationist analysis.

Once we add in the Consistency Constraints, there are a couple of different options for deriving the comparative from the positive. Wheeler (1972) proposes that the function of the comparative is to stipulate that the adjective is evaluated relative to a comparison class consisting of the target and standard of comparison, and to assert that, relative to this comparison class, it is true of the target and false of the standard. This is shown in (21), where c is the comparison class parameter:

$$(21) \llbracket \text{MORE } A \text{ than } y \rrbracket^c = \lambda x. \llbracket A \rrbracket^{\{x,y\}}(x) = 1 \wedge \llbracket A \rrbracket^{\{x,y\}}(y) = 0$$

But if this is the right meaning, it is difficult to see how implicit and explicit comparatives can be distinguished in crisp judgment contexts. According to (21), the explicit comparative *Uranus is bigger than Neptune* is true just in case Uranus counts as big relative to the comparison class {Uranus, Neptune}, and Neptune does not. But then it is not at all clear how we can rule out the corresponding implicit comparative, which seems to be saying exactly the same thing.

Klein (1980) proposes a different semantics for comparatives, given in (22):

$$(22) \llbracket \text{MORE } A \text{ than } y \rrbracket^c = \lambda x. \exists c' [\llbracket A \rrbracket^{c'}(x) = 1 \wedge \llbracket A \rrbracket^{c'}(y) = 0]$$

This analysis shares with Wheeler’s the idea that the basic function of a comparative is to relativize the extension of the adjective to a comparison class that makes it true of the target and false of the standard. The

difference is that it does not stipulate what this comparison class is; it merely asserts that there is one.

This difference is crucial, and is exploited by Robert van Rooij in the comparison class-based analysis of explicit and implicit comparatives that he develops in his chapter in this volume. In essence, van Rooij proposes that explicit comparatives be given a Klein-style analysis, and implicit comparatives be given a Wheeler-style analysis, so that *compared to* structures have the meaning in (23) (definite descriptions involving implicit comparison will be the same in the relevant respects):

$$(23) \quad \llbracket A \text{ compared to } y \rrbracket^c = \lambda x. \llbracket A \rrbracket^{\{x,y\}}(x) = 1 \wedge \llbracket A \rrbracket^{\{x,y\}}(y) = 0$$

In van Rooij's analysis, the implicit comparative *Uranus is big compared to Neptune* has the truth conditions in (24a), and the explicit comparative *Uranus is bigger than Neptune* has the truth conditions in (24b):

$$(24) \quad \begin{array}{l} \text{a. } \llbracket \textit{big} \rrbracket^{\{\textit{uranus}, \textit{neptune}\}}(\textit{uranus}) = 1 \wedge \llbracket \textit{big} \rrbracket^{\{\textit{uranus}, \textit{neptune}\}}(\textit{neptune}) = 0 \\ \text{b. } \exists c' \llbracket \llbracket \textit{big} \rrbracket^c \rrbracket^c(\textit{uranus}) = 1 \wedge \llbracket \llbracket \textit{big} \rrbracket^c \rrbracket^c(\textit{neptune}) = 0 \end{array}$$

The truth conditions in (24a) lead to unacceptability in crisp judgment contexts, according to van Rooij, because of whatever principles are ultimately determined to provide the best way to derive the Similarity Constraint and explain judgments about the inductive premise of the sorites. Van Rooij discusses a couple of options in his chapter, which I will not evaluate here; the upshot is that comparison classes that consist only of two objects that deviate by a very small degree along the compared dimension are ruled out, so the infelicity of implicit comparison in crisp judgment contexts can be viewed as a kind of presupposition failure.

The reason that (24b) does not run into the same problem is because it introduces existential quantification over comparison classes. (24a) is unacceptable because it necessarily involves the inadmissible comparison class {Uranus, Neptune}; (24b) is acceptable because it is not restricted to this comparison class, and it is true as long as there is some other, admissible comparison class relative to which *big* is true of Uranus and false of Neptune. If there is such a class, then given the Consistency Constraints, (24b) entails that the size of Uranus exceeds the size of Neptune, which is what we want.

An important feature of van Rooij's analysis is that the constraints that explain the infelicity of implicit comparison in crisp judgment contexts

(whatever they turn out to be) are general constraints on comparison classes and the meaning of the positive form, and so do in fact carry over fully to explicit comparatives. The reason explicit comparatives do not run into the same problems is because of the “extra” meaning they introduce: existential quantification over comparison classes. Van Rooij thus avoids weakening the hypothesis that the comparative is fully derived from the positive, as we saw above was the case for the supervaluationist analysis, and so provides what looks like the best candidate for a semantic analysis of the comparative in which it is fully compositionally derived from the positive.

Van Rooij’s analysis is not without its own potentially problematic features, however. In particular, the analysis necessitates assumptions about the domains of vague predicates that call into question the foundational hypothesis that the semantics of vague predicates can be captured strictly on the basis of orderings on individuals, rather than in terms of more abstract objects like degrees. Specifically, as van Rooij observes, his analysis predicts that an assertion of an explicit comparative in a crisp judgment context should entail (or maybe pragmatically presuppose) the existence of “witness” objects that, together with the compared objects, can be used to construct an appropriate comparison class. Van Rooij formalizes this condition as a constraint on models which requires that for any two objects x, y that differ by a small degree along a compared dimension δ , there should be objects w, z such that z is significantly different from y relative to δ but indistinguishable from x , and w is significantly different from y and indistinguishable from x .¹³

The problem is that it does not appear that assertions of explicit comparatives actually entail (or presuppose) the existence of such witnesses. Consider, for example, an assertion of *Uranus is bigger than Neptune*, evaluated against a common ground in which we know that the sizes of the eight planets in the solar system (plus Pluto, for nostalgia’s sake) are as shown in (25):

(25)	Jupiter	142,984 km	Earth	12,756 km	(Pluto 2,274 km)
	Saturn	120,536 km	Venus	12,104 km	
	Uranus	51,118 km	Mars	6,794 km	
	Neptune	49,532 km	Mercury	4,880 km	

This sentence is clearly true, but does not appear to entail (or presuppose) the existence of objects (planets or otherwise) with sizes just a bit bigger than Uranus or smaller than Neptune, in this solar system or any




Figure 4.3 Saturn, Uranus, Neptune and Earth

other. Even if we explicitly hypothesize the absence of such objects, the sentence is both felicitous and true:

- (26) Let us consider a universe that consists of just the nine planets listed in the table in (25). As the table shows, in this universe, *Uranus is bigger than Neptune*, just as it is in our universe.

One response might be to claim that Saturn and Earth count as the relevant witnesses in this example. But this would mean that Saturn is indistinguishable from Uranus relative to size but significantly different from Neptune, and that Earth is indistinguishable from Neptune but significantly different from Uranus. This is clearly wrong, as illustrated by Figure 4.3, which provides a graphical representation of the relative sizes of these four planets (minus Saturn's rings) in descending order.

Evidently the "witnesses" that are required in order to explain the implicit/explicit distinction, and that are required to explain our judgments about the truth and felicity of the comparative in (25) and (26), are rather abstract, and need not correspond to actual objects in the world that the sentence is about (or the model that represents it). In fact, as van Rooij points out, the structure of the model that results from the addition of a witness constraint is one that bears a distinct similarity to the abstract structures assumed in many degree-based theories. While this might provide a technical solution to the implicit/explicit distinction that allows us to maintain a semantics in which the comparative is derived from the positive, it clearly undermines one of the main selling points of the comparison class approach, which is that it can base a semantics of gradability and comparison entirely on how the (actual)

objects in the domain of a gradable predicate relate to each other, without reference to more abstract objects like degrees. Whether it more seriously undermines the analysis is a question that deserves further consideration.

4.4 Conclusion

Stassen (1985: 24) states that “a construction in natural language counts as a comparative construction if that construction has the semantic function of assigning a graded (i.e. non-identical) position on a predicative scale to two (possibly complex) objects.” Both implicit and explicit comparisons count as comparative constructions by this definition, yet as we have seen in this chapter, they differ in acceptability when used to characterize orderings between objects that differ by a small but observable amount along the relevant parameter, indicating that the two constructions have slightly different meanings. This point is significant for typological work on the grammar of comparison, but, as I have argued in this chapter, it is also relevant to questions about vagueness. I have shown that the contrast between implicit and explicit comparison follows directly from an analysis in which positive and comparative forms of gradable predicates are both derived from a more abstract source, and in which vagueness arises from a semantic feature of the positive, though the facts do not necessarily tell us exactly what the nature of this feature is (i.e. whether it comes from interest relativity or from something else). The facts are not so easily explained in a theory in which vagueness is a purely epistemic or model-theoretic phenomenon, or one in which the comparative is compositionally derived from the positive.

Notes

* Thanks to Paul Égré and Nathan Klinedinst for helpful comments on the first draft of this chapter, and to Robert van Rooij for extensive and thought-provoking discussion of explicit vs implicit comparison. This chapter is based upon work supported by the National Science Foundation under Grant No. BCS-0620247.

1. There are, of course, many other ways to talk about the size relation represented in Figure 4.1, such as sentences containing other degree morphemes (e.g. *Venus is not as big as Uranus*), collocations in which the compared dimension is provided by a noun or nominalized adjective and the ordering relation comes from a verb (*Uranus exceeds/surpasses Venus in size*), or constructions using phrases like *relative to*, *comparatively speaking* and so forth. Some of these constructions have their own unique properties, but for the most part they can be grouped either with the sentences in (3) or with those in (4) with respect to the kinds of judgments I will discuss below. Therefore, in order to keep the discussion in this chapter focused, I will limit my discussion to examples parallel to those in (3) and (4).

2. How exactly these constraints are derived is a point of variation among analyses. See Kennedy (1999) for discussion, and see Robert van Rooij's chapter in this volume for a detailed formalization of one way of deriving them.
3. Evidence that this construction is functionally a comparative, and not an arbitrary coordination of independent propositions, comes from the fact that its use need not entail that *Aquitania* is small. As noted by Marsack (1975: 66), "Even in the case of giant vessels like the *Queen Mary* and the *Aquitania* this construction would be used. To indicate that the 85,000 ton *Queen Mary* is bigger than the 45,000 ton *Aquitania*, a Samoan of the old school would say [8]." (When he says "Samoan of the old school," Marsack alludes to the fact that constructions like (8) are no longer in regular use, and have been replaced in modern colloquial Samoan by what is clearly a type of explicit comparison structure.)
4. van Rooij (this volume) claims that implicit comparisons are false in crisp judgment contexts. My own judgment about the truth or falsity of the examples in (10) in the context of Figure 4.2 is not so clear, in contrast to my judgment of their (un)acceptability. Given that there is a natural pragmatic explanation for the facts, as described in the text, I prefer to characterize the examples in (10) as infelicitous rather than false.
5. This is not to say that absolute gradable predicates are always acceptable in implicit comparison; in fact, they show a different set of restrictions, which can be traced to the fact that they make use of fixed standards of comparison. See Kennedy (2007a) and Sawada (2009) for discussion.
6. Paul Égré asks whether the two uses of *old* discussed might not involve distinct senses after all, but rather a single sense, with the difference in judgments arising from the interaction of the adjective and noun meanings: while a two-day difference in age is unlikely to be significant when considering boys, a very small difference might very well be significant in the case of computer files. To some extent, I think this is correct: I have argued in Kennedy (2007b) that relative and absolute gradable predicates denote the same kinds of properties (interest relative or otherwise), and that it is their scalar features that determine whether they define sharp or fuzzy boundaries (i.e. whether they are vague or not). But even given that, I believe that there is reason to believe that two uses of *old* discussed here actually measure different (though often related) properties: *old* in (11) measures "absolute age"; *old* in (12) measures (something like) "distance from a salient transition point." Evidence that these senses are distinct comes from the fact that they have different antonyms – *young* for the former, and *new* for the latter – and can be put in opposition without contradiction. For example, I could point at two women of different ages at a faculty reception and felicitously assert (ia) to indicate that the young woman has been serving as a department chair for a longer period of time than the old woman. If we replace *new* with *young*, however, as in (ib), the sentence sounds contradictory:
 - (i) a. The old woman is a new department chair and the young woman is an old department chair.
 - b. #The old woman is a young department chair and the young woman is an old department chair.

(ia) is fine because the adjectives that modify the two occurrences of *woman* involve the absolute age sense, while the ones that modify the two occurrences of *department chair* involve the “distance from a transition point” sense, and so imply nothing about the actual age of the two women. In (ib), however, the preferred parse (because of the parallel syntactic structure) is one in which both sets of adjectives involve the absolute age sense. Given that an old woman is older in absolute age than a young woman, and a young department chair is younger in absolute age than an old department chair, the preferred parse of (ib) is one that entails that the first woman is both older and younger than the second woman, which is a contradiction.

7. A variant of this kind of analysis is one in which gradable predicates denote relations between degrees and individuals, so that e.g. *big* is true of a pair $\langle d, x \rangle$ just in case x 's size equals (or at least equals) d (see e.g. Cresswell 1977, von Stechow 1984, Heim 2000, Schwarzschild 2005). As in the functional analysis I assume here, the relational analysis requires the adjectival root to combine with some sort of degree morphology or undergo some sort of type-shifting rule in order to saturate the degree argument and derive a property of individuals. See Bogal-Allbritten (2008) for arguments based on morphosyntactic data in Navajo that favor a functional analysis.
8. An alternative to hypothesizing a null, positive degree morpheme would be to assume a type-shifting rule that achieves the same semantic results (Kennedy 2007b; cf. Chierchia 1998 for similar points in the domain of nominal meaning, and general discussion of the significance of the morphology/type-shifting distinction). Grano (2010) presents arguments based on the morphosyntax of the positive form in Mandarin Chinese that the type-shifting option may in fact be preferable.
9. In fact, I advocate a slightly different approach in Kennedy (2007b), in which the positive form denotes a property that is true of an object just in case it “stands out” relative to the kind of measurement that the adjective encodes.
10. This is not a necessary assumption: the claim that we cannot know the precise location of the boundary between the positive and negative extension of a vague predicate is consistent with an interest-relative semantics, as Fara (2000) points out. However, if the explanatory force of the epistemic analysis of vagueness lies in its ability to maintain a “sharp” semantics for vague predicates and explain our judgments about particular uses of them based on general principles of epistemic uncertainty, then we should evaluate it relative to a semantics for the positive form that is consistent with crisp judgments.
11. On the other hand, it is not clear that we need to invoke a margin for error at all, if we adopt an interest-relative semantics. The interest-relative account of the infelicity of the definite description in (19) goes as follows. We have a standing interest in behaving in a way consistent with the conventional meanings of the expressions of our language. This interest should lead us to fix the extension of *big* in a way that makes it true of Uranus and false of Neptune, in order to accommodate the presuppositions of the definite article. However, this very same interest also requires us to accept the entailments of the positive form of *big*, which include the strong position that Uranus but not Neptune has a size that *significantly* exceeds the contextual standard of comparison, rather than the weaker position that Uranus but not Neptune has a size that (merely) exceeds the standard. If our overall interests are in

general incompatible with this entailment, as claimed by Fara, we have an irresolvable conflict, leading to anomaly.

12. It has sometimes been claimed that Mandarin Chinese has the opposite markedness relation, given contrasts like the one in (i) (see e.g. Sybesma 1999: 27):

- (i) a. *Zhangsan gao.*
Zhangsan tall
'Zhangsan is taller than some contextually salient individual.'
'*Zhangsan is tall.'
- b. *Zhangsan hen gao.*
Zhangsan HEN tall
'Zhangsan is tall.'

However, as shown in recent work by Liu (2009) and Grano (2010), the contrast in (i) is really a fact about nonnegated, matrix assertions, and in other syntactic contexts, unmarked gradable predicates can have positive form meanings. More convincing evidence for morphological marking of the positive form comes from Elizabeth Bogal-Allbritten's (2008) work on Navajo, though in this language *both* the positive and comparative forms are derived by combining a root (which cannot appear by itself) with degree morphology. This is exactly the kind of pattern that the degree-based theory outlined at the beginning of this section predicts to be a possibility.

13. There are two exceptions: objects that are maximal and minimal with respect to δ , if they exist.

References

- Bartsch, R. and Vennemann, T. (1973). *Semantic Structures: a Study in the Relation between Syntax and Semantics*. Athäenum Verlag, Frankfurt.
- Bogal-Allbritten, E. (2008). Gradability and degree constructions in Navajo. BA Honors thesis, Swarthmore College.
- Bogusławski, A. (1975). Measures are measures: in defence of the diversity of comparatives and positives. *Linguistische Berichte*, 36: 1–9.
- Chierchia, G. (1998). Reference to kinds across languages. *Natural Language Semantics*, 6: 339–405.
- Cresswell, M. J. (1977). The semantics of degree. In Partee, B. (ed.), *Montague Grammar*. pp. 261–92. Academic Press, New York.
- Fara, D. G. (2000). Shifting sands: an interest-relative theory of vagueness. *Philosophical Topics*, 20: 45–81.
- Fara, D. G. (2008). Profiling interest relativity. *Analysis*, 68: 326–35.
- Grano, T. (2010). Mandarin *hen*, universal markedness, and tense. Unpublished MS, University of Chicago.
- Heim, I. (2000). Degree operators and scope. In Jackson, B. and Matthews, T. (eds), *Semantics and Linguistic Theory 10*, pp. 40–64. CLC Publications, Ithaca, NY.
- Kamp, H. (1975). Two theories of adjectives. In Keenan, E. (ed.), *Formal Semantics of Natural Language*, pp. 123–55. Cambridge University Press, Cambridge. Reprinted in 1984.

- Kamp, H. and Partee, B. (1995). Prototype theory and compositionality. *Cognition*, 57: 129–91.
- Kennedy, C. (1999). *Projecting the Adjective: the Syntax and Semantics of Gradability and Comparison*. Garland, New York. (1997 UCSC PhD dissertation).
- Kennedy, C. (2007a). Modes of comparison. In Elliott, M., Kirby, J., Sawada, O., Staraki, E. and Yoon, S. (eds), *Papers from the 43rd Annual Meeting of the Chicago Linguistic Society*. Volume 1: *The Main Session*, pp. 139–63. Chicago Linguistic Society, Chicago, Ill.
- Kennedy, C. (2007b). Vagueness and grammar: the semantics of relative and absolute gradable predicates. *Linguistics and Philosophy*, 30: 1–45.
- Kennedy, C. and McNally, L. (2005). Scale structure and the semantic typology of gradable predicates. *Language*, 81: 345–81.
- Klein, E. (1980). A semantics for positive and comparative adjectives. *Linguistics and Philosophy*, 4: 1–45.
- Liu, C.-S. L. (2009). The positive morpheme in Chinese and the adjectival structure. *Lingua* DOI: 10.1016/j.lingua.2009.06.001.
- Marsack, C. C. (1975). *Samoan*. English Universities Press, New York.
- Putnam, H. (1975). The meaning of ‘meaning’. In Gunderson, K. (ed.), *Language, Mind, and Knowledge*, pp. 131–93. University of Minnesota Press.
- Raffman, D. (1994). Vagueness without paradox. *The Philosophical Review*, 103: 41–74.
- Raffman, D. (1996). Vagueness and context relativity. *Philosophical Studies*, 81: 175–192.
- Richard, M. (2004). Contextualism and relativism. *Philosophical Studies*, 119: 215–42.
- Sapir, E. (1944). Grading: a study in semantics. *Philosophy of Science*, 11: 93–116.
- Sawada, O. 2009. Pragmatic aspects of implicit comparison: an economy based approach. *Journal of Pragmatics*, 41: 1079–103.
- Schwarzschild, R. (2005). Measure phrases as modifiers of adjectives. *Recherches Linguistiques de Vincennes*, 35: 207–28.
- Soames, S. (1999). *Understanding Truth*. Oxford University Press, Oxford.
- Stanley, J. (2003). Context, interest relativity, and the sorites. *Analysis*, 63(4): 269–80.
- Stassen, L. (1985). *Comparison and Universal Grammar*. Basil Blackwell, Oxford.
- Stechow, A. von. (1984). Comparing semantic theories of comparison. *Journal of Semantics*, 3: 1–77.
- Sybesma, R. (1999). *The Mandarin VP*. Kluwer, Boston.
- Syrett, K., Kennedy, C. and Lidz, J. (2010). Meaning and context in children’s understanding of gradable adjectives. *Journal of Semantics*, 27: 1–35.
- Unger, P. (1975). *Ignorance*. Clarendon Press, Oxford.
- van Benthem, J. (1982). Later than late: On the logical origin of the temporal order. *Pacific Philosophical Quarterly*, 63: 193–203.
- Wheeler, S. (1972). Attributives and their modifiers. *Noûs*, 6: 310–34.
- Williamson, T. (1992). Vagueness and ignorance. *Proceedings of the Aristotelian Society*, 66: 145–62.
- Williamson, T. (1994). *Vagueness*. Routledge, London.

5

The Inhabitants of Vagueness Models

Galit W. Sassoon

ILLC/University of Amsterdam

5.1 Introduction

As part of the common practice of experimental research in cognitive psychology, objects are systematically represented by clusters of property values. Individuals correspond to points in an n -dimensional space, for some number n , where each dimension (axis) is some scalar property. For example, the set of possible individuals presented as stimuli in a given experiment can be represented with the two-dimensional space generated from the scalar properties denoted by *red* and *long*, or from the dimensions *color* and *shape*, the latter seen as nominal-scale properties, assigning to entities values such as “red,” “blue,” “square” and “circle.” The result is a set of individuals including a red square, a blue square, etc. Many other examples can be found in, for instance, Murphy (2002).

The ontology underlying these experimental designs consists of different types of instances:

- (1) a. A domain of *time points*, D_p
- b. A domain of *individuals*, D_x (including “stages,” i.e. individuals at a given time p)
- c. A domain of *property values* $D_r \subseteq D_x$ (“satisfaction extents”): measurable subparts of individuals – their heights, lengths, weights, colors, happiness extents, etc.¹
- d. A domain of *graded properties*, $D_{(x,r)}$ (mappings of individuals to values)

Now let us consider the ontology underlying language and thought more broadly, i.e. the ontology which is the core of the interpretation

of natural language expressions. Members of the domain of possible individuals within this ontology can also be characterized as points in an n -dimensional space, with intrinsic properties such as those listed in (1c, d) above forming the generating coordinates (height, length, temperature, color, happiness extents and so on). On this proposal the domain D_x is a theoretical construct, consisting of all the possible combinations of property values which we humans can conceive of. Hence, the domain consists of “possible” individuals. Upon encountering actual, real-world entities, we attempt to create a link between them and the corresponding elements of D_x ; but no empirical observation of an actual, real-world entity can pinpoint one single individual in D_x , because only partial, proper subsets of the property values of worldly entities are accessible to us. Thus, observations only isolate a subset of the domain consisting of all those possible individuals that, for all that we know, might still be the observed entity, i.e. all those that share with the observed entity all the property values accessible to us.

This chapter explores the predictive fruitfulness of this hypothesis, namely the hypothesis that:

- (2) The elements of the domain of possible individuals D_x are identified by and only by their property values.

Methodologically speaking, this study purports to contribute to the general enterprise of bridging the gap between researches in linguistics and psychology, by aiming towards a unified conception of individuals. However, more significantly, this study purports to shed new light on central problems much discussed in the philosophical literature, problems pertaining to cross-world identity and to vagueness versus ignorance.

In a nutshell, an important consequence of the proposal that individuals are identified by the sum of their intrinsic property values is that it directly predicts a certain range of our ontological or metaphysical intuitions concerning cross-world identity.

On the one hand, intuitively, an individual who is, for instance, six feet tall in this world is not the same person in an alternate world (bound to the same time) in which his height is different. This intuition directly follows. The property values of an individual per a time point are index invariant (they do not vary across worlds), for one individual in D_x cannot be identified with two different sums of property values.

On the other hand, as discussed in section 5.2, if the property values of two individuals are the same across alternative worlds (bound to the same

time) the individuals are the same. They are the same regardless of how these property values categorize under vague predicates, for example, even if six feet tall counts as *tall*, in this world, but not in the alternate realities, and so the given individual is *tall* in this world and is *not tall* in the alternate realities. This can be viewed as an instance of Leibniz's law; two individuals with all the same property values must be identical. The crux of this proposal pertains to vague predicates; membership in their denotations plays no role at all in distinguishing between individuals (Sassoon 2007, Frank Veltman, p.c.). It follows that their only use is to blur differences between many different property values so as to form secondary, coarse-grained and therefore general categories, for example the whole dense set of possible heights can collapse into general categories such as "tall," "short," and "borderline."

As discussed in section 5.3, the proposal in (2) implies a departure from the Kripkean perspective on the reference of singular terms in that, according to (2), the referent of a noun like *Aristotle* is uniquely specified in a context *c* if and only if Aristotle's sum of property values is entirely specified in *c*. This never happens. Speakers, as well as language communities, never have access to the entire set of property values of a proper name's referent. For example, if all except the referent's height is specified, many possible individuals, exactly alike except differing in height, might form the unique referent of, for example, *Aristotle*. Thus, the referent varies across the worlds consistent with *c*.

A benefit of this departure concerns how to represent truth value gaps caused by vagueness versus those caused by accidental ignorance within a supervaluationist semantics. In practice, only a subset of the property values of the referents of singular terms are known in actual contexts and therefore we do not have a unique individual to refer to, but instead have sets of possible individuals. Consequently, terms may create truth value gaps in otherwise nonvague statements like, for instance, *Aristotle is taller than Plato*. This is a clear case of ignorance about property values, for we do not know the heights of these famous people. While the ordering of individuals by their heights is fixed (world-invariant), we do not know exactly which possible individuals in D_x are being compared. For all that we know, many possible individuals exactly alike except in height may form the unique referent of, for example, *Aristotle*.

Another source of truth value gaps is the absence of linguistic conventions regarding the spatial borders of referents of terms. Consider, for example, statements of the form *Mount Etna occupies n square meters* (cf. Lewis 1988, 1993). While the size of an individual is world-invariant, *Mount Etna* may refer to many possible partially overlapping regions; thus, no convention tells us precisely which individual in D_x is to

be measured. Finally, in some cases, when the existence of a linguistic convention uniquely specifying the referent is debatable, the question remains open, whether personal ignorance or semantic indeterminacy is the case. Different conventions for defining the region of a given city or state x may be held by people with equal authority regarding the interpretation of x in the given language. Different conventions may prevail in court, in novels, and in other daily life situations. On the present proposal, vagueness and accidental ignorance are represented in one and the same model, as discussed in section 5.3. This helps deal with the given examples. At the same time, the differences between the two phenomena, when such exist, are captured.

The main consequences of the present study are therefore tightly related to vagueness. Let us call models for the interpretation of vague predicates *vagueness models*, including in particular the kind of models used by *supervaluationist* theories (to be described shortly). Vague predicates (like *tall*, *heavy*, and *happy*) are characterized by the absence of sharp boundaries. Some individuals exist, for which we cannot tell whether they fall under the predicate or not. They form a denotation gap. Vagueness models, then, include *partial information states* (“contexts”) c , representing the common knowledge of possible communities of speakers (Stalnaker 1978). The interpretation of linguistic expressions in contexts c is modeled via a set of indices T_c , the worlds (Stalnaker 1978) or completions (also called *precisifications*; van Fraassen 1969; Fine 1975; Kamp 1975) consistent with c . Worlds are alternative universes or realities, while completions are information states, which are “classical” in the sense that every statement is either true or false. Completions may include more information than worlds do; for instance, each completion determines cutoff points for vague predicates, although this type of information may well not be part of the actual world (Kamp 1975). This chapter refers mainly to completions, but where differences between completions and worlds are relevant, they are explicitly discussed. Truth of a statement S in c is defined based on these indices:

- (3) a. S is true in c if and only if S is true in every $t \in T_c$;
 b. S is false in c if and only if S is false in every $t \in T_c$, and
 c. S 's truth value is undetermined otherwise.

Vague predicates are associated with graded properties (Kennedy 1999). Intuitively, individuals can satisfy these properties to different extents, and they fall under the predicates if and only if their satisfaction extents (height, health, happiness extent, etc.) exceed the predicates' contextual cutoff point (“*standard for membership*”). Though satisfaction extents

usually vary with time, this chapter does not aim to deal with issues specifically pertaining to changes over time. Thus, reference to time indices is omitted.

The graded property of some predicates, especially multidimensional ones, is indeterminate, i.e. completion-dependent (Klein 1980). This is manifested by indeterminacy with regard to the interpretation of the derived comparative of the adjective in question. For example, unlike the interpretation of the comparative *taller*, the interpretation of, for example, *tastier* or *more similar to Dan* is indeterminate. A predicate like *tall* always assigns a given individual, x_1 , the same value per a time point (x_1 's height). By contrast, a predicate like *similar to Dan* may assign x_1 different values in different completions. Crucially, the reason for this is not that x_1 changes from one completion to another, for it does not. Rather, it is the graded property associated with *similar to Dan* that changes. The value of x_1 in *similar to Dan* in a given completion depends on what parts of x_1 (which of x_1 's property values) are taken into account in similarity comparisons (general look, hair color, facial expressions, character, etc.) Hence, formally:

- (4) a. Let f be a function associating each $t \in T_c$ and predicate P with a *graded property*, $f(P, t)$, i.e. a function from individuals $x \in D_x$ to values $r \in D_r$.
- b. P is true of an individual $x \in D_x$ in t if and only if x 's value in P , $f(P, t)(x)$, exceeds P 's standard in t .
- c. x 's value in P is determined in c if and only if it is completion-invariant: $\forall t_1, t_2 \in T_c, f(P, t_1)(x) = f(P, t_2)(x)$.

For example, two completions t_1 and t_2 may differ along the graded property of a predicate like *beautiful*, $f(\text{beautiful}, t_1) \neq f(\text{beautiful}, t_2)$, representing different interpretations. The values of individuals in *beautiful* are determined (index-invariant) if and only if the context uniquely selects one possible interpretation (one measure of beauty) for this domain of individuals. Again, the individuals' sums of property values remain fixed across completions, for, as stated in (1c), property values are simply subparts of individuals; two individuals in two different worlds cannot have different subparts yet be one and the same (cf. Lewis 1986). It is the interpretation (graded property) of *beautiful* that varies, taking different aspects of an individual into account across different completions.

These notions generalize to expressions other than one-place predicates. For example, a context c may fail to determine whether some

individuals, x_1 and x_2 , stand in a relation like *love*, perhaps because they do not possess much love for one another, or perhaps because in c we do not yet know what *love* is – that is, which property values the graded property of *love* represents. Then in c the pair $\langle x_1, x_2 \rangle$ falls in the gap of the two-place predicate *love*. Thus, in a vagueness model, a vague predicate's graded property, as well as its denotation (set of instances) may vary across indices of evaluation.

To recapitulate, this chapter studies the consequences of the proposal that individuals are identified by and only by their property values. Section 5.2 explores our intuitions concerning whether, in different cases, individuals inhabiting different indices can be identified. Section 5.3 aims to capture the distinction between vagueness and accidental ignorance while representing the two phenomena within one and the same model. Section 5.4 briefly addresses additional implications of the model developed in this chapter, pertaining to demonstratives, identity statements and the interaction of terms with modals.

5.2 Cross-completion identity

Following Kripke (1980), linguists usually assume that a proper name like *Dan* rigidly designates an index-invariant individual, say – x_1 . This practice is based on the (perhaps wrong, but highly intuitive) idea that individuals may be identified across different indices. We can imagine an alternative reality in which Dan is taller than he is in this one. Why? At first we say things like *Let's call our baby "Dan,"* thereby associating the name with a particular individual; then this usage is passed on. Dan grows, his height changes, and yet it is intuitively the same Dan. Consequently, statements about what might or might not have been true of a name like *Dan* (unlike a definite description) are always thought to be about *that particular person* in various contexts, rather than about whoever would have happened to be called *Dan* in those contexts.

In opposition to that view, Lewis (1986) extensively argues that a name like *Dan* denotes different individuals in different indices, with these individuals being bound to each other by a counterpart relation. The intuition is that if in one index t_1 the proper name *Dan* denotes an individual x_1 who is 1.87 meters tall, and in t_2 it denotes an individual x_2 who is 1.86 meters tall, it does not make sense to say that these are the same individual. Intuitively, individuals differing in some of their property values (e.g. their height) cannot be considered the same. If the extension of *Dan* consists of two different individuals in t_1 and t_2 in T_c , it is undetermined in c , even if the two are bound by a counterpart

relation. Thus, absence of information about Dan's height is absence of information about the individual to which *Dan* refers.

This intuition becomes more obvious when entities other than persons are considered. For example, if it is unsettled whether the name *Princeton* denotes just the borough or any one of countless somewhat larger regions, the extension of *Princeton* consists of many different individuals in different completions (Lewis 1988). Likewise, if it is unsettled whether the referent of *John's house* includes the garage or not, *John's house* denotes different individuals in different completions. Notice also that a single cat may stand alone on a mat, although the cat minus one hair is also a cat. So, are there many cats on the mat? No; the definite description *The cat on the mat* denotes many different cats in different completions, yet a singular cat in each one (Lewis 1993). And if in *c* it is unsettled which, say, half-falling hairs are part of a given cat at a given time, the cat's name, e.g. *Tibbles*, denotes several different cats in different completions, yet a singular cat in each one (Lewis 1993).

Lewis (1986) views individuals and alternative realities as mereological sums, thus not objecting to the idea that two worlds (or completions, for that matter) may share a part (an individual) in common, in just the same way that Siamese twins may share a hand. But, of course, the hand that two twins have in common cannot have a different number of fingers for each twin. Similarly, the common part of two worlds or completions (a given individual) cannot have different properties (e.g. heights) in them. Likewise, if an individual's temporal stages differ by height, they are not one and the same. The intuition that cross-index identity between individuals differing in property values is implausible directly follows from a representation of individuals as the sums of their property values.

Moreover, the converse intuition follows for all other circumstances, i.e. it is harder to claim that individuals with identical property values are "different" in different indices.

Importantly, Lewis distinguishes between *intrinsic* properties (e.g. height, shape and number of fingers), which depend on the way an individual *x* itself is in an alternative *t*, and *extrinsic* ("relational") properties, which depend on *x*'s relations to others in *t*. If, as this chapter suggests, individuals are the sum of their property values, property values are obviously intrinsic. Conversely, the properties denoted by vague predicates like, for example, *tall* are relational. Whether *x* is tall or not is a matter of *x*'s relations to those heights representing the cutoff point between *tall* and *not-tall*. This proposal has the following intuitive consequences. Let us suppose that in two completions the proper name *Dan* denotes an individual who is 1.87 meters tall, but in one completion this individual

is considered *tall* and in the other *not-tall*, and these individuals are identical in all other respects. The intuition now is that we definitely can say that this is the same individual. It is only our interpretation of the word *tall* that has changed.

The crucial difference between this case and one in which the individuals' heights vary between two completions lies in whether intrinsic properties are being compared across completions, or only the ways in which we categorize these properties (as exceeding the cutoff points of vague predicates or not). Therefore, according to the present account, two individuals in two completions are the same if and only if they have the same length, width, color, intelligence, and so on and so forth, and regardless of whether or not these property values are sufficient for them to count as *long*, *wide*, *red*, *intelligent*, etc. in these completions.

In fact, Lewis's main objection against cross-world identity regards accidental intrinsic properties, i.e. properties that tend to vary across alternatives and time points (e.g. height and shape; Lewis 1986, section 4.2). But to the extent that these can be kept equal across two alternatives (or time points), the two alternatives may share parts.

Arguably, completions may differ only with respect to the standards of vague predicates like *tall*. If speakers' interests and goals vary, the standard may vary, even if the set of heights under discussion (the comparison class) remains fixed (Fara 2000: 57–9). Equal differences in height may be seen as salient (potential cutoff points) given one goal, and not salient given another goal. Hence, we have some leeway in our choice of standards for vague predicates, considering a fixed set of heights.

If two completions differ only with respect to the standards of vague predicates, these two completions share all the individuals. All property values of all individuals are identical in the two completions, while the status of individuals in standard-based properties varies, for example, the heights forming the clear cases of *tall* and the pairs of heights saliently similar or dissimilar (i.e. the cutoff point of *tall*) vary. Completions differing in this way can represent different speakers' goals.

If cutoff points of vague predicates are not part of the actual world (or worlds in general), worlds cannot differ along predicates' standards. But, surely, completions contain information about cutoff points. That cutoff points affect interpretation is assumed in linguistics across the board, whether these are parts of worlds or only of completions. Thus, different indices of evaluations (at least completions) may have individuals in common.

One of the advantages of a representation of individuals by their property values is that it captures these intuitive distinctions. Individuals

differing in height across completions are different, while individuals not differing in any property value are one and the same individual. An additional advantage is discussed in the following section.

5.3 Two sources for truth value gaps – vagueness versus ignorance

The adjectives in (5) below are all vague. As a result, even speakers who are familiar with all the relevant facts (Dan's height, Sam's IQ, etc.), may not be able to determine the truth values of the statements in (5a–c). Typically, supervaluationist theories distinguish between *vagueness* and *ignorance*, arguing that these statements lack a truth value. They are neither true nor false, so there is nothing there to know or to be ignorant of. Individuals fail to be in the positive or negative denotation of, for example, *tall*, when they qualify neither clearly as tall, nor clearly as nontall, not because it is not *known* whether they are tall or not (Kamp 1975, Kamp and Partee 1995: 148). Conversely, as the predicates in the statements in (6) are nonvague, they have precise application conditions and consequently determined truth values. Speakers who are not familiar with all the relevant facts (e.g. Dan's height) do not *know* these truth values. They are, therefore, in a state of ignorance.²

- (5) a. Dan is tall.
 b. Sam is intelligent.
 c. Mary is happy.
- (6) a. Dan is two meters tall.
 b. Dan is taller than Sam is.
 c. Dan is twice as tall as Sam.

Yet the same formal means, supervaluations, are used both for the representation of vagueness (Kamp 1975) and epistemic ignorance (Veltman 1984). Thus, the two phenomena are often not formally distinguished. The intuitive distinction between them is also under debate. According to the epistemic approach (Sorensen 1988, Williamson 1994, Fara 2000), vagueness is a purely epistemic phenomenon, consisting of the absence of knowledge or possibility to gain knowledge concerning denotation membership, despite the existence of a true answer. For every individual, it is determined whether it is tall or not, but the cutoff point is “unknowable,” so, still, for some individuals we cannot tell whether as a matter of fact they are tall or not (Williamson 1994).

Rivals of this approach doubt the plausibility of an inherent impossibility of knowing what the standards are. If they are out there in the world, why can we not discover them? Intuitively, the truth value gaps related to the statements in (5) and those related to (6) stem from different sources. A representation of individuals by their property values allows a representation of these two sources, thus capturing the intuitive distinction between vagueness and accidental ignorance, while representing both in the same formal model.

Consider, for example, speakers who presume to be familiar with Dan's height, but not his other property values. The individuals *Dan* denotes in the completions of the context c presupposed by these speakers differ in weight, mood, etc., but not in height. The statement *Dan is tall* is true in c if and only if Dan's completion-invariant height exceeds the completion-dependent standard height in every completion; it is false in c if and only if Dan's height fails to exceed the standard height in every completion, and the statement is undetermined otherwise (Kamp 1975). Thus, a predicate like *tall* is vague in c because the cutoff point of *tall* varies in T_c .

However, variance in cutoff points across indices cannot explain the truth value gaps pertaining to the statements in (6). The graded property of *tall*, $f(\text{tall}, t)$, in each t , maps individuals to index-invariant heights. So the ordering and ratios between individuals' heights are also index-invariant. Hence, no denotation gaps are admitted by comparison predicates like *taller* (which relate to ordering), or measure predicates like *two meters tall* (which relate to ratios, meaning, roughly, "twice as tall as a meter"). But, then, how can we possibly not know the truth value of statements like *Dan is two meters tall* or *taller than Sam*? Given a representation of individuals by their property values, it follows that this is possible due to absence of information concerning the interpretation of the arguments in these statements (*Dan, Sam, she, I, etc.*)

Importantly, this chapter takes individuals to be real entities, identified by their real properties. So, the height values are invariably determined for all individuals in D_x .³ However, when we use proper names, we do not know exactly which individuals they refer to, since we do not know all of their property values.

Let us explicitly distinguish between the notion *discourse entities*, namely the partially known referents of terms, and *individuals*, namely elements of D_x . Each discourse entity corresponds to one and only one individual x in D_x in each completion. However, the fact that a speaker observes a discourse entity x (*Dan, that girl over there, you, etc.*) does

not help her select that unique x in D_x that this discourse entity is. That is so because most of x 's properties are not accessible (x 's precise height, weight, mood, etc.) Our partial knowledge about a discourse entity is limited to a proper subset of the property values of the worldly individual it is. Thus, in practice, discourse entities reduce to subparts of individuals (cf. Landman's 1986, 1990 notion of *partial individuals*, and related notions discussed therein). A name conventionally used to refer to an entity is in practice doomed to refer to many individuals in D_x , each sharing that subpart. So when we say things like *Let's call our baby "Dan,"* in actuality, we only associate the name *Dan* with a set of possible referents (individuals in D_x that may agree on many property values – e.g. their names, IDs, addresses, family, etc., but not on, for example, their height). In practice, *Dan* is linked to no particular individual. We say *Dan*, but, not being omniscient, we do not and cannot know exactly to which member of the set of possible individuals we are referring. So *Dan's* denotation should be represented nonrigidly, as completion-dependent. For the same reasons, no linguistic community can specify the unique referent of the proper names in the given language. So *pace* Kripke (1980), denotations of proper names like *Dan*, *Princeton*, or *University of Amsterdam* should be represented nonrigidly, as completion-dependent.⁴

When we do not know, for example, the heights of two discourse entities, say *Dan* and *Sam*, we do not know to which two individuals these names refer in c . When this happens, we may easily not know how their heights compare. If *Dan's* height is not accessible to the speaker (its referent is 1.87 m tall in t_1 , 1.86 m tall in t_2 , etc.), she may not know whether *Dan is taller than Sam, twice as tall as Sam* (if, say, *Sam* is a child) or *1.87 meters tall* is true or not.⁵

Furthermore, consider again proper names of mountains or cities like *Etna* and *London*. We do not know the exact piece of land they denote in every time point. We are either ignorant about this matter or perhaps their linguistic meaning fails to determine a precise index-invariant denotation per each time point (or both). Who can tell? According to the present proposal, we need not have an answer to this question. A unified supervaluationist account of both phenomena readily captures these cases (see Lewis 1988, 1993 for a critical analysis of alternative solutions to problems concerning multiple referents of singular terms).

Denotation gaps caused by ignorance might differ from those found in vague cases with respect to higher-order vagueness. In vague expressions, the gap itself lacks sharp boundaries. In proper names, the gap often appears to have sharp boundaries. Higher-order vagueness is

often located and accounted for in terms of *definitely* operators (see e.g. Barker 2002), for example, often some heights are considered *definitely tall* while others are *tall* but not definitely so. However, *definitely* also applies to proper names and other terms when either semantic indeterminacy or ignorance regarding the reference are at stake, i.e. when some individuals appear to be more plausible referents than others (in just the same way as the plausibility of exceeding *tall's* cutoff point is bigger for some heights).⁶ If a given piece of land may be the referent of *Mount Etna*, so is any piece of land differing in one centimeter, two centimeters, etc., but at a certain point when the region is too big or too small it becomes harder to say that this is *definitely Mount Etna*. Certain pieces of land are more plausible referents than others. Furthermore, in the same way, the domain of individuals D_x may be ordered by plausibility of forming *Dan's* referent. Given Dan's known property values (say, his being one year old), other property values (say, weight of 5, 10, ..., 100 kilograms) may be plausible to different degrees; assume 11 kg is the average weight of one-year-olds. Then, individuals whose weight is 80 kg are *definitely not Dan*, but individuals whose weight is 10 kg (and all their other property values are consistent with what we know of Dan) are not *definitely not Dan*. These degrees coincide with truth value judgments of modal statements such as *Dan can weigh 15 kilograms*, *Dan can't possibly weigh 100 kilograms*, etc. (see Sauerland and Stateva 2007 and this volume for more regarding the distribution of modifiers across vague versus epistemically indeterminate predicate interpretations).

The present proposal diverges from standard proposals in the following respects. First, typically, formal semantic analyses build on the assumption that facts pertaining to the height of all the entities in the domain are given in actual contexts c . If you measure them, you find out what they are, if you do not it does not have anything to do with the fact that measure and comparison statements have determined truth values. But in actuality, many property values cannot as yet be accessible to a whole community of speakers. Such facts are represented in the present proposal.⁷

Second, typically in formal models, proper names are rigid designators, and individuals are often identified across different indices even if their property values vary. In order to represent the fact that the interpretation of *Dan is taller than Sam* may be unknown in c , *tall's* graded property must assign the unique referents of *Dan* and *Sam*, x_1 and x_2 , values whose rank order is different in different completions t in T_c . Only by virtue of that would the pair $\langle x_1, x_2 \rangle$ be an element of $\llbracket \text{taller} \rrbracket_{t_1}$ in one index $t_1 \in T_c$, rendering the statement *Dan is taller than Sam* true in t_1 , while

the reversed pair $\langle x_2, x_1 \rangle$ be an element of $\llbracket \text{taller} \rrbracket_{t_2}$, in another index $t_2 \in T_c$, rendering the statement false in t_2 . How can that come to be the case?

The assumption that x_1 and x_2 stand for real-world entities (and as such cannot have two different heights in two indices) can be rejected. On such theories, we may not know whether *Dan is taller than Sam* is true or false simply because we may be ignorant with regard to the heights of the unique referents of names like *Dan* and *Sam*. But that means that individuals in D_x are vague, partial objects, just like discourse entities are in the present proposal. But, then, this fact deserves an explicit representation, as in the present proposal.

Alternatively, if real-world entities cannot have two different heights in two indices, perhaps one and the same real-world individual x_1 would be considered, for example, *1.85 meters tall* in t_1 and *1.87 meters tall* in t_2 . Or perhaps, 1.87 meters would be considered *taller than* 1.86 meters in t_1 but not in t_2 . These options appear highly unintuitive.

Either way the intuitive distinction between *tall* and *taller* is blurred. *Tall* is vague; we intuitively regard it as admitting gaps. But *taller* is not; it should, intuitively, rarely admit denotation gaps. Rigid designation for proper names means that in both predicate types the denotations vary considerably across indices. Thus, vagueness models typically (e.g. Kamp 1975) abstract away from cases of accidental ignorance about property values, representing comparative and measure predicates as possessing an index-invariant (nonvague) interpretation. The representation of truth value gaps in these cases, it is claimed, requires the incorporation of an epistemic modal base. But the question of how a representation of epistemic modality may be incorporated into a vagueness model (or the other way around), remains unresolved, because if one actually combines the two representations, one risks losing the distinction between vague and nonvague predicates.

On the present proposal, this problem is solved. While predicates like *two meters tall* (or *taller*) denote the same set of individuals (or pairs of individuals) in every completion $t \in T_c$, they do not relate to the same discourse entities (or pairs thereof) in every t . If x_1 has more height than x_2 , $\langle x_1, x_2 \rangle \in \llbracket \text{taller} \rrbracket_t$ in every $t \in T_c$. Truth value gaps are nonetheless possible. In indices in which the referent of *Dan* is x_1 (the discourse entity “Dan” is associated with x_1), and the referent of *Sam* is x_2 (the discourse entity “Sam” is associated with x_2), we regard the statement *Dan is taller than Sam* as true, and in indices in which it is the other way around (x_2 is taller), we regard the statement as false, despite the fully determined interpretation of *taller* and due to the partial

interpretation of *Dan* and/or *Sam* – our partial knowledge about their property values.

So the denotation of *taller than* is always completely known, but this knowledge is trivial, it is not empirical. It is about imaginable individuals, indices in D_x . Concerning actual things that we see or speak about (discourse entities), we can only say that they stand in the relation *taller than* if and only if we have knowledge about their heights (if, say, we measure them with a ruler). Then, we can tell that in any t , all the individual pairs which are not in $\llbracket \textit{taller than} \rrbracket_t$ are not the referents $\llbracket \textit{Sam} \rrbracket_t$ and $\llbracket \textit{Dan} \rrbracket_t$, respectively.

To wrap up, the present proposal can represent vagueness and accidental ignorance side by side, distinguishing between them by their source – absence of information about interpretations of vague predicates (undetermined cutoff points) versus absence of information about interpretations of arguments (discourse entities), respectively. Perhaps vagueness is inherent ignorance concerning the denotations of vague predicates, but the impossibility of knowing the denotations of proper names is also inherent (there are too many property values to specify). Thus, to the extent that vagueness is ignorance, vagueness is not ignorance of a special *nature* (“inherent”), except in as much as its nature depends on its special *source*.

5.4 Demonstratives as nonrigid designators, identity statements as informative

Here are a few comments and examples as to how a formal model with partial interpretations for terms works.⁸

First, recall that observations of worldly entities (referents of proper names) only isolate subsets of the domain, the sets of those possible individuals consistent with the information that we have about the observed entity (its color, height, length, character, and so on). Thus, like predicates, also the interpretation of proper names can and should be modeled by means of both positive and negative denotations, as explained below.

The partially determined *positive denotation* of a predicate P in a context c , $\llbracket P \rrbracket_c$, equals the intersection of P 's totally determined positive denotations, $\llbracket P \rrbracket_t$, in c 's completions $t \in T_c$ (7a). The *negative denotation* $\llbracket \textit{not } P \rrbracket_c$ equals the intersection of P 's totally determined negative denotations, the complements of P 's positive denotations in c 's completions, $\llbracket \textit{not } P \rrbracket_t = D_x - \llbracket P \rrbracket_t$ (7b). The same applies of any proper name A , as stated in (8), except that the fully determined positive denotations are always singletons: $\llbracket A \rrbracket_t = 1$, for any $t \in T_c$.

- (7) a. $\llbracket P \rrbracket_c = \cap \{ \llbracket P \rrbracket_t \mid t \in T_c \}$
 b. $\llbracket \text{not } P \rrbracket_c = \cap \{ D - \llbracket P \rrbracket_t \mid t \in T_c \}$
 c. P 's gap in c is the set $D_x - \llbracket P \rrbracket_c - \llbracket \text{not } P \rrbracket_c$
- (8) a. $\llbracket A \rrbracket_c = \cap \{ \llbracket A \rrbracket_t \mid t \in T_c \}$.
 b. $\llbracket \text{not } A \rrbracket_c = \cap \{ D - \llbracket A \rrbracket_t \mid t \in T_c \}$
 c. A 's gap in c is the set $D_x - \llbracket A \rrbracket_c - \llbracket \text{not } A \rrbracket_c$

If we know that Sam's height is 1.86, but we do not know the extent to which Sam loves Dan, or what her ID number is, the positive denotation, $\llbracket \text{Sam} \rrbracket_c$, is empty, but the negative denotation, $\llbracket \text{not Sam} \rrbracket_c$, consists of all those individuals whose height is not 1.86. Thus, the negative denotation encodes information about the proper name's referent. Everything in it cannot possibly be Sam, while everything outside it might still be Sam. If a property P (say, *human*, *happy*, *60 kg*, etc.) holds true of everything outside it, we can positively say that the referent is P , despite the fact that we do not know it (the positive denotation, $\llbracket \text{Sam} \rrbracket_c$, is empty).

Second, given the present proposal, identity statements such as *Dan is Mr Cohen* can be highly informative. The question whether two referents of proper names are the same or not remains open as long as there are completions in which the given proper names denote the same individual, and completions in which they do not. Additionally, when we learn that *Dan is Mr Cohen*, we can add all the entities that are known not to be Dan to $\llbracket \text{not Mr Cohen} \rrbracket_c$ and all the entities that are known not to be Mr Cohen to $\llbracket \text{not Dan} \rrbracket_c$, which means that we may gain a lot of information about both.

Third, statements about what might or might not be true (or have been true) of *Dan* can only be based on the accessible subpart of Dan. Importantly, the present account diverges from descriptive theories, whereby we fundamentally understand proper names the way we understand descriptions in the sense that the reference of, for example, *Aristotle* has a descriptive content such as "the master of Alexander" or "the student of Plato" (Frege 1892; Searle 1958). This view has been rejected for a variety of reasons (see e.g. Kripke 1980). In particular, proper names, unlike definite descriptions, tend not to allow *de dicto* readings, as illustrated in (9). Assume Dan is the dean. Intuitively, (9a) can relate to whoever could be *the dean* (rather than to the dean in the index of evaluation, e.g. Dan), while (9b) cannot relate to whoever could be the dean, or whoever could be called *Dan*.

- (9) a. Mary could have married the dean.
 b. Mary could have married Dan.

I take it that in intensional contexts (e.g. contexts of use of *could*), some meaning components of terms like *the dean* or *Dan* have to remain fixed across the indices of evaluation quantified over (the modal base of *could*), while others have to change. Predicates have two meaning components – their denotation and their entailments (e.g. *dean* entails a person, with an authority over an academic unit or area). These entailments form the descriptive content of definite descriptions like *the dean*. This content remains fixed in *de dicto* readings, while the referent varies. In contrast, on the present account, proper names like *Dan* have no descriptive content whatsoever, no entailments to be kept fixed, only a referent, which in partial contexts is only partially known. In principle, every individual can bear any name. No meaning component other than the known subpart of the referent (its accessible property values) can identify *Dan* across the indices of evaluation quantified over. Therefore no variance in the property values is allowed except the minimum necessary to imagine, e.g. Mary marrying the referent. Otherwise, statements like (9b) would have been too trivial, conveying that, for example, Mary could have married an individual. Thus, the fixed set of property values must be substantive enough for them to be indicative of a counterpart of *Dan* (cf. Lewis 1986). In this way, what could or could not be true (or have been true) of *Mary* and *Dan* can only be based on the accessible subpart of *Mary* and *Dan*. Some properties may intuitively seem more compatible with it than others. So modal statements are about possible individuals who share a substantial subpart with *Mary* and *Dan*, not just about any possible individual called *Dan*, unlike modal statements about, for example, *the dean*, which are precisely about any possible referent of *dean*. The deans in the counterfactual worlds (or completions) considered for (9a)'s *de dicto* reading need not share any substantial set of property values with the actual dean. They only need to share with it the few property values which are necessary for falling under *dean*.

Fourth, for any predicate *P*, $\llbracket \textit{the } P \rrbracket_c$ denotes an element $x \in D_x$ if and only if x is known to be the unique element in *P*'s positive denotation in any t in T_c . But this is rarely if ever the case. For example, even when someone points at a given entity stating that this is the dean, many fully determined individuals may still correspond with this discourse entity, rendering the positive denotation of the definite description, $\llbracket \textit{the dean} \rrbracket_c$, empty in c .⁹ Still, the negative denotation, $\llbracket \textit{not the dean} \rrbracket_c$, encodes all the information about the referent available in the context, as explained in the first point above. In *de re* readings of statements such as (9a), a substantial part of this information remains fixed, substantial

enough for it to be indicative of a counterpart of the referent, as explained in the third point above.

Fifth, we may be aware of the existence of an individual, while some of its values may not be accessible to us, if, say, it is covered by a blanket. This means that we know that the statement *An individual exists under this blanket* is true in the given context c , but we do not know which one of the possible individuals it is. We do not know which one of the possible individuals the discourse referent *this individual* denotes. We can refer to this individual by saying *this* or *the individual under the blanket* but the interpretation of these expressions will still be empty. So even variables and pronouns, whose interpretation appears to directly pick an individual, at least in deictic contexts, are indeterminate in the sense that we cannot tell the exact x in D_x to which they refer. We can only access a subpart of that individual. So, when we say "I" (or *this*, pointing to something), different assignments of semantic values are still plausible. For example, if we point at a ball and say *this ball is red* that does not allow the hearer to put any particular individual in D_x into the denotation of *red*. After all, we do not know which individual the ball we were shown is. However, it does allow the hearer to accommodate two things. One, if something has the color of this individual (no matter which individual it actually is) the hearer can now consider that thing to be red (so the hearer has a better idea concerning how *red* divides the domain to sets of red, not red and borderlines in the given context). Two, whatever the reference of *this ball* is, it is an object which is red. In other words, the hearer can add everything that is not red to the negative denotation $\llbracket \textit{not this ball} \rrbracket_c$.

According to the present proposal, even the denotation of *I* is empty in contexts c , but here is an optimist remark in this regard. It is true that in the absence of any information regarding the denotations of terms one cannot gain information about the world. Even complete information about the truth values of all possible statements of the form $P(a)$ (for any predicate P and term a) tells one next to nothing if every possible individual may form the referent of every possible term (Groenendijk et al. 1996). Demonstratives, therefore, play an invaluable role of unambiguously establishing a connection between a term and a worldly entity. Even if the use of a demonstrative, as in, for example, *this is a ball*, establishes a connection to no unique individual in D_x , partial information about the denotations of terms may yield much information about the world. A term whose referent we feel that we know (e.g. *this* in any appropriate use of this word) is associated with a unique set of possible referents, not overlapping with the set of possible referents of any other

noncoreferring term. The individuals in this set share a unique subset of property values. Hence, the subset accessible to us may be partial, but still different from that of any other discourse entity we know. Thus, given the present account the entities we know do not correspond to any known individuals, but they do correspond to nonoverlapping subsets of individuals. Our information, then, consists of more than just relations between words. Despite its partiality, it suffices to identify and distinguish between most of the individuals we encounter.

Last but not least, notice that there are many senses in which one can *know who x is* for any proper name *x* (Groenendijk et al. 1996, note 20). For example, if we both observe a line of four individuals and I point at the third and tell you *this is Sam*, in a certain sense of the words you now *know who Sam is*. Certainly, one may not know, for example, Sam's precise height, but still feel that she knows Sam. The truth value of statements such as *Sam is 1.87 meters tall* is something we typically ignore in situations in which we know someone by acquaintance or friendship. Thus, knowing *x* by acquaintance does not entail knowing the referent, only a subset of its property values. Unfortunately, or perhaps fortunately, the same is true also with regard to knowing a friend, knowing a partner, or even knowing oneself. At least according to the present proposal, life can never be dull; there is always something new to discover about each and every one we know.

5.5 Conclusions

While individuals are not usually represented by their property values, this option seems to yield intuitively correct results regarding the way proper names and other terms are used. In addition, it directly captures the intuition that cross-index (world or completion) identity between individuals differing in their property values is implausible, while cross-index identity between individuals not differing in their property values is not implausible. Also, the identification of individuals with their property values captures the intuitive distinction between epistemic and vagueness-based truth value gaps, and the fact that identity statements are informative. This proposal draws a rather surprising and interesting picture concerning the semantics of demonstratives. Finally, the discussion in section 5.4 briefly suggests that this account can provide insights concerning interactions of terms with modals.

All in all, the proposal studied in this chapter proves to be a fruitful source of new ways for dealing with philosophical problems pertaining to the inhabitants of vagueness models.

Acknowledgments

Most of this work was carried out in the project “On vagueness – and how to be precise enough,” funded by the Netherlands Organization for Scientific Research (NWO 360-20-201). I have benefited a lot from insightful comments by Nirit Kadmon (back in 2007), Frank Veltman, Paul Égré and Nathan Klinedinst; needless to say, any mistakes are clearly mine.

Notes

1. Back to Russell (1905), linguists assume that adjectives like *tall* map individuals to degrees. Clearly, *tall*'s degrees reflect the ordering between entities' heights, but little is agreed about them beyond this fact (Klein 1991). The notion of degrees as property values (heights, lengths, etc.) resembles Moltmann's (2006) notion of “tropes” and Kennedy's (2001) “degree sorts.”
2. Truth value judgments of statements like those in (6) admit variation given variance in the granularity level contexts presuppose. For instance, assume that Sam's height is 1.80 cm and Dan's height is 1.80 and a millimeter. Contexts in which we care about millimeter differences render statements such as *Dan is taller than Sam* true and statements such as *Dan is 1.80 cm tall* false, while contexts which are only sensitive to coarser differences (say, centimeters) render the former false and the latter true. However, when the contextual granularity level is given, the statements in (6) differ from those in (5) in having a determinate truth value. It is this point that is crucial for us.
3. Of course, individuals not inhabiting the real world are “real” only to the extent that other possible worlds are real.
4. Even if it is not a necessary feature of any property value to be unknowable, there are simply too many property values for us to know all of them.
5. The truth value of comparatives such as *Sam is taller than Jim* is something we commonly know even when we are ignorant of the truth value of statements of the form *Sam is 1.87 meters tall*, simply because we are able to make a comparison between the heights of Sam and Jim, whatever they might precisely be. In such situations, in different completions, the referents of both *Sam* and *Jim* have different heights, but the referent of *Sam* is always taller than the referent of *Jim*.
6. There is no consensus as to whether and how supervaluationist models can account for higher-order vagueness. An interesting possibility is that each context of utterance u is associated with a cloud of partial contexts C_u , instead of a unique context c_u . After all, speakers are never sure about every precise detail concerning their background assumptions. In particular, they are not sure where the boundaries of gaps of vague expressions are placed. A cloud C_u determines a whole set of possible cutoffs between the positive denotation and the gap and a whole set of cutoff points between the negative denotation and the gap, i.e. the desired second-order vagueness. Likewise associating contexts of utterance with clouds of clouds creates third-order vagueness, and so on and so forth.
7. Intuitively, one says that Sam must be hungry if and only if one is almost, but not completely, sure that Sam is hungry; this happens occasionally, because no

one but Sam has access to Sam's hunger and its extent; one can only inductively infer (cf. Veltman 1984) that this value is above the standard of *hungry*, by using indirect, inconclusive symptoms, such as Sam's being pale or nervous.

8. For a more complete formal model with partial interpretations for terms see Sassoon (2007, Ch. 6).
9. For any P , $\llbracket \textit{the } P \rrbracket$ is well defined (although unknown) if and only if P 's positive denotation is still empty but some elements in D_x may still be the unique P . $\llbracket \textit{the } P \rrbracket$ is undefined (denotes \perp) if and only if more than two elements in D_x are already known to be P , or all the elements in D_x are already known not to be P .

References

- Barker, C. (2002). The dynamics of vagueness. *Linguistics and Philosophy*, 25(1): 1–36.
- Fara, D. (2000). Shifting sands: an interest-relative theory of vagueness. *Philosophical Topics*, 28(1): 45–81 (originally published under the name Delia Graff).
- Fine, K. (1975). Vagueness, truth and logic. *Synthese*, 30: 265–300.
- Frege, G. (1892). Über Sinn und Bedeutung. In *Zeitschrift für Philosophie und philosophische Kritik*, 100: 25–50. Translated as “On sense and reference”, in Geach, P. and Black, M. (eds), *Translations from the Philosophical Writings of Gottlob Frege*, 3rd edn, 1980, Blackwell, Oxford.
- Groenendijk, J., Stokhof, M. and Veltman, F. (1996). Coreference and modality. In Lappin, S. (ed.), *The Handbook of Contemporary Semantic Theory*. Blackwell, Oxford.
- Kamp, H. (1975). Two theories about adjectives. In Keenan, E. (ed.), *Formal Semantics for Natural Language*, pp. 123–55. Cambridge University Press.
- Kamp, H. and Partee, B. (1995). Prototype theory and compositionality. *Cognition*, 57: 129–91.
- Kennedy, C. (1999). *Projecting the Adjective: the Syntax and Semantics of Gradability and Comparison*. Garland, New York. (1997 UCSC doctoral dissertation.)
- Kennedy, C. (2001). Polar opposition and the ontology of “degrees.” *Linguistics and Philosophy*, 24: 33–70.
- Klein, E. (1980). A semantics for positive and comparative adjectives. *Linguistics and Philosophy*, 4: 1–45.
- Klein, E. (1991). Comparatives. In von Stechow, A. and Wunderlich, D. (eds), *Semantik/Semantics, an International Handbbook of Contemporary Research*, pp. 673–91. Walter de Gruyter, Berlin and New York.
- Kripke, S. (1980). *Naming and Necessity*. Harvard University Press.
- Landman, F. (1986). Pegs and alecs. In Landman, F., *Towards a Theory of Information. The Status of Partial Objects in Semantics*. GRASS 6, Foris, Dordrecht; *Linguistics and Philosophy*, 97–155.
- Landman, F. (1990). Partial information, modality, and intentionality. In Hanson, P. P. (ed.), *Information, Language, and Cognition*. University of British Columbia Press, Vancouver.
- Lewis, D. (1986). *On the Plurality of Worlds*. Blackwell.
- Lewis, D. (1988). Vague identity: Evans misunderstood. *Analysis*, 48(3): 128–30.

- Lewis, D. (1993). Many, but almost one. In Campbell, K. Bacon J. and Reinhardt, L. (eds), *Ontology, Causality, and Mind: Essays on the Philosophy of D. M. Armstrong*. Cambridge University Press. Reprinted in David Lewis (1999). *Papers in Metaphysics and Epistemology*. Cambridge University Press.
- Moltmann, F. (2006). Comparatives without degrees. A new approach. A manuscript for the Workshop on Scalar Meaning, University of Chicago.
- Murphy, G. (2002). *The Big Book of Concepts*. The MIT Press, Cambridge, Mass.
- Russell, B. (1905). On denoting. *Mind*, 14, 479–93. Reprinted in Russell, B. (1973). *Essays in Analysis*, pp. 103–19. Allen and Unwin, London.
- Sassoon, W. G. (2007). Vagueness, gradability, and typicality – a comprehensive semantics analysis. Unpublished doctoral dissertation, Tel Aviv University.
- Sauerland, U. and Stateva, P. (2007). Scalar vs epistemic vagueness: evidence from approximators. In Friedman, T. and Gibson, M. (eds), *SALT XVII*, pp. 228–45. Cornell University, Ithaca, NY.
- Searle, J. (1958). Proper names. *Mind*, 67(266): 166–73.
- Sorensen, R. (1988). *Blindspots*. Clarendon Press, Oxford.
- Stalnaker, R. (1978). Assertion. In Cole, P. (ed.), *Syntax and Semantics 9: Pragmatics*, pp. 315–32. Academic Press.
- van Fraassen, B. C. (1969). Presuppositions, supervaluations and free logic. In Lambert, K. (ed.), *The Logical Way of Doing Things*, pp. 67–91. Yale University Press, New Haven.
- Veltman, F. (1984). Data semantics. In Groenendijk, J., Janssen, T. and Stokhof, M. (eds), *Truth, Interpretation and Information, Proceedings of the 3rd Amsterdam Colloquium*. Floris Publications, Dordrecht.
- Williamson, T. (1994). *Vagueness*. Routledge, London and New York.

Part II

Approximators and Intensifiers

This page intentionally left blank

6

Two Types of Vagueness*

Uli Sauerland and Penka Stateva

6.1 Introduction

This chapter is about vagueness in natural language semantics. More specifically, we discuss lexical means of making vague assertions more or less precise in compositional semantics. Examples of expressions that have this effect are *approximately*, *absolutely*, *definitely*, and *roughly speaking*. While many of these expressions are modifiers and adverbs, some such expressions are neither. Hence, for the purposes of this chapter we call expressions that make vague assertions more or less precise *approximators*. Our main claim is that the distribution of such expressions provides evidence for the view that vagueness in language comes in at least two varieties, which we call *scalar vagueness* and *epistemic vagueness*.¹

We assume that vagueness is characterized by giving rise to the sorites paradox. Consider the three examples in (1) through (3). In each case a sorites paradox reasoning goes through: if we accept the premise that “If Harry is bald, it’s always true that, even if he had one more hair than he does, he would still be bald,” which sounds innocent, it follows that (1) does not tell us anything about the number of hairs on Harry’s head. Similarly, (2) would not tell us how many grains of sand there are in the heap if we accept the premise that “if something is a heap of sand and we remove one grain from it, it is still a heap of sand.” Finally, (3) would not tell us anything about the time of John’s arrival if we accept the premise that “if John arrived at 6 o’clock, then (3) would still be true if he arrived one second later than he actually did.”

- (1) Harry is **bald**.
- (2) This is a **heap** of sand.
- (3) John arrived at **6 o’clock**.

In our intuition, all three cases are acceptable though for each of them there is a slight feeling of discomfort associated with the inductive premise of the sorites argument, perhaps because we anticipate the paradoxical conclusion at this point. A very natural example of an underlying sorites premise in a case similar to (3) was pointed out to us by Stephanie Solt (p.c.). The reasoning in this case is part of a joke about a young paleontologist:

A young paleontologist is engaged in excavating a dinosaur bone when a tourist comes by and is looking over his shoulder. After a while, the tourist asks the paleontologist: "How old is the dinosaur you're excavating?" "100 million and three years," the young paleontologist replies. The tourist is amazed: "How can you know so precisely?" The paleontologist explains: "When I started working at this site three years ago, my boss told me the dinosaur was 100 million years old."

Note that the point of the joke, in essence, is that the young paleontologist is not obeying the inductive premise of a sorites argument. An experienced paleontologist would have answered that the dinosaur is 100 million years old because he or she would have known the following principle: if *100 million* was the appropriate response at some time, then it is also the appropriate response 3 years later. That is exactly the inductive premise of the sorites argument.

We contrast two views of vagueness: the monistic and the dualistic view. The monistic view of vagueness assumes that there is one general mechanism of vagueness that underlies all phenomena exhibiting vagueness in natural language. Adherents to the uniform view do not all agree on what the general mechanism of vagueness is, but they agree that there is just one mechanism. For instance, the following treatments of vagueness advocate a uniform view: Lakoff (1973), who analyses all forms of vagueness using fuzzy logic, Kamp (1981), who uses supervaluations, Lasersohn's (1999) "Pragmatic Halos," and to some extent also the epistemic view of vagueness of Williamson (1994). The dualistic view assumes that there are at least two mechanisms that can give rise to vagueness. Adherents to the dualistic view like Pinkal (1995) and Kennedy (2007) have, as far as we are aware of, generally distinguished between (2) and (3). Intuitively, the distinction is that there seems to be no precise concept of *heap*, while there is a precise concept of *6 o'clock*: a point in time. However, the expression *6 o'clock* never really refers to this precise concept, but rather to a broader, vague one – otherwise,

a sentence like *I arrived at 6 o'clock* could never be true.² This intuitive distinction underlies the dualistic view of vagueness.

This chapter seeks to make two contributions to the debate between the two views. One contribution is a novel argument for the dualistic view. The intuitive distinction between vague concepts based on underlying precise concepts and vague concepts without such a basis does not *prima facie* necessitate two different accounts of vagueness in language. We contribute a new argument based on linguistic, distributional evidence for the dualistic view. The second is a formal account of the dualistic view spelling out mechanisms for both scalar and epistemic vagueness and also the approximator expressions that regulate both kinds of vagueness. New tests distinguishing scalar and epistemic vagueness are one application of our work. This is useful because there are cases where the distinction is intuitively unclear: existing work that advocates the dualistic view does not say where expressions like *bald* in (1) are classified. In fact, *bald* is generally considered a core case of vagueness in language and is often used to exemplify the concept (for instance, by Kamp (1981) and Williamson (2000: 102)), but it also is intuitively related to a precise concept – that of having no hair whatsoever on the scalp. The evidence from approximators we discuss in this chapter argues that *bald* exhibits scalar vagueness, i.e. really has *having no hair* as its basic meaning.

6.2 Initial support for the dualistic theory: distributional differences between approximators

The example in (4) illustrates that approximators have a limited distribution: while *exactly* and *approximately* easily combine with *fifty*, they are unacceptable with *Beef Stroganoff*, as in (4b):

- (4) a. What John cooked were exactly/approximately fifty tapas.
 b. # What John cooked was exactly/approximately Beef Stroganoff.

The contrast in (4) between different vague expressions is unexpected from the point of view of the monist theory of vagueness. Specifically, this point is shown below for the analysis of Lasersohn (1999), which is the only monist analysis we are aware of that provides an account of some approximators. On our dualist account, however, the difference between *fifty* and *Beef Stroganoff* follows from that between scalar and epistemic vagueness.

There are also differences between approximators. Consider, for example, the expressions *definitely* and *maybe*. They can be used to express

approximation: for example, (5) when used while we are eating the dish John cooked. In such contexts, it is the meaning of the phrase *Beef Stroganoff* that is discussed, rather than the ingredients and method of preparation of the dish in front of us.³

(5) What John cooked is definitely/maybe Beef Stroganoff.

We take these distributional differences to suggest that there are two groups of approximators: scalar approximators and epistemic approximators. The two classes correspond to the two kinds of vagueness, scalar and epistemic vagueness.⁴

There are several other approximators that we consider in detail below. Specifically, we further distinguish between two kinds of approximators in the scalar domain: for instance, *absolutely*, *completely*, and *totally* are more or less in complementary distribution with *exactly*. We show below that *exactly* is used when the associated scalar expression denotes a mid-point of a scale while *absolutely*, *completely*, and *totally* are used only with endpoints (Kennedy and McNally 2005):

- (6) a. *What John cooked was absolutely/completely/totally fifty tapas.
 b. What John cooked was absolutely/completely/totally appropriate.

For the English examples, it is necessary to mention a further property of *exactly* right away. In the scope of negation, *exactly* can combine with any kind of predicate and patterns in that respect with the epistemic approximators.

(7) Red wine isn't exactly healthy.

This seems to cast doubt on the validity of the argument for distinguishing among kinds of vagueness. However, there are two reasons to believe that (7) does not threaten the suggested classification. For one, consider the German and Bulgarian counterparts of *exactly*. In (8a) and (9a) these are literal translations of English *exactly*. Unlike English *exactly*, however, both Bulgarian *točno* in (8b) and German *genau* in (9b) are not acceptable with epistemically vague predicates in negative contexts.⁵

- (8) a. Červenoto vino e točno trigodisno.
 red-the wine is exactly three-year-old

- b. #Červenoto vino ne e točno zdravoslovno
 red-the wine is not exactly healthy
- (9) a. Der Rotwein ist genau drei Jahre alt.
 the red-wine is exactly three years old
- b. #Der Rotwein is nicht genau gesund.
 the red-wine is not exactly healthy

In addition, even *precisely*, which is a near-equivalent of *exactly* in English, does not pattern like *exactly*:

- (10) #Red wine isn't precisely healthy.

We conclude then that English has a second lexical entry for *exactly*, which is a strong negative polarity item and an epistemic approximator, and leave it aside in the following.

Example (11) provides a further case where a scalar approximator can be combined with an epistemically vague expression. These examples can be translated literally into German and Bulgarian, and hence we have no evidence that we are looking at a separate lexical entry for *approximately* in (11). Nevertheless, this ambiguity analysis presently still seems to be the simplest analysis for (11) to us.

- (11) a. Red wine isn't even approximately healthy.
 b. What John cooked isn't even approximately Beef Stroganoff.

6.3 Proposal

The dualistic theory of vagueness that we advocate is based on the observation that markers of approximation come in two classes, depending on the kind of vagueness they make more precise or more vague. Vague predicates, we suggest, should be described as scalarly or epistemically vague.

6.3.1 Scalar vagueness

The first kind of vagueness can only be related to expressions that denote a point on a scale. Numerals can serve as a prototypical example. We have already mentioned in section 6.1 above that numerals allow for an interpretation that is consistent with a larger segment of the scale: they can denote an interval (Krifka 2007). For example, the expression *5 meters* could, in a given context, be a good description of the length




Figure 6.1 Alignment of length terms with scale intervals

of a rod that we know to have an actual length in the interval between 4.5 m and 5.5 m. In that case, the scale is partitioned in segments of 1 m.

Following Krifka (2007), we assume the model for the alignment of point-denoting scalar terms with granularity intervals of a scale that is illustrated by Figure 6.1. A scale can be simultaneously divided up into intervals of varying granularity. In principle, each point-denoting term is a candidate for denoting any interval that contains this point. However, general pragmatic principles entail that the shortest expression must be used for each interval.⁶ This yields the alignment of terms and scale intervals shown in Figure 6.1. For instance, *5 meters* denotes the interval from 4.50 to 5.50 at the 1 m granularity, while *4 meter 50* and *4 meter 90* are blocked from denoting this interval because they are longer expressions than *5 meters*. Instead, *4 meter 50* denotes an interval at the half-meter granularity, and *4 meter 90* at the 10 cm interval granularity. The grey intervals in Figure 6.1 are not denoted by any expression of the form *x meters y* because the shortest expressions denoting a point in these intervals already denote a larger interval.

6.3.2 Granularity functions

We propose that granularity is a contextual parameter of interpretation. Formally, we assume that a granularity function maps each point of a scale to an interval that contains it. Here are examples of different extensions of *5 meters* which vary because of a different setting of scale granularity:

- (12) a. $\text{gran}_{\text{fine}}(5 \text{ m}) = [4.95 \text{ m}, \dots, 5.00 \text{ m}, \dots, 5.05 \text{ m}]$
- b. $\text{gran}_{\text{mid}}(5 \text{ m}) = [4.75 \text{ m}, \dots, \dots, 5.00 \text{ m}, \dots, \dots, 5.25 \text{ m}]$
- c. $\text{gran}_{\text{coarse}}(5 \text{ m}) = [4.50 \text{ m}, \dots, \dots, \dots, 5.00 \text{ m}, \dots, \dots, \dots, 5.50 \text{ m}]$

Normally, several granularities are under consideration simultaneously, and hence the *gran* parameter of interpretation contains more than one granularity function. When a scalar point expression is evaluated, it is

mapped by the coarsest granularity such that the expression is the shortest expression that could denote the resulting interval. Hence, we end up with the mapping in (13):

- (13) a. $\llbracket 5 \text{ meters} \rrbracket^{\text{gran}} = \text{gran}_{\text{coarse}}(5 \text{ m}) = [4.50 \text{ m}, 5.50 \text{ m}]$
 b. $\llbracket 4 \text{ meters } 50 \rrbracket^{\text{gran}} = \text{gran}_{\text{mid}}(4.5 \text{ m}) = [4.25 \text{ m}, 4.75 \text{ m}]$
 c. $\llbracket 4 \text{ meters } 90 \rrbracket^{\text{gran}} = \text{gran}_{\text{finest}}(4.9 \text{ m}) = [4.85 \text{ m}, 4.95 \text{ m}]$

Our approach also applies to nonnumerical expressions like *the middle*. We assume that, if A is actually the center of the circle, the denotations of *the middle* for different granularity functions could be the following:

- (14) $\llbracket \text{the middle of the circle around A} \rrbracket$
 a. $\text{gran}_{\text{finest}}(\text{S})(\text{A}) = \text{the } 1 \text{ mm circle surrounding A}$
 b. $\text{gran}_{\text{med}}(\text{S})(\text{A}) = \text{the } 1 \text{ cm circle surrounding A}$
 c. $\text{gran}_{\text{coarse}}(\text{S})(\text{A}) = \text{the } 2 \text{ cm circle surrounding A}$

Our principles need to be generalized to many other scales. For closed scales, granularity functions seem to divide a space into equidistant intervals. Therefore, we assume the following definitions for *granularity function* and the notions *finer/coarser*: a granularity function γ for scale S has the following properties:⁷

- (15) a. $\forall s \in S : s \in \gamma(s)$
 b. $\forall s \in S : \gamma(s)$ is convex
 c. $\forall s, s' \in S : \max(\gamma(s)) - \min(\gamma(s)) = \max(\gamma(s')) - \min(\gamma(s'))$

A granularity function γ is *finer* (or *coarser* when $<$ is replaced by $>$) than γ' if:

- (16) $\forall s \in S : \max(\gamma(s)) - \min(\gamma(s)) < \max(\gamma'(s)) - \min(\gamma'(s))$

These definitions will be sufficient for our purposes in this chapter. Ultimately, though, it would be desirable to extend the approach to open scales, which Hobbs and Kreinovich (2006) argue to have logarithmic granularity.

6.3.3 Scalar approximators

What is the role of scalar approximators on this view? Let us go back to an example from the introduction (repeated from (4)):

- (17) What John cooked were approximately/exactly fifty tapas.

We propose that scalar approximators reset the granularity parameter to the coarsest granularity. For this reset, the following new composition principle is needed (assuming the general framework of composition of Heim and Kratzer 1998):

- (18) *Granularity modifying composition*: If $\llbracket A \rrbracket^\Gamma$ has in its domain functions that take sets of granularity functions as arguments, then the new composition rule $\llbracket A B \rrbracket^\Gamma = \llbracket A \rrbracket^\Gamma (\lambda\Gamma' \llbracket B \rrbracket^{\Gamma'})$ must be applied.

The lexical entries for *approximately* and *exactly* can then be given as in (19). Both set the granularity parameter for the evaluation of their complement to a singleton set; *exactly* to the finest and *approximately* to the coarsest granularity.

- (19) a. $\llbracket \text{exactly} \rrbracket^{\text{gran}}(G) = G(\{\text{finest}(\text{gran})\})$
 b. $\llbracket \text{approximately} \rrbracket^{\text{gran}}(G) = G(\{\text{coarsest}(\text{gran})\})$

This approach predicts that the use of *exactly* makes it possible to denote some of the grey intervals in Figure 6.1. For instance, *exactly five meters* denotes the interval from 4.95 m to 5.05 m in a context where exactly the three granularities drawn in Figure 6.1 are under consideration.

Our account also predicts the oddity of (20) in a general context: on the coarsest scale, 49 and 50 belong to the same interval. But then, 50 must be used over 49 to denote this interval.

- (20) # What John cooked were approximately 49 tapas.

We give lexical entries for other scalar approximators in section 6.4 below, and also discuss the restriction to midpoints of *exactly* and *approximately*. We argue there that at least the following expressions belong to this class:

- (21) a. Scalar more precise approximators: *exactly, absolutely, completely, precisely, perfectly*.
 b. Scalar less precise approximators: *approximately, about, partially, sufficiently, roughly*.

6.3.4 Epistemic vagueness and approximators

We observed above that expressions like *definitely* and *maybe* can be used as approximators. In contrast to scalar approximators, they can combine with any predicate that does not have a precise meaning or at least it is

not known. For such predicates we use the term *epistemically vague* (cf. Williamson 1994). We propose that epistemically vague predicates differ in their extensions even across worlds where physical object properties (i.e. the number of sand grains in a heap) do not differ.⁸ Let us take *heap* as a prototypical example of an epistemically vague predicate. As illustrated in (22), we assume that the minimum amount of sand that constitutes a heap can differ: the extension of *heap* may include in a possible world w_1 any pile of more than 20 grains, but in a possible world w_2 it may include only objects consisting of more than 30 grains, even though the two worlds are indistinguishable in terms of the location and size of objects.

- (22) $\text{heap}(w_1) = \langle \text{twenty grains, twenty-one grains, } \dots \rangle,$
 $\text{heap}(w_2) = \langle \text{thirty grains, thirty-one grains, } \dots \rangle$

This approach predicts that epistemic approximation arises as a side effect of general epistemic quantification. Since the epistemic approximators also all have general epistemic uses, this is a desirable result. Approximation is most clearly intended when there is no other epistemic uncertainty:⁹

- (23) This perfectly cone-shaped pile of 17 sand-grains on the table in front of us is maybe/definitely a heap.

Our assumptions naturally lead to a view that *maybe* and *definitely* express existential and universal epistemic quantification, respectively, and that the effect of approximation is a consequence of this use.

Distributional properties, different from those of scalar modifiers, as well as the quantificational nature of epistemic modifiers, are sufficient to establish the typology of vague predicates in line with the main goal pursued in this study. However, we present the semantics of epistemic modifiers in a significantly simplified manner. As previously observed (Barker 2002 among others), treating *definitely* as a modifier with universal quantificational force fails to acknowledge its potential to be involved in higher-order vagueness phenomena. In other words, an analysis that does not recognize that *definitely* is itself a vague expression that can be further modified is bound to predict the semantic equivalence between (24a) and (24b), contrary to native speaker's intuition:

- (24) a. John is definitely tall.
 b. John is definitely/clearly definitely tall.

A more careful definition will, in our view, incorporate these facts. Barker (2002) makes such a proposal, which is couched in terms of dynamic semantics,¹⁰ and suggests that the vague component in the meaning of *definitely* is akin to the meaning of the modifier *very*. Interestingly, Barker's implementation of this idea involves universal quantification over standards of comparison associated with the modified vague predicate. In the context of our distinction between scalar and epistemically vague predicates, this would imply that vague predicates are all intrinsically associated with some point-denoting concept. In the case of scalar vagueness this concept is epistemically accessible; it can be named and made precise by scalar approximators. The point-denoting concept that comes along with epistemically vague predicates, on the other hand, is in the default case a respective standard of comparison, which is epistemically inaccessible but its possible fluctuations are measurable in quantificational terms by epistemic modifiers. Another possible direction to pursue in expanding on our definition of *definitely* is to allow for second-order vagueness by incorporating Bayesian, probabilistic models (Lassiter 2009).

We think that our approach to *maybe* and *definitely* should also be extended to other epistemic approximators, such as the following:¹¹

- (25) a. Epistemic more certain approximators: *definitely, positively, for sure, certainly*.
 b. Epistemic less certain approximators: *maybe, -ish, like, if you will*.

English epistemic modals can also be used as approximators, but since they have an evidential component requiring indirect evidence (von Stechow and Gillies 2007), this requires a special context. For example, (26) could be used by me in a context where I just heard a third person, who speaks English natively, refer to the sand accumulation in front of me as a heap:

- (26) This cone-shaped pile of 17 sand-grains must be a heap.

6.3.5 Combinatorics: further support for the dualistic view

Stacking approximators provides further support for the dualistic view. Our theory predicts that scalar approximators could not be stacked for the following reasons: recall that we suggested that these approximators restrict the granularity parameter of their complement to one granularity function. A second scalar approximator in the scope of the first

is vacuous. The facts, as we see in (27a) and (27b), coincide with the predictions.¹²

- (27) a. #John is exactly/precisely approximately 30.
 b. #John is approximately exactly/precisely 30.


The monistic account of Lasersohn (1999), which addresses some approximators, makes different predictions for (27):¹³ Lasersohn (1999) represents vagueness uniformly as a *halo* of values close enough to the truth to be considered true, but not really true. For example, the halo of 30 in a given context could be the set of all points representing real numbers between 29.5 and 30.5. In the diagrams (28) through (31) truth is indicated by black, the halo area by grey, and falsity by white.

(28) '30':



In this theory *exactly* narrows the halo. Because of this function, Lasersohn calls it a *slack regulator*.

(29) 'exactly 30':



As a less precise approximator, Lasersohn discusses only *loosely speaking*. We assume *approximately* would receive the same analysis: it makes the predicate actually true for all the values that are otherwise only in its halo:

(30) 'approximately 30':



Lasersohn's analysis correctly predicts (27a) to be odd because *approximately 30* has no halo that could be narrowed. However, (27b) is predicted to be without blemish by Lasersohn's analysis: as (31) shows, the resulting predicate should be true of values in its narrowed halo.

(31) 'approximately exactly 30':



We therefore conclude that the halo theory cannot account for the distribution of approximators correctly.

6.4 Subclasses of scalar approximators

6.4.1 Distribution with adjectives

Within the class of scalar expressions, not all can combine with any scalar approximator. For example, *exactly closed* and *approximately wet* both do not sound natural. Our main claim in this section is that the scalar approximators should be divided into two subclasses: those that combine with endpoints of a scale, and those that combine with nonendpoints of a scale. *Absolutely* is an example of an approximator that makes the endpoint more precise, while *exactly* is an approximator that makes nonendpoints more precise. This claim can be examined by combining approximators with adjectives in English for the following reason: Rotstein and Winter (2004) as well as Kennedy and McNally (2005) argue that the scales associated with adjectives can be either open or closed, and that this affects the denotation of adjectives. Specifically, adjectives associated with closed scales denote the endpoints of the scale, while those associated with open scales cannot denote endpoints. Hence, our proposal predicts, on the one hand, that the endpoint-associated approximators should combine only with the closed-scale adjectives. Following Kennedy and McNally (2005), we explain the distribution of *absolutely* in (32) on this basis. Approximators like *exactly*, which are not restricted to endpoints, on the other hand, can also only combine with point-denoting scalar expressions. The semantic value we gave in (19a) for *exactly* initially predicts that it can combine with any point-denoting scalar expression. But, we expect furthermore that *exactly* should be blocked whenever a more specialized expression such as *absolutely* is applicable. Therefore, *exactly* is expected to combine only with expressions that denote nonendpoints on a scale. In this way, we explain the distribution of *exactly* in (32).¹⁴


- (32) a. The glass is *absolutely/#exactly* full.
 b. The glass is *exactly/#absolutely* half full.

Intuitions on the use of approximators are sometimes fleeting. For instance, (32a) with *exactly* improves after pondering it for a while. We think this is due to some coercion of the scale structure where *full* does not denote an endpoint. This is supported by an intuition that Chris Tancredi (p.c.) pointed out to us: *exactly full* could be used to describe a glass in which the liquid is exactly level and aligned with the upper rim of the glass; *absolutely full*, though, could also be used to describe a glass where the surface of the liquid it contains is bulging upwards higher than the brim of the glass.

To circumvent the effect of scalar coercion, we also tested our distributional claims with a corpus study using the British National Corpus (BNC). The result of this study is reported in Table 6.1. We measured how frequently an approximator–adjective sequence occurs relative to the individual frequency of approximator and adjective. The numbers we report in the table are such that if approximator and adjective were randomly distributed in the corpus, the value 1 should appear.¹⁵ The columns of the table show for each approximator the number of individual occurrences in the corpus altogether in the second row. Beginning from the third row, a relative frequency score is shown for specific scalar expressions. This value is calculated according to the formula at the top of the table: the number of occurrences of the string of the approximator followed by the scalar expression divided by the number of individual occurrences of each word. The rows contain for each of the adjectives we looked at, the number of individual occurrences in the first column and the frequency scores for each approximator in the following columns.

The results for the most part confirm our expectations. We have marked the unexpectedly high values with a box. We did not mark unexpected 0 values since we assume that the corpus is too small to conclude anything from these. The three most notable departures from our predictions are the frequencies of *approximately true*, *exactly right*, and *exactly satisfactory*. We assume that each of these is explained by reference to a coerced scale structure, but that two different processes of coercion are at work. For *exactly satisfactory*, we think that a scale is available where *satisfactory* denotes a nonendpoint between *unsatisfactory* and *great*. For *right* and *true*, we assume that the scale of truth can be adjusted to have truth as a midpoint between two poles of falsity as shown in (33).

(33)



We assume that the scale in (33) is available when the truth of a non-endpoint answer to a scalar question asking is relevant. Hence, we expect the contrast in (34):

- (34) a. A: How tall is Tony?
 B: 190 cm.
 A: You're exactly right./That's approximately true.
- b. A: How dry do you think the shirt is?
 B: It should be dry by now.
 A: #You are exactly right.

Table 6.1 Relative frequencies of approximator–adjective strings. British National Corpus, UK (1980–93), 100 million words interface by Mark Davis/Brigham Young University, <http://view.byu.edu>, the values reported are:

		100 mill. * Freq.(approximator–adjective-diagram) Freq.(approximator) * Freq.(adjective)					
Approximator	Individual frequency (approx.)	<i>approximately</i>	<i>exactly</i>	<i>precisely</i>	<i>completely</i>	<i>absolutely</i>	<i>partially</i>
		2,826	10,307	3,424	8,339	5,672	1,286
Adjective	Individual frequency (adj.)						
<i>true</i>	17,737	15.9	2.2	3.3	5.4	42.7	48.2
<i>false</i>	3,584	0	0	0	46.8	9.8	0
<i>right</i>	84,904	0.8	9.3	2.1	1.7	53.2	1.8
<i>wrong</i>	15,505	0	1.9	1.9	48.7	15.9	0
<i>clean</i>	6,281	0	0	0	13.4	8.4	0
<i>dirty</i>	2,666	0	0	0	0	0	0
<i>wet</i>	3,820	0	0	0	9.4	0	0
<i>dry</i>	6,443	0	0	0	46.5	0	12.1
<i>complete</i>	12,594	0	0	0	0	0	6.2
<i>incomplete</i>	703	0	0	0	0	0	0
<i>open</i>	29,268	0	0	0	8.6	1.8	26.6
<i>closed</i>	9,389	0	0	0	11.5	1.9	24.8
<i>satisfactory</i>	2,161	0	4.5	0	66.6	16.3	0
<i>unsatisfactory</i>	760	0	0	0	0	0	0
<i>bald</i>	611	0	0	0	294.4	0	0
<i>hairly</i>	392	0	0	0	0	0	0
<i>ten</i>	19,645	37.8	8.4	1.5	0	0.9	0
<i>half</i>	29,003	85.4	9.0	3.0	0	0.6	0

6.4.2 More precise scalar approximators

Before looking at other approximators, recall again the distribution of *exactly* as shown in (35): it cannot combine with the endpoint-denoting scalar expressions in (35a). With the nonendpoint-denoting scalar expressions in (35b), however, *exactly* is fine. These include numerals, time descriptions, spatial boundaries, *same*, and equatives. Finally, non-NPI (nonnegative polarity item) *exactly* cannot combine with the non-scalar expressions in (35c).

- (35) a. #*exactly* dry/pure/full/empty/white
 b. *exactly* three/?north/the same
 c. #*exactly* beef Stroganoff/a heap of wood

Our lexical entry for *exactly* in (19) at this point only explains the restriction of *exactly* to scalar vague terms: *exactly* is vacuous with epistemically vague predicates because granularity is not made use of by them. We assume that the restriction to nonendpoints follows from the fact that there are expressions like *absolutely* and *completely*, which specifically make the endpoints more precise.

For the semantic entries of further approximators, we adopt the convention that d be the type of degrees. Furthermore, we assume that adjectives are of the type $\langle\langle d, t \rangle, \langle e, t \rangle\rangle$. For example, for *dry* we assume the lexical entry in (36):

- (36) $\llbracket \text{dry} \rrbracket^{\text{w,gran}} = \lambda I \in D_{\langle d, t \rangle} \lambda x \in D_e. I$ holds of the degree to which x is dry.

Absolutely, *completely*, and *totally* are all three in complementary distribution with *exactly* in the scalar domain. As Kennedy and McNally (2005) have already observed, *completely* can combine only with predicates which denote endpoints of a scale as shown in (37a). It does not combine with nonendpoints in (37b) or with epistemically vague expressions in (37c). The distribution of *absolutely* and *totally* is very similar to that of *completely* as far as we can tell, and in the following we only talk about *completely*.¹⁶

- (37) a. *completely* dry/pure/white
 b. #*completely* three/north/the same
 c. #*completely* Beef Stroganoff/a heap of wood

The lexical entry in (38) is similar in spirit to Kennedy and McNally's (2005) suggestion, but it refers to the function of the modifiers as granularity setters.

$$(38) \llbracket \text{completely} \rrbracket^{\text{gran}} = \lambda f \in D_{(dt,et)} \lambda x \in D_e \\ \exists D \in D_{dt} [D = \text{finest}(\text{gran})(\max(\text{domain}(f))) \& f(D)(x)]$$

This entry takes an adjective as argument and returns a property. Hence, the lexical entry explains why *completely* must be adjacent to the scalar expression. It is distinct from *exactly* as shown by (39). Since it makes reference to the endpoint, (38) also explains why *completely* requires closed scales.

- (39) a. Mary arrived (*completely) with (completely) dry shirts.
 b. Mary arrived (exactly) at (exactly) noon.

Precisely has a distribution similar to that of *exactly*, but is preferably combined with anaphoric expressions denoting nonendpoints of a scale, as illustrated in (40). At present, we assume that *precisely* has the same lexical entry as *exactly* in (19).

- (40) a. #precisely dry/pure/full/empty/white
 b. ?precisely three/that amount/#north/there
 c. #precisely Beef Stroganoff/a heap of wood

Perfectly also belongs to this group of scalar approximators, and its meaning and distribution are closely related to those of *completely*:

- (41) a. perfectly dry/straight/clean/safe/appropriate/possible
 #perfectly full/empty/white
 b. #perfectly three/north
 c. #perfectly Beef Stroganoff/a heap

Perfectly has, however, an intensional component. Intuitively, an expression like *perfectly dry* makes reference to a desired point of dryness.¹⁷ Given that, it should be analyzed in a way similar to modal superlatives like *the driest possible*. Building on this intuition, we can suggest a lexical entry for *perfectly* that follows closely, in the relevant respect, Schwarz's (2006) definition of the modal superlative operator *-est-possible*. We assume that like *-est-possible*, *perfectly* moves to take sentential scope. It leaves behind a trace of the type $\langle d, t \rangle$, and combines with the lambda-abstract created by movement. The lexical entry for

perfectly in (42) applied to *The shirt is perfectly dry* can be paraphrased as: The shirt is completely dry and complete dryness of the shirt is desirable.

$$(42) \llbracket \text{perfectly} \rrbracket^{w, \text{gran}} = \lambda R \in D_{(dt, st)} \forall w' \in D_s \exists D \in D_{(d, t)} [\text{Acc}_w(w') \ \& \ D = \text{finest}(\text{gran})(\max(\text{domain}(R))) \ \& \ R(D)(w) = 1 \ \& \ R(D)(w') = 1]$$

The similarity in meaning between modal superlatives and *perfectly* explains the oddity of (43):

(43) #The shirt is perfectly/completely driest possible.

6.4.3 Less precise scalar approximators

Approximately clearly patterns with *exactly* in its distribution. It, too, combines only with scalar expressions not associated with endpoints as shown in (44):

- (44) a. #approximately dry/pure/white
 b. approximately three/north/the same
 c. #approximately Beef Stroganoff/a heap of wood

The lexical entry in (19) explains why the use of *approximately* with epistemically vague expressions is blocked in the same way as for *exactly*. For the restriction to endpoints, we claim that the use of *approximately* is blocked by *more or less*, just as *exactly* was blocked by *completely*.

We analyze *more or less* as the counterpart of *completely*, making reference to the coarsest granularity that combines with endpoints as shown in (45):¹⁸

- (45) a. more or less dry/pure/white
 b. #more or less three/north
 c. #more or less Beef Stroganoff/a heap of wood

Our lexical entry for *more or less* is given in (46):¹⁹

$$(46) \llbracket \text{more or less} \rrbracket^{\text{gran}} = \lambda f \in D_{(dt, et)} \lambda x \in D_e \exists D \in D_{dt} [D = \text{coarsest}(\text{gran})(\max(\text{domain}(f))) \ \& \ f(D)(x)]$$

Note that *more or less* must be adjacent to the scalar expression it is construed with, as is the case for *completely*.

(47) John arrived (#more or less) with (more or less) clean clothes.

Pretty has a use that is in many ways similar to *more or less*, but we still need to investigate it in more detail.

About is similar to *approximately* but can only be construed with numerals and temporal expressions as shown in (48):

- (48) a. about three, at about noon, at about midnight, at about the same time
 b. #about clean/open/north

We assume that *about* directly applies the coarsest granularity function to its complement. This results in the same semantic effect as *approximately* does, though the mechanism is slightly different. We assume that *approximately* is not blocked by *about* because the resulting interpretation is identical.

$$(49) \llbracket \text{about } D \rrbracket^{\text{gran}} = \text{coarsest}(\text{gran})(\llbracket D \rrbracket)$$

Some, weak disjunction, and approximative inversion in Russian are scalar approximators that can only be construed with numerals. (50a) illustrates the use of *some* we refer to, while (50b) illustrates what we call *weak disjunction* because the disjunction *or* must be unstressed for the approximator interpretation.

- (50) a. some fifteen boys
 b. He is gone for two or three days.

In German, the equivalent of weak disjunction does not involve disjunction at all, as illustrated in (51):

- (51) Er ist für zwei drei Tage weg.
 he is for two three days gone

Approximative inversion in Russian is illustrated by (52a). The postnominal position of the cardinal indicates that an approximative interpretation is intended.

- (52) a. Ja vstretil studentov desjat'.
 I met students ten
 'I met approximately 10 students'

- b. Ja vstretil desjat' studentov
 I met ten students
 'I met ten students'

We assume that all these approximators involve the same lexical entry as the one we gave for *about*. This is shown for approximative *some* in (53). For approximative inversion and weak disjunction we assume that the construction licenses a silent counterpart of approximative *some*.

$$(53) \llbracket \text{some } D \rrbracket^{\text{gran}} = \text{coarsest}(\text{gran}) \llbracket D \rrbracket$$

Partially has the same distribution as *completely* and *more or less*: as Kennedy and McNally (2005) have already observed, it can combine with predicates that introduce closed scales:²⁰

- (54) a. partially dry/pure/empty
 b. #partially impure/three/north/white
 c. #partially beef Stroganoff/a heap of wood

We perceive *partially open* to be equivalent to *not completely closed*. This follows from the lexical entry for *partially* in (55):

$$(55) \llbracket \text{partially} \rrbracket^{\text{gran}} = \lambda f \in D_{\langle dt, et \rangle} \lambda x \in D_e \\ \exists D \in D_{dt} [D = \text{finest}(\text{gran})(\min(\text{domain}(f))) \ \& \ \neg f(D)(x)]$$

6.4.4 Further potential scalar approximators

The framework we develop for approximative modifiers calls for some discussion of *almost*, *nearly*, and *barely*. While we think that the analysis of these expressions also needs to make appeal to granularity, we do not have a fully worked out proposal to offer at this point. Briefly consider *almost*. It seems that *almost* does involve scales, but the scale can often be a derived temporal scale as examples (56a) and (57a) illustrate (cf. Rapp and von Stechow 1996, Penka 2006). *Approximately* in (56b) and (57b) and the other scalar approximators we talked about above seem to be unable to access this scale.

- (56) a. John almost killed Harry. (*With intended interpretation: John might have killed Harry.*)
 b. #John approximately killed Harry.

- (57) a. Charles is almost King by now. (*With intended interpretation: Charles will be King soon.*)
 b. #Charles is approximately King by now.

In examples like (58), however, we can readily analyze *almost* as referring to the interval one below the maximum on the coarsest granularity. To write out this lexical entry for *almost* in (58), we use the notation $prev_{gran}(D)$ to denote the granularity interval on the scale of D that precedes D .

(58) The shirt is almost dry.

- (59) $[[almost]]^{gran} = \lambda f \in D_{(dt,et)} \lambda x \in D_e$
 $\exists D \in D_{dt}[D = coarsest(gran)(\max(\text{domain}(f))) \&$
 $f(prev_{gran}(D))(x)]$

The examples in (60) show that other scalar approximators can themselves be modified by *almost*. In this case, it seems that the resulting interpretation makes reference to the interval one below the maximum on the finest granularity.

- (60) The shirt was almost perfectly/completely dry.

6.5 Conclusion

In this chapter, we argued for the necessity to distinguish between two kinds of vagueness, scalar and epistemic. We based this conclusion on a study of the semantic and distributional properties of approximators like *exactly*, *completely*, *definitely*, *approximately*, *more or less*, and *maybe*. One kind of vagueness relates precise expressions to intervals: we argued for a granularity parameter. This parameter determines for expressions, like *five o'clock*, which denote points on a scale, an interval these expressions are associated with. Scalar approximators like *exactly*, *completely*, *approximately*, and *more or less* regulate the size of the granularity intervals. We argued in section 6.4 that the scalar approximators are further subdivided into endpoint and nonendpoint-oriented ones. The other kind of vagueness is genuinely epistemic, and epistemic expressions like *definitely* and *maybe* regulate this kind of vagueness.

One surprising result of our work is that *heap* and *bald*, two often cited examples of vague predicates, actually belong to two different classes of vague expressions: *heap* is epistemically vague, while *bald* must be scalarly vague since it combines with the scalar approximators *completely* and *more or less*.

There are a number of things we had to put aside for reasons of space and time in this chapter. This includes the analysis of the wider distribution *not exactly*, which, as we noted, is a phenomenon restricted to English, and of *not even approximately*, which seems cross-linguistically more frequent. Furthermore, we did not actually provide an analysis of the approximator *loosely speaking*. This is one case that Lasersohn (1999) discusses, but his analysis is incompatible with our framework for the analysis of vagueness. Therefore, we seem to be under an obligation to provide an alternative analysis. However, we argue instead that *loosely speaking* cannot straightforwardly be captured by either analysis because we believe that *loosely speaking* is speech-act oriented. Our reason for this assumption is provided by the data in (61): embedded under a speech-act verb in (61a), *loosely speaking* is part of the content of the embedded clause and interpreted from John's perspective. But, embedded under the nonspeech-act verb *think* in (61b), *loosely speaking* is interpreted from the speaker's perspective and not part of the content of the embedded clause. The speech act dependency is not captured by Lasersohn's (1999) proposal and at present is beyond the scope of our proposal as well.

- (61) a. John said that, loosely speaking, a heap of sand is on his table.
 b. John thinks that, loosely speaking, a heap of sand is on his table.

Approximators are a large class of expressions and many aspects of their semantic treatment are still not fully explicit here. Nevertheless, we hope to have shown that studying their distribution and meaning more closely can provide valuable data for understanding the nature of vagueness in language and the modeling of vagueness in compositional semantics.

Notes

* Portions of this work were presented at the 2008 conference on Vagueness at the ENS Paris, at the Tokyo Semantics Circle and at the Center for General Linguistics in Berlin and we are grateful to the audiences for their comments. In particular, we thank Manfred Bierwisch, Chris Kennedy, Stephanie Solt, Chris Tancredi, Makoto Kanazawa, Eric McCready, Gillian Ramchand, Chris Cummins, an anonymous reviewer, and the editors for their input on this chapter, and Lisa Hartmann for editorial work. Some results presented here were presented in a proceedings paper before (Sauerland and Stateva 2007). Needless to say, all remaining errors are solely our responsibility. This work was carried out within the Emmy-Noether Research Group *Interpretation von Quantoren*, which is funded by the German Research

Council (DFG SA 925/1-2) with additional support from the Project VAAG (DFG SA 925/4) within the ESF LogicCC Eurocores program.

1. The term *imprecision* in work by Pinkal (1995) and Kennedy (2007) partially overlaps with what we refer to as *scalar vagueness*. For example, *bald* is usually regarded as vague rather than imprecise. Hence, we introduce two new terms in this chapter. Our terminology also reflects the fact that we regard both phenomena as a kind of vagueness.
2. It may be that in comparatives like *after 6 o'clock* we access the precise meaning of *6 o'clock*, as Chris Cummins has suggested to us. Intuitively, there is some overlap between *at 6 o'clock* and *after 6 o'clock* – e.g. if the arrival was at 6:01, one could use either phrase to describe the time of the arrival. However, only further empirical investigation can confirm whether this is best captured by assuming that the comparative accesses the precise meaning of *6 o'clock*.
3. See also Barker (2002).
4. Epistemic approximators can also be used with expressions that are scalar vague as in (i). A purely approximating interpretation of (i) is available, though only in scenarios like the following: John volunteered to cook fifty tapas. He cooked a lot, but only exactly forty-nine tapas. Now we are discussing whether John honored the contract. If I believe that John's original utterance "I will make fifty tapas" implied only that we would make about fifty, I could use (i) to state my belief. Note though that this meaning is still quite different from *exactly fifty*.
 - (i) The number of tapas John cooked is definitely fifty.
5. In German, the expression *gerade* ("straight"/"just") approaches the distribution and meaning of *exactly*: the negated example (ii) is fully acceptable. In the positive example (iii), on the other hand, *gerade* can only be interpreted as a discourse particle, indicating an opposition between (iii) and a prior statement in the discourse.
 - (ii) Rotwein ist nicht gerade gesund.
red wine is not straight/just healthy.
 - (iii) Rotwein ist gerade gesund.
red wine is straight/just healthy.
6. Krifka shows that in some cases other considerations of cognitive efficiency override the use of the shortest expressions. For example, *18 months* denotes a bigger interval than *20 months* when describing a child's age.
7. For two- or more-dimensional concepts like *middle* a partial ordering needs to be defined appropriately.
8. The proposals of Barker (2002) and ongoing work by Lassiter (2009) seem to us to go in a similar direction as our approach.
9. Makoto Kanazawa (p.c.) points out that example (iv) also brings about an approximation interpretation and that, furthermore, there is an interesting contrast between (iv) and (v). The contrast is explained if a felicitous use of *the heap* requires that in all worlds of the common ground there must be

a salient referent for it, and furthermore presupposes that they are satisfied when evaluated against the common ground.

- (iv) This heap is definitely a heap.
- (v) #This heap is maybe not a heap.

10. Barker (2002) actually defines *clearly*, suggesting that this definition is appropriate also for *definitely*, by focusing on three degrees:

- (vi) $\iota(\max(\lambda d.c[d/\alpha] \in \alpha(x)))$, the maximal degree d to which an individual x has a property α in a world under consideration c .
- (vii) $\iota(\max(\lambda d.c[d/\alpha] \in C))$, the maximal degree d out of all standards of comparison for the property α in the worlds under consideration in a context C .
- (viii) $d(c)([[clearly]])$, the degree (standard of comparison) that results from applying a delineation function d to the meaning of the gradable adjective *clearly* in c .

For John to be definitely tall, it must be the case that the degree of John's tallness derived from (vi) must exceed the respective standard of comparison derived through (vii) by a certain amount defined by (viii).

11. Gillian Ramchand (p.c.) points out that in some dialects of British English *ish* can be used not just as an affix as in *greyish* or *beef Stroganoffish*, but also as an independent sentence-final morpheme. A Linguist List posting by Margaret Fleck from February 1992 points out the following example (<http://linguistlist.org/issues/3/3-129.html>):

- (ix) A: Is your algorithm working?
B: Yes, it's working. Ish.

The affix *-erly* seems to be restricted to compass directions as in *northerly* vs *northern* and, hence, seems to be scalar. Judging from the data reported in Siegel (2002), English *like* behaves also like an epistemic approximator in that it is not restricted to scalar predicates. Geoff Pullum, in a November 22, 2003 Language Log internet post, points out that *if you will* is similar in distribution and meaning to *like* (<http://itre.cis.upenn.edu/%7EEmyl/language/000138.html>).

12. An anonymous reviewer pointed out to us that examples like *Weigh exactly approximately 35 mg of a dried brass sample into a 250 mL Erlenmeyer flask* can be found with a Google search. However, all the examples of this type we found are in writings about chemistry like the example just cited. We surmise that this may be a special jargon within chemistry.

13. Lasersohn's (1999) paper does not address negation at all. In our view, negation is problematic for Lasersohn's "halo theory," but other monist accounts do not suffer from this problem. The differences between approximators with regard to stacking and other phenomena, however, are a problem for all monistic theories.

14. Our explanation is still incomplete in the following way: it is not clear how the approximator selects the appropriate scalar structure. We hope this will provide a direction for future research.
15. Technically, the value reported for a random distribution should be close to 1, but not necessarily exactly 1.
16. We put aside for now uses of *absolutely* with quantifiers like *every* and *no*. Furthermore, *absolutely* can combine with *zero*, while *completely* and *totally* cannot.
17. Eric McCready (p.c.) points out uses like *perfectly awful* and *perfectly horrible* where perfection is not related to desirability. We put these cases aside for now.
18. One exception to the complementarity of *approximately* and *more or less* is *more or less the same*. We put this aside for now.
19. As in the case of *absolutely*, we restrict attention in this chapter to uses of *more or less* with adjectives.
20. Manfred Bierwisch (p.c.) points out that the German counterpart of *partially*, *teilweise*, must always receive a mereological interpretation. A door that is *teilweise* open, must be a door that consists of at least two independently movable pieces, only one of which is open.

References

- Barker, C. (2002). The dynamics of vagueness. *Linguistics and Philosophy*, 25: 1–36.
- Fintel, Kai von and Gillies, Antony S. 2007. An opinionated guide to epistemic modality. In Szabó Gendler, T. and Hawthorne, J. (eds), *Oxford Studies in Epistemology*, vol. 2, pp. 32–62. Oxford University Press.
- Heim, I. and Kratzer, A. (1998). *Semantics in Generative Grammar*. Blackwell, Oxford, UK.
- Hobbs, J. and Kreinovich, V. (2006). Optimal choice of granularity in commonsense estimation: why half-orders of magnitude. Unpublished manuscript.
- Kamp, H. (1981). The paradox of the heap. In Mönnich, U. (ed.), *Aspects of Philosophical Logic: Some Logical Forays into Central Notions of Linguistics and Philosophy*, pp. 225–77. Reidel, Dordrecht.
- Kennedy, C. (2007). Vagueness and grammar: the semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy*, 30: 1–45.
- Kennedy, C. and McNally, L. (2005). Scale structure and the semantic typology of gradable predicates. *Language*, 81.
- Krifka, M. (2007). Approximate interpretation of number words: a case for strategic communication. In Bouma, G., Krämer, I. and Zwarts, J. (eds), *Cognitive Foundations of Communication*, pp. 111–26. Koninklijke Nederlandse Akademie van Wetenschappen.
- Lakoff, G. (1973). Hedges: a study in meaning criteria and the logic of fuzzy concepts. *Journal of Philosophical Logic*, 2: 458–508.
- Laserson, P. (1999). Pragmatic halos. *Language*, 75: 522–51.
- Lassiter, D. (2009). The algebraic structure of amounts: evidence from comparatives. In Icard, T. (ed.), *Proceedings of the 14th Student Session of ESSLLI*, pp. 51–61. ESSLLI, Bordeaux, France.
- Penka, D. (2006). Almost there: the meaning of *almost*. In *Proceedings of Sinn und Bedeutung 10*. Berlin.

- Pinkal, M. (1995). *Logic and the Lexicon*. Kluwer, Dordrecht, Netherlands.
- Rapp, I. and von Stechow, A. (1996). Fast “almost” and the visibility parameter for D-adverbs. *Journal of Semantics*, 16(2): 149–204.
- Rotstein, C. and Winter, Y. (2004). Total adjectives vs partial adjectives: scale structure and higher-order modifiers. *Natural Language Semantics*, 12: 259–88.
- Sauerland, U. and Stateva, P. (2007). Scalar vs epistemic vagueness: evidence from approximators. In Gibson, M. and Friedman, T. (eds), *Proceedings of SALT 17*, pp. 228–45. Ithaca, NY.
- Schwarz, B. (2006). Modal superlatives. In *Proceedings of SALT 15*. CLC Publications, Cornell University, Ithaca, NY.
- Siegel, M. (2002). *Like*: the discourse particle and semantics. *Journal of Semantics*, 19: 35–71.
- Williamson, T. (1994). *Vagueness*. Routledge, London, UK.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press, Oxford, UK.

7

Degree Modifiers and Monotonicity

Rick Nouwen*

7.1 Introduction

In a short squib, Zwicky (1970) wonders what could explain the assignment of certain adverbial functions.¹ He observes that there exist pairs of expressions where the positive of the pair is a sentence adverbial, while the negative one is a degree modifier. For instance, *unusually* in (1) is a degree adverbial. That is, the example expresses that the children are noisy to a degree that is unusual. No such similar reading is available if we replace *unusually* with its positive counterpart *usually*. The example in (2) is instead interpreted as saying that it is usual for the children to be noisy.

- (1) The children are unusually noisy.
- (2) The children are usually noisy.

This contrast appears a general feature of positive–negative pairs. Compare, for instance, *typically* with *atypically* or *possibly* with *impossibly*.

Another observation made by Zwicky in the same squib is that there is a resemblance between negative adverbs like *unusually* and evaluative predicates, like *surprising*, *amazing*, *terrible*, etc.² Adverbs based on such predicates also have degree functions. For instance, in its most salient reading, (3) expresses that the children are noisy to a degree that is amazing, not that the fact that they are noisy is amazing.

- (3) The children are amazingly noisy.

This is not to say that negative adverbs and adverbs based on evaluative predicates do not have an ad-sentential use. This becomes apparent by

using comma-intonation around the adverb in (1) and (3) or to place the adverbs in sentence-initial position. Note, however, that an example like (2) completely lacks a reading where the adverb is taken to modify degree. It is this contrast that I will try to explain in what follows. I will suggest that the reason behind it is semantic in nature and that my explanation is relevant to a broader range of degree modification phenomena.

In section 7.2, I will discuss Zwicky's observation in somewhat more detail, present several ways of extending it and put forward a descriptive generalization that captures the data. Section 7.3 discusses aspects of two approaches to evaluative degree modifiers like *surprisingly* and introduces a proposal for their semantics which, when extended to degree modifiers in general, accounts for Zwicky's observation. In section 7.4, I elaborate on this proposal by investigating the crucial role played by monotonicity. Finally, section 7.5 discusses some of the assumptions I make and gives pointers for further research.

7.2 Degree modification and markedness

Zwicky's observations are not restricted to English. In Dutch, for instance, we see a similar distribution of adverbial functions. One difference is that Dutch marks ad-sentential adverbs. Gradable adjectives like *verrassend* (surprising) combine with *genoeg* (enough) to form a sentence modifier *verrassend genoeg*. Only the short form *verrassend* can be used as a predicate modifier.³ Thus, *verrassend* in (4) is a degree modifier, while *verrassend genoeg* in (5) is a sentence modifier.

- (4) Jasper is *verrassend lang*.
 Jasper is surprising tall.
 "Jasper is surprisingly tall."
- (5) Jasper is *verrassend genoeg lang*.
 Jasper is surprising enough tall.
 "It is surprising that Jasper is tall."

The combination of a short form adverb and a nongradable expression is infelicitous.

- (6) Het is *verrassend* [#](*genoeg*) 3 uur.
 It is surprising (enough) 3 hour.
 Only: "It is surprisingly 3 o'clock."

This suggests that predicate-modifying adverbials need to operate on a degree argument. There are, however, additional, syntactic restrictions, as becomes apparent from the modification of nonadjectival gradable predicates. An indefinite like *a weirdo* can be degree-modified in various ways, but not by means of an evaluative degree adverb.

- (7) a. Jasper is very much a weirdo.
 b. Jasper is such a weirdo.
 c. Jasper is quite a weirdo.
 d. Jasper is more (of) a weirdo than Crazy Carl.
- (8) Jasper is surprisingly a weirdo.

The only reading available for (8) is that it is surprising that Jasper is a weirdo, not that Jasper is a weirdo to a surprising degree. We would have to assume that the position *surprisingly* in (8) is unavailable for degree adverbs. Such an assumption is supported by Dutch data.

- (9) Jasper is verrassend #(genoeg) een weirdo.
 Jasper is surprising enough a weirdo
 Only: "It is surprising that Jasper is a weirdo."

It thus looks like the unavailability of a degree reading for (8) is due to a syntactic rather than a semantic restriction. This is further supported by the fact that the adjectives the degree adverbs are based on can modify degree *within* a noun phrase. That is, Zwicky's observations extend to the degree-modifying function of adjectives themselves. Consider (10) and (11):

- (10) Jasper is an unbelievable weirdo.
 (11) Jasper is a believable weirdo.

In (10), a degree-modified reading is salient, where Jasper is said to be a weirdo to a high degree. No such degree reading is available for (11). Instead, it says that Jasper is a weirdo and that one can (easily) believe that he is that.⁴

It would thus seem that Zwicky's observation was not really about adverbs per se, but rather about what kind of predicates can be assigned a degree-modifying function in general. Let us have a closer look at the contrasting sets of expressions at issue. Zwicky mentions the following adverbs in his short squib:

(12) *Degree adverbs*: unusually, atypically, abnormally, uncharacteristically, impossibly, uncommonly, unnaturally, extraordinarily, particularly, especially, surprisingly, amazingly, disgustingly, alarmingly, bothersomely, shamefully, fantastically, incredibly, unbelievably, marvelously, dreadfully, awfully, preposterously, terribly

Sentence adverbs: usually, typically, normally, characteristically, possibly, commonly, naturally, ordinarily, generally

A simple observation, and as I will argue a crucial one, is that the degree adverbs are all based on predicates that express some form of markedness. Something that is *unusual*, or *atypical*, or *uncommon*, or *preposterous*, or *fantastic*, etc. will stand out in a way that *usual* things, *typical* things, or *common* things do not. I will call this the *markedness generalization*:

Markedness generalization: Degree modifiers tend to be based on predicates that express some form of markedness. In other words, the objects that satisfy these predicates, in some respect, stand out in their domain.

Note that, in contrast to the generalization that Zwicky seems to suggest, this means that two opposite predicates could both function as a degree modifier. For instance, *marvelous* and *terrible* seem each other's opposites, yet they both express a markedness feature and so they both have a degree function.

(13) Jasper is a terribly nice man.

(14) Polystyrene is marvelously useful.

Support for the markedness generalization comes from a few further observations on possible degree modifiers. Firstly, expletives like *fucking* or *damned*, by which a speaker may signal a marked emotional attitude to what is said, as in (15), have a role as a degree modifier, as in (16):

(15) Watch out! That's fucking dynamite!

(16) Lasagne takes a fucking long time to prepare.

Related to such expletives are interjections like *man* (McCready 2009), *boy*, or *gosh*. Such expressions also seem to be able to do two things: either mark a proposition in a certain way, or modify degree within that proposition. Take, for example, *man*. McCready (2009) observes that the core

function of *man* is to express the speaker's positive or negative emotion with respect to the proposition s/he is asserting. In (17), for instance, the speaker emphasizes what a good or bad thing (depending on his or her political inclination) it is that Obama got elected:

(17) Man, Obama got elected.

Interjections like *man* fit the markedness generalization. As (17) shows, their function is to express a form of markedness, and, as (18) shows, it turns out that at the same time they have a function as a degree modifier:

(18) Man it's hot.

There are two available readings for (18), depending on intonation. (See McCready 2009 for details.) With comma-intonation separating *man* from *it*, a reading surfaces in which the speaker expresses his or her emotional involvement in it being hot. Without comma-intonation, (18) expresses that it is hot to an intensified degree.

In general, the markedness generalization observes that operations that express the markedness of a proposition double as (degree) intensifiers. This double function is also familiar from exclamative intonation, which is standardly associated with the expression of attitudes like surprise, disbelief, elation, etc. The example in (19) is comparable to (17) and expresses a positive or negative emotion towards the fact that Obama got elected. In (20), however, it is the degree to which the pie is nice that is marked.

(19) Obama got elected!

(20) What a nice pie Jack baked!

In sum, I have shown in this section that Zwicky's observation can be generalized as a condition on expressions that act as degree modifiers, be it in adverbial, adjectival or some other form.

7.3 Analysis

In this section, I will introduce the main considerations that ultimately lead to my analysis of the data. I will do so by focusing on a single running example, (21):

(21) Jasper is surprisingly tall.

7.3.1 Predicate modification and sentence modification

It is important to be clear about how (21) is semantically different from an example where *surprisingly* is a sentence modifier, as in (22). Consider (23) as a paraphrase of (22):

(22) Surprisingly, Jasper is tall.

(23) *Jasper is tall and the speaker is surprised about this.*

For (21), such an analysis seems less promising. The distinguishing context is one in which I expected Jasper to be tall (with respect to the contextual standard of comparison), but in which Jasper is still taller than I expected. In such a context, (21) is true and (22) is false (cf. Morzycki 2008). In other words, (22) expresses surprise at the fact that Jasper is among the tall, while (21) expresses surprise with respect to how tall he is.

Another difference between the ad-sentential and ad-adjectival use of expressions like *surprisingly* has to do with the interpretation of the gradable adjective. Whereas *tall* in (22) is interpreted with respect to the standard of comparison, the interpretation of *tall* in (21) is much more complex. At first sight, it appears that it follows from (24a) that (24b):

(24) a. Jasper is surprisingly tall.

b. Jasper is tall.

However, as Katz (2005) first noticed, a similar entailment is absent from examples with absolute, rather than relative standards. For example, (25a) does not entail (25b):

(25) a. The lecture hall was surprisingly full.

b. The lecture hall was full.

Katz suggests that the examples in (24) are also not in an entailment relation, but that (24b) is more likely to have the status of implicature. If something is taller than expected, then it will very probably be taller than the standard of comparison given the major role expectation seems to play in establishing the standard. As support for this, Katz provides (26) (Katz 2005: 194). If *surprisingly tall* did entail *tall*, then such examples would be infelicitous.⁵

(26) Although he is quite short, Peter is surprisingly tall, given his background.

In what follows, I will assume, with Katz, that (24a) does not entail (24b).

7.3.2 Morzycki (2008)

Morzycki (2008) presents a thorough analysis of the syntax and semantics of degree modification by evaluative adverbs. His approach to the semantics of adverb–adjective combinations like the one in *Jasper is surprisingly tall* is based on an analogy with (embedded) exclamatives. That is, *Jasper is surprisingly tall* is likened to *It is surprising how tall Jasper is*. Based on the proposal of Zanuttini and Portner (2003) that exclamatives involve an operation of domain widening, Morzycki’s analysis of (27a) is (roughly) as in (27b):

- (27) a. Jasper is surprisingly tall.
 b. It is surprising that there is a degree to which Jasper is tall such that this degree is in a widened domain but not in the original domain.

According to Morzycki’s proposal, domain widening makes it the case that someone who is surprisingly tall is tall to a degree that is not contextually salient. Crudely put, it involves being tall to a degree that is off the scale. Domain widening is thus responsible for the intensifying function of the adverb. It is, however, not part of the lexical semantics of adverbs like *surprisingly*. Morzycki argues that such adverbs combine with adjectives through a mediating null degree morpheme, [R]. While *surprisingly* has a simple semantics like $\lambda p.\text{surprising}(p)$, the feature [R] is responsible for domain widening, and so it is ultimately [R] that is the intensifier. In my view, this aspect to Morzycki’s approach is a serious disadvantage. By disconnecting degree intensification from the lexical semantics of the adverb, it will be impossible to get an explanation of why only a certain kind of predicate can be turned into a degree adverb. The semantics I will propose, in contrast, attributes the degree-boosting role of adverbs like *surprisingly* directly to their lexical semantics.

7.3.3 Explaining the markedness generalization

An important observation in Morzycki (2008) is that there is a difference between being *surprisingly tall* and having a height that is surprising. Someone who was expected to be quite tall, but turns out shorter than expected, is clearly not *surprisingly tall*. Moreover, someone who was expected to be tall but not expected to be tall to his or her exact height is also not *surprisingly tall per se*. For instance, following an example from Morzycki (2008: 6), I might be surprised to find out that Jasper’s height in centimeters equals my bank account number, but that would not make me assert (28):

(28) Jasper is surprisingly tall.

In the approach of Katz (2005), examples like (28) are excluded from contexts in which it is merely Jasper's height that is surprising by quantifying over degrees. Katz's analysis of (28) is roughly as in (29):

(29) Jasper is tall to a degree d and every degree $d' \geq d$ is such that it would be surprising were Jasper tall to degree d' .

This says not only that Jasper's height is surprising, but moreover that, had he been taller, we would have been equally surprised.

Katz did not attempt to connect his proposal to Zwicky's observation. (Like Morzycki, Katz only considers evaluative adverbs.) However, I do think there is a way to exclude the degree modifier use of certain adverbs in this approach. For instance, if we attempt to generalize (29) to apply to an example like (30), we arrive at (31):

(30) Jasper is usually tall.

(31) Jasper is tall to a degree d and every degree $d' \geq d$ is such that it would be usual were Jasper tall to degree d' .

Clearly (31) is false, no matter how tall Jasper is, for it commits us (for instance) to finding it normal had Jasper been taller than anyone alive. In general, we could say that the reason we cannot use unmarked predicates as degree modifiers, is that such use would commit us to extending the unmarkedness to higher points on the scale. Even though 1 meter 75 is a usual height, this does not license (30), for according to (31) this entails that a version of Jasper that is 3 meters tall is also *usually tall*.

In many respects, this reasoning is already close to my final proposal. The main idea is that unmarked predicates are excluded from modifying degree because, as degree modifiers, they would license inferences that make the construction useless. However, as I will now explain, Katz's assumption that degree modification involves universal quantification over degrees is unnecessary.

7.3.4 The proposal in a nutshell

A major assumption I will make is that gradable predicates are *monotone*. To give an illustration, an effect of this assumption is that someone who is tall to a certain degree, is tall to all lower degrees too. That is, the set

of degrees to which tall people are tall includes all the degrees to which shorter people are tall.⁶

I believe that the explanation for Zwicky's observations lies in the inferences triggered by the adverbs about the degrees to which certain predicates hold. Say we call John *unusually tall*. What I propose this means is that there exists a degree to which John is tall that is not usual (for someone like John). Had John been taller, he would also have been tall to this unusual degree (by monotonicity), and so we infer that had John been taller, we would have also called him *unusually tall*. This is why we can only use *unusually tall* to refer to someone who is *taller than (what is considered) normal* and why we cannot use it to refer to someone who is just of *an unusual height*.

Now take an expression like *usually tall*, which lacks an interpretation of *usually* as a degree modifier. But, say, we try to interpret it like that anyway, and we claim John to be *usually tall*. This then means that John is tall to some degree which is "usual." The problem now is that such a statement is not informative. Given the assumption I made about gradable adjectives, above, anyone is tall to some usual degree. Take the minimal degree: everyone is tall to the minimal degree. Consequently, it is very usual to be tall to that degree. The result is that *usually tall* is a trivial property. This, I will argue in more detail below, is why adverbs like *usual* are not degree adverbs.

Crucial to this explanation is the role of inferences. Statements about degrees license downward-directed inferences. If John is tall to degree *d*, then he is also tall to any degree lower than *d*. An adverb like *unusually* reverses such inferences and thereby licenses inferences that are upwards directed: if John is unusually tall, then had he been taller, he would also be unusually tall. Evaluative predicates fit neatly in this reasoning. Take *surprisingly*. If John is *surprisingly tall* then this means that John is tall to some degree that was unexpected. Had John been taller, then he would still have been tall to this unexpected degree. As a consequence, his height would still be surprising. Thus, being surprisingly tall comes to mean *taller than expected*. An adverb like *expectedly* could not have a use as a degree modifier. We expect anyone to be tall to the minimum degree and, so, *being expectedly tall* is a trivial property.

Since I claim that the possibility of degree modification crucially depends on the monotonicity of gradable predicates, it is predicted that the contrasts observed by Zwicky disappear once we consider constructions that contain reference to specific degrees only. For instance, it is fine to say that Jasper's height is *usual*. Similarly, something can have a normal width, but that does not make it *normally wide*.

7.4 The proposal in detail

7.4.1 Adjective semantics

According to what Beck (2009) calls the *standard theory*, gradable adjectives are relations between individuals and degrees, that is of type $\langle d, \langle e, t \rangle \rangle$ (Cresswell 1976, von Stechow 1984, Heim 2000). One particular way of making this precise is as in (32):

$$(32) \quad \llbracket \text{tall}_w \rrbracket = \lambda d \lambda x. x \text{'s height in } w \geq d$$

Such a definition entails that such adjectives are monotone in the following sense:

$$(33) \quad \forall w, x, d, d' : \text{tall}_w(x, d) \ \& \ d' < d \rightarrow \text{tall}_w(x, d')$$

The result is that someone who measures 180 cm in height will have all the degrees of height that someone who measures 170 cm has.

In the absence of any form of degree modification, the degree slot of the adjective is saturated by a covert existential operator *pos*.

$$(34) \quad \llbracket \text{pos} \rrbracket = \lambda A. \lambda x. \exists d [A(x, d) \ \& \ d \geq \text{the contextual standard for } A]$$

In analogy to *pos*, I will assume that overt degree modifiers are of a similar type. However, I propose that degree modifiers start out as propositional modifiers of type $\langle \langle s, t \rangle, \langle s, t \rangle \rangle$. That is, an adverb like *surprisingly* is given the following semantics:

$$(35) \quad \llbracket \text{surprisingly} \rrbracket = \lambda p. \lambda w. p(w) \ \& \ \text{surprising}_w(p)$$

Adjectival modifiers are of type $\langle \langle d, \langle e, \langle s, t \rangle \rangle \rangle, \langle d, \langle e, \langle s, t \rangle \rangle \rangle \rangle$ and can be derived from propositional modifiers by means of a simple type shift Δ :

$$(36) \quad \Delta = \lambda \mathcal{P}_{\langle \langle s, t \rangle, \langle s, t \rangle \rangle} \lambda A_{\langle d, \langle e, \langle s, t \rangle \rangle \rangle} \cdot \lambda d_d \cdot \lambda x_e \cdot \mathcal{P}(A(x, d))$$

I furthermore assume that the degree variable is existentially closed after application of the modifier. (Alternatively, we could make this existential closure part of the semantics of the degree modifier, cf. Katz 2005, Morzycki 2008.) Here is a worked out example.

$$(37) \quad \llbracket \text{surprisingly tall} \rrbracket \\ = \Delta [\lambda p \lambda w. p(w) \ \& \ \text{surprising}_w(p)] [\lambda d. \lambda x. \lambda w. \text{tall}_w(x, d)]$$

$$\begin{aligned}
&\rightsquigarrow \lambda A.\lambda d.\lambda x.\lambda w.A(d, x)(w) \ \& \ \text{surprising}_w(A(d, x)) \\
&\quad \times [\lambda d.\lambda x.\lambda w.\text{tall}_w(x, d)] \\
&\rightsquigarrow \lambda d.\lambda x.\lambda w.\text{tall}_w(x, d) \ \& \ \text{surprising}_w(\lambda w'.\text{tall}_{w'}(x, d)) \\
&\rightsquigarrow (\text{existential closure}) \\
&\quad \lambda x.\lambda w.\exists d[\text{tall}_w(x, d) \ \& \ \text{surprising}_w(\lambda w'.\text{tall}_{w'}(x, d))] \\
(38) \quad &\llbracket \text{Jasper is surprisingly tall} \rrbracket \\
&= \lambda w.\exists d[\text{tall}_w(j, d) \ \& \ \text{surprising}_w(\lambda w'.\text{tall}_{w'}(j, d))]
\end{aligned}$$

At first sight, it might appear that the semantics in (38) is problematic for Morzycki's case of someone who happens to have a freakish height. Say it turns out that Jasper's height in centimeters is exactly my bank account number. Now it seems, against our intuitions, that (38) is true. Obviously, there exists a degree to which Jasper is tall such that it is surprising that Jasper is tall to that degree, namely the degree corresponding to Jasper's height.

It is instructive to provide a careful explanation of why (38) nevertheless provides the correct truth conditions. Call the degree corresponding to Jasper's height d_j . This is the degree, when expressed in terms of centimeters, that corresponds to my bank account number. Notice that there is a difference between (39a) and (39b):

$$\begin{aligned}
(39) \quad &\text{a. } \text{surprising}_w(\lambda w'. \max_d(\text{tall}_{w'}(j, d)) = d_j) \\
&\text{b. } \text{surprising}_w(\lambda w'. \text{tall}_{w'}(j, d_j))
\end{aligned}$$

In the situation described, (39a) is true, for (39a) states that it is surprising (in w) that Jasper's height is d_j . However, (39b) is false. This is because having d_j as your height is surprising, but not having d_j as one of your degrees of tallness. Imagine, for instance, that after having measured Jasper, we measure his neighbor and find out that he is 12 centimeters taller than Jasper. This entails that Jasper's neighbor is tall to degree d_j , just like Jasper is, but I think it is easy to agree that this fact is hardly cause for surprise.

7.4.2 Predictions

One result of the way the semantics was set up above is that we predict that the monotonicity characteristics of the predicate that acts as a degree modifier matter. Take a propositional operator O which is upward monotone in the sense that $p \rightarrow p' \Rightarrow O(p) \rightarrow O(p')$. Let T be a monotone degree relation (like that corresponding to *tall*). If we now construct a degree-modified property from T using O , we arrive at $\lambda x.\exists d[T(x, d) \ \& \ O(T(x, d))]$. (I am omitting world variables here for the sake of readability.) Given

the monotonicity of T , for any x , if x has a degree of T , then x has the bottom element on T 's scale as a degree of T -ness. Given the monotonicity of O , for any x , if for some d it holds that $O(T(x, d))$, it follows that x being T to the bottom element of scale satisfies O . Since anyone who has a degree of T -ness is T to the bottom element of the scale, it follows that the degree-modified property $\lambda x.\exists d[T(x, d) \ \& \ O(T(x, d))]$ is a nondiscriminant property. Consequently, combinations of upward monotone predicates and degree predicates are infelicitous.^{7,8}

To illustrate, take modal operators like *possible* and *necessary*. Clearly if p entails p' , then both $\diamond p$ entails $\diamond p'$ and $\square p$ entails $\square p'$. The proposition that x is tall to degree d entails that x is tall to the bottom element on the scale. Consequently, if it is possible/necessary that x is tall to some degree, then it is entailed that it is possible/necessary that s/he is tall to the bottom degree. As a result, anyone with a height is *possibly/necessarily tall* in the degree-modified sense. This explains why (40a) and (40b) lack a degree reading:

- (40) a. Jasper is possibly cute.
 b. Jasper is necessarily cruel.

In contrast, the examples in (41) do have a degree reading:

- (41) a. Jasper is impossibly cute.
 b. Jasper is unnecessarily cruel.

This can be explained by the fact that the inferences are reversed once we consider downward monotone operators. This allows such operators to construct discriminating properties in their role as degree modifiers. For instance, $\neg\diamond$ is downward monotone: if p is impossible, then anything that entails p will be impossible too. So, if it is deemed impossible that x is tall to degree d , then it is equally deemed impossible that x is tall to degree $d+n$. It could be, however, that x being tall to degree $d-n$ is not deemed impossible. The property $\lambda x.\exists d[T(x, d) \ \& \ \neg\diamond T(x, d)]$ thus describes individuals who map to some upper part of the scale associated with T .

I predict then that Zwicky's observation amounts to a contrast of monotonicity. Let me illustrate with some examples. In the pair *usual/unusual*, the former is upward monotone, while the latter is downward monotone. To see this, consider first of all the fact that (42a) entails (42b):

- (42) a. Jasper usually wears a black sweater to work.
 b. Jasper usually wears a sweater to work.

Since Jasper wearing a black sweater to work entails that Jasper wears a sweater to work, (42) shows that *usually* preserves this entailment, and consequently that it is upward monotone. In contrast, if it is unusual for Jasper to wear a sweater to work, then it will be also unusual for Jasper to wear a black sweater.

For the case of evaluative predicates, the downward monotonicity is observable in that they license negative polarity items.⁹ For instance, Kadmon and Landman (1993) observed the following:

- (43) a. She was amazed that there was any food left.
 b. I was surprised that he budged an inch.
 c. We were astounded that she lifted a finger.
- (44) #She expected that there was any food left.

It is not easy to check for the monotonicity characteristics of such predicates on the basis of entailment relations. For instance, *John read a boring book* entails that *John read a book* and so we expect *She was amazed that John read a book* to entail *She was amazed that John read a boring book*. This entailment does not arise, however, and one of the reasons is that these examples are factive. *She was amazed that John read a book* presupposes that John read a book. Similarly, *she was amazed that John read a boring book* presupposes that John read a boring book. But the latter presupposition is not entailed by the assertion that someone was amazed at John reading a book. What is needed for such cases, then, is the notion of *Strawson entailment*. This modification of standard entailment relations was proposed by von Stechow (1999) as a way of working around presuppositions in tests for monotonicity patterns. A sentence *S* Strawson-entails *S'* if and only if *S* entails *S'* on the premise that all presuppositions of *S'* hold.

Unfortunately, von Stechow's suggestion is not completely helpful with evaluative predicates like *surprising*, *amazing*, etc. This is because *surprise* can be directed at certain specific aspects of the world. For instance, following von Stechow, the following would have to hold if *surprise* were downward monotone:

- Premise 1: John read a boring book. (*by presupposition of the conclusion*)
- Premise 2: It is surprising that John read a book.
- Premise 3: John read a book \Leftarrow John read a boring book.
-
- Conclusion: It is surprising that John read a boring book.

Intuitions might not be clear about an example like this, for even if we did not expect John to read a book, we might know that John is an extremely boring fellow and that he has no taste whatsoever. So, if he were to read a book, we would actually expect him to read a boring one. This makes it the case that we will be hesitant to accept the conclusion of the above argument. Nevertheless, what is crucial here is that there is an element of truth in the conclusion. If the news that John read a book leads to surprise, then so would the news that he read a boring book, even though the boring aspect of this latter bit of news is expected.

Kadmon and Landman say about examples like this that there is evidence for downward monotonicity if we enforce a constant perspective throughout such tests. In their discussion they compare *I'm sorry he bought a car* to *I'm sorry he bought a Honda*, where the former should entail the latter if *sorry* were downward entailing. They explain:¹⁰

If I'm sorry he bought a car, I clearly wish he had bought no car. What ought I to feel, then, about his buying a Honda? I ought to be sorry about it, qua car. In fact, [...] I MUST be sorry about it, qua car. That is because refraining from buying a Honda is an absolute requirement for satisfying my wish. I cannot prefer for my wish to be satisfied in "another way." Hence, *sorry* is [downward monotone] (on a constant perspective). (Kadmon and Landman 1993: 383)

A similar reasoning exists for *surprise*. If I did not expect Jasper to read a book, then I will be surprised to find out that he did, irrespective of what kind of book he read. The upshot is that evaluative predicates are presumably downward monotone. According to my proposal, this explains why they can function as degree modifiers.¹¹

7.5 Discussion

In this short chapter, I have proposed a way of capturing the generalization that only a specific subset of predicates can be assigned a degree-manipulating function, a generalization that, to my knowledge, has never been under semantic scrutiny before. My approach assumes that, in principle, any predicate can modify degree, but that many such modifications would not yield an informative interpretation. The distribution of adverbial functions to predicates is governed by the inferences, in particular by the monotonicity inferences, triggered by the predicate. There are a few questions and connections I have left unattended, however. I will discuss some of them briefly.

To start, the adverbs *fortunately* and *unfortunately* are apparent counterexamples to the markedness generalization, for it is difficult to assign them a degree-modifying function:

(45) Jasper is fortunately tall.

(46) Jasper is unfortunately tall.

The preferred reading for (46) is not one where Jasper is so tall that it is unfortunate. Instead, it can only be interpreted as saying that it is unfortunate that he is tall. To make matters worse, there seems to be some regularity behind this exception, for the same exception occurs in Dutch:

(47) #Jasper is onfortuinlijk lang.

Jasper is unfortunate long.

Even if we were to find natural occurrences of *unfortunately* in a degree-modifying role, we would still need to explain what makes degree readings with this adverb generally dispreferred. I will leave such complex matters to further research.

A further topic that deserves attention is a phenomenon from the psychological literature called *framing*. So-called *framing* is exemplified by contrasts like that between (48) and (49) (Sanford et al. 2007). Even though *very few people died* and *a few people died* can be truthfully used in similar situations (say where there were five fatalities), their evaluative effects are very different:

(48) Very few people died, which is #terrible/marvelous.

(49) A few people died, which is a terrible/#marvelous.

The semantic mechanism behind the typology of degree-modifying predicates that I presented above was based on the inferences such predicates give rise to. Interestingly, this makes the proposal applicable to framing. In examples like (48) and (49) monotonicity properties might provide an explanation. Given that *a few* is upward monotone, it follows that *a few people died* would have remained true had more people died. Conversely, *very few people died* remains true in case fewer people die. In line with the reasoning explained above, evaluating the fact that *a few people died* as *marvelous* is infelicitous because it triggers the inference that the speaker would have found it marvelous too had more people died. Similarly, calling the fact that very few people died *terrible* commits the speaker

to finding it terrible had nobody died. See Geurts (2010) for a recent account of framing that, although arguing that the semantic mechanism behind framing is more intricate than simple monotonicity, is based on inferences in much the same way as the above explanation of the typology of degree modifiers is.

A more general issue has to do with the characterization of those predicates that can act as degree modifiers. A further crucial difference between degree-modifying adverbs and adverbs that lack such a function is that the former but not the latter are factive.¹² This is illustrated by the ad-sentential use of both classes of adverbs: (50a) entails that the children are noisy, (50b) does not:

- (50) a. Unusually, the children are noisy.
 b. Usually, the children are noisy.

As I discussed in section 7.3.2, Morzycki (2008) considers combinations like *surprisingly noisy* as analogous to embedded exclamatives. In the approach of Zanuttini and Portner (2003), exclamatives are assumed to be factive (as well as involving a mechanism of domain widening). However, this connection by itself does not *explain* why nonfactive adverbs cannot modify degree; it is merely part of the observation that factivity appears to be somehow relevant. The factivity of exclamatives and degree modifiers like *surprisingly* and *unusually* form a parallel generalization to the markedness generalization I made in section 7.2. Apparently, markedness signals tend to be factive predicates (but, obviously, not vice versa). For instance, emotives like *man* or *gosh* are clearly factive, while adverbs that do not fit the markedness generalization like *allegedly* are not.

Another issue that I have left open is how the discussion above relates to expressions like *clearly*, an adverb that has earned a prominent place in the vagueness literature (cf. Cohen and Wolf, this volume; Barker, this volume). This adverb, however, behaves more like a sentence adverb than like a degree modifier. Truth conditionally, there appears to be no difference, for instance, between saying *Clearly, Jasper is tall* and *Jasper is clearly tall*. Furthermore, *clearly* can precede intensified adjectives, as in *Jasper is clearly very tall*, while evaluative adverbs cannot: *#Jasper is surprisingly very tall*. Finally, being *clearly tall* is not about clarity with respect to the degree of tallness, but rather about clarity regarding the category membership. So, saying *Jasper is clearly tall* is not the same as saying that it is clear *how tall Jasper is*.¹³ If we were to try to interpret *clearly* as an adverb of degree in my theory, then *Jasper is clearly tall* would end up meaning that there is a degree to which Jasper is tall such that it is clear that Jasper

is tall to that degree. To be sure, this could be true even when it is not clear how tall Jasper is. Since Jasper is human, it is clear that he is taller than 5 centimeters. Thus, there is a degree to which he is tall, for which it is clear that he is tall to that degree. This, however, would make anyone *clearly tall*. On a degree modifier reading, in other words, *clearly* parallels *usually* and *possibly* in its inability to form discriminant properties.

Notes

* This work grew out of my 2005 Amsterdam Colloquium paper (Nouwen 2005). Since then, I have presented versions at several events in the Netherlands and ultimately at the Vagueness and Language Use workshop in Paris in April 2008. My thanks go to the audiences of these talks for stimulating discussion. In particular, I would like to thank Paul Égré, Bart Geurts, Nathan Klinedinst and an anonymous reviewer for comments on this chapter. This work was supported by a grant from the Netherlands Organization for Scientific Research (NWO), which I hereby gratefully acknowledge.

1. Thanks to Chris Kennedy for directing my attention to this piece.
2. The expressions I call evaluatives here were called “psychological predicates” in Zwicky (1970).
3. Ad-sentential adverbs are derived from nongradable adjectives by the suffix *-erwijze*. For instance, the sentence modifier *mogelijk-erwijze* (possibly) is based on the adjective *mogelijk* (possible).
4. There is a considerable amount of idiosyncrasy in the use of adjectives to modify degree. For instance, it is not easy to get a degree-modified reading for something like *Jasper is an unusual weirdo*. However, in line with Zwicky’s observations, only those adjectives that have degree-modifying adverbial counterparts are such that they are capable of modifying degree within a noun phrase.
5. I do not think that examples such as (26) provide definitive evidence against *surprisingly* modification entailing the positive, for it seems to me that riders like *given his background* could be interpreted as shifting the comparison class. Still, there is a clear contrast between (26) and (i).

(i) ??Although he is quite short, Peter is tall, given his background.

In general, it seems to me that, given the context-dependence of vague predicates, we lack a fully reliable method for testing whether occurrences of such predicates are interpreted akin to the positive form or not.

6. Due to this assumption, the present proposal is rather tied to a particular framework, namely the degree semantics approach to adjectives. (In particular, the approach that takes adjectives to be relations between individuals and degrees.) It remains to be seen whether a similar approach could be successful in, say, a delineation approach such as that of Klein (1980). However, such approaches often end up stating constraints on gradable predicates that are reminiscent of monotonicity. See for instance Doetjes et al. (2008) for a proposal for degree semantics that lacks degree arguments but is otherwise very

close to the assumptions I am making in this chapter. I will leave a detailed comparison to further research, however.

7. See Morzycki (2009: 197) for related reasoning with respect to a specific group of adjectival modifiers. As Morzycki notices, some size adjectives can modify degree, but only those that express *largeness*, as opposed to *shortness*:

- (ii) a. He is a {big/huge/enormous/gigantic/mammoth} idiot.
 b. #He is a {small/tiny/minute/microscopic} idiot.

There is an obvious and interesting connection to Zwicky's observations. Here, too, the function of degree modification is assigned only to one part of a polar division of predicates. However, adjectives like *big* and *small* differ from predicates like *surprising* and *usual* in that they lack propositional uses. I will therefore leave the interesting question of how a contrast like (ii) fits in my proposal for further research.

8. For the proposal to work, it is important to assume a strong relation between *uninformativeness* and *ungrammaticality*. I will refrain from spelling out the details of this relation, but see Gajewski (2002) and Fox and Hackl (2006) for discussion.
 9. Note that nonevaluative downward monotone predicates do not license negative polarity items, as shown in (iii):

- (iii) #It is unusual that he budged an inch.

I do not have an explanation for why this is the case.

10. Bart Geurts (p.c.) suggested to me that a more straightforward way of testing for monotonicity with these predicates is by using sentences that eliminate factivity and minimize the role perspectives might play. For instance, the test would be whether *It would be surprising if John found 10 marbles* entails that *It would be surprising if John found 11 marbles*.
 11. The proposal further predicts that nonmonotone operators, like downward monotone operators, form discriminating properties if they are used as degree modifiers of gradable predicates. Bart Geurts (p.c.) suggested to me that adverbs like *pleasantly*, *cosily*, etc. would be a case in point, as in (51):

- (iv) a. The water is pleasantly warm.
 b. The city centre of Utrecht is cosily small.

We might also find support in cases of modification by size adjectives (see note 8):

- (v) I'm an enormous fan of Motörhead, (upward)
 a. ...#but a small fan of Mötley Crüe. (downward)
 b. ...but only a small fan of Mötley Crüe (nonmonotone)

12. Thanks to an anonymous reviewer for stressing this point.
 13. Thanks to Paul Égré for suggesting these data to me.

References

- Beck, S. (2009). Comparison constructions. In Maienborn, C., von Stechow, P. and Portner, P. (eds), *Semantics: an International Handbook of Natural Language Meaning*. Mouton de Gruyter, Berlin. Publication in preparation.
- Cresswell, M. (1976). The semantics of degree. In Partee, B. (ed.), *Montague Grammar*, pp. 261–92. Academic Press.
- Doetjes, J., Constantinescu, C. and Součková, K. (2008). A neo-Kleinian approach to comparatives. In Ito, S. and Cormany, E. (eds), *Proceedings of Semantics and Linguistic Theory XIX*. Ithaca, New York.
- Fox, D. and Hackl, M. (2006). The universal density of measurement. *Linguistics and Philosophy*, 29(5): 537–86.
- Gajewski, J. (2002). *L-Analyticity in Natural Language*. MIT.
- Geurts, B. (2010). Goodness. In Aloni, M., Bastiaanse, H., de Jager, T., van Ormondt, P. and Schulz, K. (eds), *Proceedings of the 17th Amsterdam Colloquium*, pp. 277–85.
- Heim, I. (2000). Degree operators and scope. In *Proceedings of SALT 10*. CLC Publications, Ithaca, NY.
- Kadmon, N. and Landman, F. (1993). Any. *Linguistics and Philosophy*, 16: 353–422.
- Katz, G. (2005). Attitudes toward degrees. In Maier, E., Bary, C. and Huitink, J. (eds), *Proceedings of SuB 9*. Radboud Universiteit Nijmegen.
- Klein, E. (1980). A semantics for positive and comparative adjectives. *Linguistics and Philosophy*, 4: 1–45.
- McCready, E. (2009). What man does. *Linguistics and Philosophy*, 31(6): 671–724.
- Morzycki, M. (2008). Adverbial modification of adjectives: evaluatives and a little beyond. In Dölling, J. and Heyde-Zybatow, T. (eds), *Event Structures in Linguistic Form and Interpretation*. Mouton de Gruyter, Berlin.
- Morzycki, M. (2009). Degree modification of gradable nouns: size adjectives and adnominal degree morphemes. *Natural Language Semantics*, 17(2): 175–203.
- Nouwen, R. (2005). Monotone amazement. In Dekker, P. and Franke, M. (eds), pp. 167–72. *Proceedings of the Fifteenth Amsterdam Colloquium*. ILLC.
- Sanford, A., Dawydiak, E. and Moxey, L. (2007). A unified account of quantifier perspective effects in discourse. *Discourse Processes*, 44(1): 1–32.
- von Stechow, K. (1999). NPI licensing, Strawson entailment, and context dependency. *Journal of Semantics*, 16(2): 97–148.
- von Stechow, A. (1984). Comparing semantic theories of comparison. *Journal of Semantics*, 3: 1–77.
- Zanuttini, R. and Portner, P. (2003). Exclamative clauses: at the syntax–semantics interface. *Language*, 79(1): 39–81.
- Zwicky, A. (1970). *Usually and unusually*. *Linguistic Inquiry*, 1(1): 145.

8

Clarity as Objectivized Belief*

Lavi Wolf and Ariel Cohen

Ben-Gurion University of the Negev, Israel

8.1 Clarity and belief

Intuitively, if something is clear, then it is not vague. As the dictionary (*Collins COBUILD Advanced Learner's English Dictionary*) succinctly puts it: “vague \neq clear.” But this is of course quite compatible with the fact that *clear* itself is a vague predicate. As Barker and Taranto (2003: 17) point out:

the vagueness of *clear* is easy to prove, since it is possible to explicitly talk about the degree to which a proposition is clear.

- [(1)] a. It is becoming clear that Mary is a doctor.
b. It is reasonably clear that Mary is a doctor.
c. It is very clear that Mary is a doctor.
d. It is painfully clear that Mary is a doctor.

Barker (2009) argues further for the vagueness of clarity by demonstrating that it is susceptible to the sorites paradox.

But what does clarity mean?

Barker and Taranto argue that *It is clear that p* does not, in fact, entail *p*, and take this fact as evidence for basing the interpretation of clarity on belief. Specifically, they propose that (2) is true iff the degree of belief of the speaker and hearer in the truth of the proposition that Abby is a doctor is greater than some vague standard.¹

- (2) It is clear that Abby is a doctor.

Their explanation for the fact that (2) does not entail (3) is simple: believing something does not make it true.

- (3) Abby is a doctor.

8.2 Clarity and justification

In a later paper, Barker (2009) argues that belief is not, in fact, appropriate as an account of clarity. He points out that belief is neither a necessary nor a sufficient condition for clarity.

Belief is not necessary for clarity, according to Barker, since both speaker and hearer might believe that there is life on Mars, yet assent to (4):

(4) It is reasonably clear that Mars is barren of life.

Belief is not a sufficient condition either, because both speaker and hearer may believe in the existence of God, yet still deny (5):

(5) It is clear that God exists.

Instead of a theory based on belief, Barker proposes an account based on *justification*. Unlike belief, justification is objective: given some body of evidence, and some relevant standards, it is an objective fact whether some conclusion is justified.

Barker's formalization of his theory is inspired by Kratzer's (1991 and elsewhere) account of epistemic modals. According to Kratzer:

(6) *must p* is true iff *p* holds in all epistemically accessible worlds whose "degree of normality" is sufficiently high.²

However, Barker notes that clarity is not the same as epistemic modality. One difference between clarity and *must* that is particularly relevant here is that the former, but not the latter, requires the relevant knowledge to be publicly available. Suppose there are no publicly available facts indicating that Abby is a doctor, but the speaker has some relevant private knowledge about her profession, say because she has seen the medical tools Abby is carrying in her briefcase. In this case (7) would be perfectly acceptable, in fact true.

(7) Abby must be a doctor.

In contrast, Barker claims that (2) would be false in the situation described.

We agree with Barker that (2) would not be unproblematically true in such a case, though, in our judgment, (2) would actually be odd rather

than simply false, so we believe that the requirement for the evidence to be publicly available is a presupposition (we will return to this issue in section 8.6.1.3 below).

Given the difference between clarity and epistemic modality, Barker modifies Kratzer's definition as follows. Rather than expressing quantification over epistemically possible worlds, he proposes that clarity expresses quantification over worlds compatible with publicly available evidence.

In order to account for the vagueness of clarity,³ he proposes a vague *standard of skepticism*, $d(\text{clear})$. According to Barker's proposed definition, (2) is true iff in all worlds consistent with publicly available evidence whose "degree of normality" is greater than $d(\text{clear})$, Abby is a doctor. This formalization captures Barker's suggestion that (2) is true iff it is justified, on the basis of publicly available evidence, to conclude that Abby is a doctor.

8.3 Back to belief?

Barker uses examples (4) and (5) to argue that belief is too subjective to account for clarity. We agree with his judgments concerning these examples: indeed, a very similar (if slightly weaker) example is attested:⁴

- (8) Most people, including most theists, agree that even if God exists, it is not clear that God exists.

However, there are at least two reasons why an account based on belief is attractive nonetheless: personal and adverbial clarity.

8.3.1 Personal clarity

Barker distinguishes between *simple* clarity, exemplified by (2), and *personal* clarity, as in (9):

- (9) It is clear to me/you/John that Abby is a doctor.

Of course, something can be clear to one but not clear to another. But what does it mean for some proposition to be clear to X but not to Y?

Barker claims that this can happen in one of two cases. One such case is when X and Y have access to different bodies of (publicly available) evidence. Although the evidence is publicly available, it is still possible that X was exposed to it, but Y was not. In this case, X may be justified in concluding that Abby is a doctor, while Y may not be.

The second case in which, according to Barker, a proposition may be clear to X but not to Y, is when X and Y have different standards of skepticism, i.e. different values for the threshold $d(\text{clear})$. X may be credulous, while Y may be harder to convince, and this is why a proposition would be clear to one but not to the other.

However, these cases do not exhaust the space of possibilities. Suppose both Alex and Bill see a photograph of Abby wearing a stethoscope and smoking, and they know nothing else about her: the photograph consists of all their relevant knowledge concerning Abby. The following exchange is not unnatural, and neither of them can be accused of uttering a falsehood:

- (10) Alex: It is clear to me that Abby is a doctor (because she is wearing a stethoscope).
 Bill: It is clear to me that she's not (because she has an unhealthy habit).

Note that, by hypothesis, Alex and Bill share the same publicly available evidence – the photograph. Thus, their disagreement cannot be explained by a difference in the publicly available evidence that Alex and Bill have access to.

Can their disagreement be explained as a difference in their respective standards of skepticisms? Well, it is possible that they might differ on their respective values for $d(\text{clear})$, but this still would not explain their disagreement. Suppose Bill is more skeptical than Alex; then the proposition that Abby is a doctor may not be clear to Bill, but its negation will not be clear to him either. Hence, he may be justified in uttering (11), but not (10).

- (11) It is not clear to me that Abby is a doctor.

How, then, can we explain the disagreement between Alex and Bill? Note again that neither of them can be accused of uttering a falsehood; this means that, in general, a disagreement on clarity cannot really be resolved.

Consider the following attested example, in which Kyburg (1976: 366) discusses his disagreement with Levi:

- (12) We have here a simple conflict of intuitions, and in the last analysis it may be that you pays your money and you takes your choice. It seems quite clear to me that [Kyburg's claim]. The opposite is quite clear to Levi.

Barker notes this property of personal clarity, and points out its similarity, in this respect, to predicates of personal taste: disagreement on taste is a *faultless disagreement* (Lasersohn 2005; see also Kölbel 2003). In the following discourse, just as in (10), it is hard to accuse any of the interlocutors of uttering a falsehood:

- (13) Alex: This chili is tasty.
 Bill: This chili isn't tasty.

Barker argues that, in both (10) and (13), the disagreement comes from the fact that Alex and Bill have different relevant standards: "Even assuming that two discourse participants agree on the facts in the world, including the degree of tallness, clarity, or tastiness of some object, they can still differ on what they consider to be the prevailing vague standard" (2009: 270).

However, degrees of tastiness or clarity, unlike degrees of tallness, are not "facts in the world." Thus, two people may agree on Mary's height, but disagree on whether this height is above the cutoff point for being considered "tall"; in contrast, if they agree on the degree of tastiness, fun, or clarity of something, they cannot disagree on whether it is tasty, fun, or clear.

Lasersohn (2008: 308) makes this point explicitly:

What seems crucial for disagreements over taste is not the location of the cutoff point, but the assignment of degrees. Different people may assign markedly different degrees of fun or tastiness to the same items, and may differ radically in the relative order of these items on the fun or tastiness scale; but no objective "matter of fact" would seem to select any one of these assignments or orderings as the correct one. John and Mary may disagree whether skydiving is fun, not because they both realize it is fun to degree d and differ as to whether d is sufficiently high to count as fun, but because John (who enjoys a good thrill) assigns it a high degree of fun, while Mary (who is terrified of falling) does the opposite.

In other words, what Alex and Bill disagree on, when they disagree on tastiness or clarity, is not merely the defining standard (often called *the cutoff point*) of the vague predicates *tasty* or *clear*, but on the degrees of tastiness or clarity assigned.

We conclude, then, that Barker's justification-based account does not really explain the puzzle of personal clarity. In contrast, a belief-based account can provide a natural explanation (cf. Barker and Taranto

2003): Alex and Bill reason in different ways, and come to believe different things on the basis of the same evidence.

The meaning of personal clarity, then, is based on belief. It would be strange indeed if the meaning of simple clarity were based on radically different principles; hence the meaning of simple clarity ought to be based on belief too. An account of clarity, then, both simple and personal, must make crucial use of belief.

8.3.2 Adverbial vs adjectival clarity

Barker uses adverbial (*clearly*) and adjectival (*it is clear that*) clarity interchangeably, and does not distinguish between their meanings. But it turns out that subtle differences between the two do exist.

Piñon (2006) discusses adverbial and adjectival modals, and he notes three differences between them. If we apply Piñon's observations to clarity, it turns out that adverbial and adjectival expressions of clarity behave differently. Moreover, *clearly* patterns with modal adverbs, while *clear* patterns with modal adjectives.

The first difference discussed by Piñon is the fact that modal adjectives can be negated, while modal adverbs cannot, as shown in (14). The examples in (15) demonstrate that exactly the same can be said about clarity:

- (14) a. It's improbable/impossible that Abby is a doctor.
 b. *Improbably/Impossibly, Abby is a doctor.
- (15) a. It's unclear whether Abby is a doctor.
 b. *Unclearly, Abby is a doctor.

The second difference involves conditionals: modal adjectives, but not modal adverbs, can occur in the protasis of a conditional. Again, the same applies to clarity:

- (16) a. If it's possible/probable/clear that Abby will be a doctor, then I, too, should apply to medical school.
 b. *If Abby will possibly/probably/clearly be a doctor, then I, too, should apply to medical school.

The third difference concerns questions: modal adverbs are generally bad in questions, while modal adjectives are markedly better. Once again, *clearly* patterns with modal adverbs, *clear* patterns with modal adjectives:

- (17) a. Is it possible/probable/clear that Abby is the best doctor in the hospital?
 b. ??Is Abby possibly/probably/clearly the best doctor in the hospital?⁵

In order to account for these differences, Piñon suggests that modal adverbs modify the strength of the speaker's belief in the assertion (the sincerity condition of assertion, see Vanderveken 1990–91), rather than its content.

Additional evidence for his claim comes from the fact that overt indicators of illocutionary strength, e.g. *certainly* or *presumably*, preclude the use of modal adverbs, but not modal adjectives (cf. Krifka 2007):

- (18) a. Certainly/presumably it is possible/probable/clear that Abby is a doctor.
 b. *Certainly/presumably Abby is possibly/probably/clearly a doctor.

We accept Piñon's account and apply it to clarity. But if *clearly* modifies the speaker's belief in the assertion, its account must make essential use of belief. And since the meaning of *clear*, while not the same as *clearly*, is obviously related to it, its account must involve belief too. Hence, both adverbial and adjectival clarity must be based on belief.

Let us put this point somewhat differently: even if Barker's theory successfully accounted for the truth conditions of clarity, it is hard to see how it would be able to explain the differences between adjectival and adverbial clarity. If we are right and this difference depends crucially on belief, it follows that an account of clarity must involve belief.

8.3.3 Desiderata

We conclude, therefore, that since there are good reasons to propose a theory that uniformly accounts for both simple and personal clarity, and since there are also good reasons to account uniformly for both *clear* and *clearly*, the desired theory must be based on belief.

However, we cannot afford to ignore Barker's persuasive arguments that belief is too subjective to explain clarity. We seem to be facing a dilemma: our account must be based on belief, yet maintain the sense of objectivity that is provided by Barker's justification-based theory. In other words, we need to "objectivize" belief.

How can this be done? To answer this question, let us look more closely into the meaning of belief.

8.4 Belief and probability

8.4.1 The lottery paradox

We adopt the Bayesian view of belief, and interpret degrees of belief in terms of probability. More precisely, the degree of belief in ϕ is the probability of ϕ over an epistemic modal base.

Things are more complicated, however. Barker (2009: 264), in fact, considers the possibility of accounting for clarity in terms of probability, and rejects it:

Certainly clarity does not involve degrees of probability, as shown by the standard lottery scenario (Kyburg 1961). If your chances of winning a lottery are inversely proportional to the total number of tickets, I cannot assert *You will lose* (at least, I can't assert it on the basis of knowledge), no matter how many other lottery tickets have been sold. Nor can I appropriately say *It is clear that you will lose*, even if you and I both believe that you will lose.

Indeed, any theory that treats belief as a degree of probability needs to face up to Kyburg's lottery paradox. Suppose there are n tickets in a fair lottery, where n is a large number. Then, the statement that ticket #1 will lose is extremely probable, hence can reasonably be believed. The same holds for the statement that ticket #2 will lose, and so on, all the way to ticket # n . However, the conjunction of all those beliefs is equivalent to saying that no ticket will win, and this, of course, is false and cannot rationally be believed.

This paradox is usually taken to indicate the difference between *belief* and *acceptance*. While belief is graded, and one can believe a proposition to a high or low degree, acceptance of a proposition is all-or-nothing, and while dependent on belief, is not determined by it.

Taking this perspective, the lottery paradox is not without a solution. In fact, several well-known solutions to the paradox have been proposed. This chapter is not the place to debate the relative merits of these solutions, but whichever solution is adopted, one should be able to apply it to a belief-based theory of clarity, hence overcoming Barker's objection.

For example, according to several popular approaches to the paradox (e.g. Pollock 1995, Ryan 1996, Douven 2002) acceptance is *defeasible*: a high probability statement is accepted unless it is defeated by some other proposition (or set of propositions) that is at least equally probable. In the case of the paradox this means that the belief that ticket # i will

lose, while highly probable, will be defeated by the equally probable propositions that other tickets will lose and the (nondefeasible) knowledge that one ticket will win; consequently, this statement will not be accepted.

If we make the plausible assumption that assertion presupposes acceptance, and, in particular, assertion of clarity presupposes acceptance, it follows that, indeed, both *You will lose* and *It is clear that you will lose* will be unassertable.

Having a high level of belief, then, is a necessary condition for expressing clarity, but not a sufficient one. *It is clear that ϕ* may be unassertable if ϕ is not accepted: if there is a belief ϕ' that is incompatible with ϕ but at least as likely as it.

Note that when the statement is not defeated, its high probability is sufficient for acceptability, and clarity is assertable, as attested on the Web:⁶

- (19) Stop Losing Money at the Casino and Instantly Discover PROVEN WINNING Formulas ... If you don't have this knowledge, it is clear that you will lose.

In this context, the sentence is a habitual: it does not say that the hearer will lose on a specific occasion, but rather that the hearer will lose in general. As is well known, habituals allow exceptions: thus, the sentence allows for the possibility of an occasional win. Therefore, the statement is not defeated: it is not contradictory to say that *all* gamblers will habitually lose, even if on every lottery some gambler is guaranteed to win.

Consequently, (19) is assertable, in fact true: since the odds are always in favor of the house, the probability that if you do not use the "proven winning formulas" you will habitually lose is quite high.⁷

8.4.2 Probability and normality

Recall that Barker's theory relies crucially on the notion of a "normal" world. One of the problems with normality-based theories is that it is notoriously hard to define what counts as normal. Consequently, it is very hard to know what the predictions of a normality-based theory are, in order to confirm or refute it.⁸ Still, we can try, following the literature on normality, to make some reasonable assumptions.

Barker makes it clear that normality and high probability are distinct. Granted, what is normal is usually also probable, and what is probable is usually also normal. But the two can diverge, and something can be normal yet unlikely, or vice versa.

For example, consider how Barker accounts for the lottery scenario. He argues that (2009: 269)

in the lottery example, *It is clear that you will lose* is correctly predicted to be false.⁹ The reason is that that world in which you improbably win the lottery is just as normal as the many worlds in which you do not win. It follows that there is no way for there to be only losing worlds above the clarity threshold.

The same ought to apply to (19): the worlds in which you do not lose are much less likely than worlds in which you do, but just as normal. Hence, the sentence would wrongly be predicted to be false.

Suppose, for another example, that two students submitted identical exams. It is, in fact, possible that this was a coincidence. But while coincidence is, by definition, improbable, it is not abnormal. A normal world, in which both students worked in perfectly normal ways on their exams, does not become less normal if, by coincidence, their answers happen to be identical. Thus, Barker's approach would predict that (20) is false:

(20) Clearly/It is clear that you cheated on the exam.

According to our judgment, however, (20) is true, and we would be quite justified in saying it to the students.

One of the characteristics of normal worlds that is usually agreed on is that, in such worlds, everything happens as it should. For example, machines perform as directed. Indeed, Krifka et al. (1995), who favorably discuss a normality-based approach to another phenomenon (genericity), interpret sentences involving a machine "with respect to a modal base and ordering source, where the machine performs the action for which it was designed" (p. 54).

There are cases, however, when, based on our past experience, it is very likely that a machine will *not* function in such a "normal" manner. In such cases, Barker would predict that we cannot use *clearly*, yet such use is attested. Consider the following passage from a post on The DVD Forum:¹⁰

(21) Bitterly disappointed – my Xbox has started crashing quite regularly. I'm not sure what to do now – this is the 2nd Xbox I'll have to send back in 10 months ... It will clearly go wrong again.

Although it is not a normal thing for machines to break down regularly, the speaker believes, based on past experiences, that it is likely. Of course,

one might claim that, in this context, worlds in which machines break down somehow *are* normal. But this would be tantamount to stripping the notion of normality of any meaning: if anything can be considered normal, then a theory of normality has no predictive power.

8.4.3 Probability and logic

In our formalization, we need to incorporate probability into our logical formalism. We use a variation on Halpern's (1990) logic, which combines first-order logic with probability. A probability structure is a tuple $\langle D, W, \pi, F \rangle$. D is a domain and W is a set of possible worlds. π is a valuation function such that for each world $w \in W$, $\pi(w)$ assigns to the symbols of the language appropriate extensions. $F = \{f_1, f_2, \dots\}$ is a set of discrete probability functions over W .

We introduce distinguished propositional functions $P_i(\phi)$, whose intended interpretation is the probability of ϕ as judged by i . A probability judgment without an index defaults to the probability judgment of the speaker.

Formally, for any proposition ϕ , set of worlds W , model M , world w and assignment function g :

$$(22) \quad \|P_i(\phi)\|_{M,w,g} = f_i(\{w \in W \mid (M, w, g) \models \phi\})$$

But why would the speaker assert a probability judgment? And why would the hearer be interested in such a statement? To answer this question, let us, following Wolf and Cohen (2009), consider the role of assertion in context update.

8.5 Assertion and context update

8.5.1 Assertion

Recall from section 8.3.2 above that we follow Vanderveken (1990–91) in assuming that an assertion is composed of (at least) two parts: its propositional content, and the degree of belief in the propositional content.

We propose an assertion operator with two arguments: $A(C, S)$. The first argument, C , is the content of the asserted proposition. The second argument, S , is its degree of strength in terms of a probabilistic inequality.

For example, consider the assertion of (23):

(23) Abby is a doctor.

This sentence does not involve clarity, a modal adverb, or any other element that may modify the sincerity conditions of the assertion. In such cases, S receives a default value. Following Piñon (2006), we assume that the default value is “at least **HIGH**,” for some constant value **HIGH**. Therefore, we formalize the assertion of (23) as (24):

(24) $A(\text{doctor}(a), P(\text{doctor}(a)) \geq \text{HIGH})$

As discussed above, the requirement that the belief in ϕ be higher than some threshold is not sufficient for the assertion of ϕ : it must be accepted, i.e. not defeated by a proposition that is at least as likely.

But why does the speaker indicate her strength of belief, and why would the hearer care about it? Our answer is that the representation of the strength of belief is necessary in a realistic theory of context update.

8.5.2 Context update

The cornerstone of context update theories is Stalnaker (1978), who argues that when an assertion is accepted, it affects the context in which it is made: its propositional content is added to the common ground, and worlds in which this proposition is false are eliminated from the context set.

Stalnaker describes the context set as “the set of possible worlds recognized by the speaker to be the ‘live options’ relevant to the conversation” (pp. 321–2). But a proposition can be a relevant “live option” even before it is accepted. Consider a simple example of a dialog:

(25) John: Bill left.

Marsha: Are you sure?

John: I saw his car drive away.

Marsha: But I think I heard his voice upstairs ... (Ginzburg 1996)

The proposition *Bill left* is neither accepted nor rejected, but it is nonetheless a valuable contribution to the discussion.

Bruce and Farkas (2007) argue that theories of context update should not disregard the *propositional nature* of assertion. Asserted propositions affect the context in significant ways even if they are not yet accepted, but rather are under *negotiation*.

Bruce and Farkas offer such a revision in the form of a transitory stage, a *table*, upon which are placed discourse items that the speaker publicizes her commitment to, before they are accepted by other participants in conversation. Placing an item on the table is a way of proposing a new addition to the common ground. This action affects the context by

projecting a set of the future context state called a *projected set*, which has the potential of replacing the current context state (thus becoming the new common ground). Bruce and Farkas do not discuss the mechanism by which interlocutors decide whether or not to accept propositions that are on the table. While we do not have a complete solution to this problem either, some preliminary work we have carried out indicates a promising approach. Moreover, our view of discourse update assigns an important role to the strength of belief of the speaker in a proposition.

We use a device which we call the *Negotiation Zone*. It is similar to Bruce and Farkas's table, except that the Negotiation Zone hosts propositions coupled with degrees of strength, the second argument of the assertion operator.

The context update process, then, involves a speaker who makes an assertion – essentially, places a proposition plus a degree of strength in the Negotiation Zone. Once the speaker has asserted her proposition, the hearer has to decide whether to accept, reject, further discuss, or agree to disagree on it.

The hearer's decision process takes into account various sources of evidence at the hearer's disposal. These sources may include direct knowledge (e.g. perception), deductive processes (e.g. inference), or reported information (e.g. hearsay).¹¹

The evidence provided by each source is not, in general, conclusive, for two reasons. One is that some sources are more reliable than others. For example, the accuracy of perception will depend on the reliability of the sense organ (e.g. eyes) or sensing device (e.g. camera); the weight of hearsay evidence will depend on the reliability of the reporter, etc.

The second reason why the evidence is not conclusive is that the evidence provided by the source is generally, by its nature, graded and not certain. For example, one possible source of evidence is inference – but inference is often probabilistic. Even if the calculation of probabilities is entirely accurate, the result is still only a probability measure, not a definite result. What is less often acknowledged is the fact that reported evidence is also graded: as we have seen in section 8.3.2, the speaker may believe what she says to various degrees, and, consequently, make statements with varying illocutionary strengths. A speaker who utters (26a) expresses more confidence than a speaker who utters (26b), who, in turn, expresses more confidence than a speaker who utters (26c):

- (26) a. Abby is certainly a doctor.
 b. Abby is a doctor.
 c. Abby is presumably a doctor.

The different illocutionary strengths of the utterances must be taken into account by the hearer in deciding whether or not to accept their propositional content.

In order to formalize the two factors that affect the relative strengths of different sources of evidence, we assign two values to each source of evidence i : a probability measure, $P_i(\phi)$, indicating the strength of the evidence for ϕ according to source i , and a weight w_i , indicating the reliability of the source. Of particular interest here is a specific type of reported information: the assertion made by the speaker. Thus, if the speaker asserted (23), the probability value will be **HIGH**, and the weight will reflect how reliable the speaker is considered.

The probability that the hearer assigns to a proposition is a weighted sum of the probabilities assigned to it by the various sources. The sum of the weights is equal to 1, so that the result can be easily shown to be a probability measure itself:

$$(27) \quad P(\phi) = \sum_{i=1}^n w_i P_i(\phi)$$

This type of probability measure is often referred to as a *mixture model*.¹² Note that because the sum of the weights is 1, the value of the probability does not necessarily increase if the number of sources increases.

If $P(\phi)$ exceeds the hearer's *threshold of acceptance*, the proposition is accepted; if $P(-\phi)$ exceeds this threshold, the proposition is rejected. Otherwise it is left in the Negotiation Zone for further discussion.

8.6 The interpretation of clarity

8.6.1 Simple clarity

8.6.1.1 Another mixture model

Consider again (2), repeated below:

(2) It is clear that Abby is a doctor.

What does the speaker who utters this sentence mean? We have argued that she is reporting a probable belief that is not defeated by a belief that is more or equally probable. But Barker has shown that this cannot be the speaker's belief, or the belief shared by the speaker and hearer. *Whose* belief is reported, then?

We argue that the speaker is referring to the beliefs of people with sound judgment, people who are knowledgeable about the matter, who can reason rationally, and whose judgment is to be trusted. Thus, when someone utters (2), she is saying that people with sound judgment would come to believe, on the basis of the available evidence, that Abby is a doctor.

Note that although what the speaker is reporting is a belief, it is not her subjective belief. In fact, it is not the subjective belief of any one person. In this sense, belief is “objectivized.”

Of course, if the speaker says that people with sound judgment believe something, she would usually believe it herself too. It would normally be odd for someone to acknowledge that people whose judgment she trusts believe *P*, yet deny that she believes *P* herself. Such a statement would be odd, in fact quite similar to Moore’s paradox:

(28) #It is clear that Abby is a doctor, but I don’t believe it.

However, there are cases in which such utterances are acceptable. This is the case when people choose (not) to believe something despite their better judgment, while acknowledging the irrationality of such a belief (and in such cases the word *still* is often used, highlighting the contrast between the objectivized belief and the speaker’s personal belief). The following sentences involve both adverbial and adjectival clarity:

- (29) a. It was clear that they were going to kill him, but we still didn’t believe it.
 b. I’m such a fool. She’s clearly not into me, but I still believe I have a chance.
 c. I realize I’m a big fuddy-duddy. This young woman is clearly the right material for our college, but I still have my doubts.

To further demonstrate this point, here is an attested example:¹³

- (30) It was clear that emotions were mounting, but every feeling was denied and every intervention considered incorrect, negated or ignored.

We account for Barker’s (4) and (5), repeated below, in the same way:

- (4) It is reasonably clear that Mars is barren of life.
 (5) It is clear that God exists.

The speaker who utters (4) may personally believe that there is life on Mars, while acknowledging that the experts deny this. Conversely, someone who denies (5) may believe in God, without necessarily believing that people who reason rationally would conclude, on the basis of the evidence, that God exists. In the former case, the probability value of the mixture will be higher than $d(\text{clear})$, while the strength of belief of the speaker in the proposition that Mars is barren of life may be lower than **HIGH**.¹⁴ In the latter case, the value of the mixture will be lower than $d(\text{clear})$, while the strength of belief of the speaker in the proposition that God exists may be higher than **HIGH**.

We formalize the notion of “good reasoners” using another mixture model. We define a probability mixture over the judgments of possible individual reasoners. Each individual reasoner $1 \leq j \leq m$ is assigned a weight, v_j , indicating how good a reasoner he or she is. And we define $P_{\text{justification}}$ to be the weighted sum of these individual probabilities:

$$(31) \quad P_{\text{justification}}(\phi) = \sum_{j=1}^m v_j P_j(\phi)$$

Of course, the probability $P_{\text{justification}}$ then participates in the mixture in (27), helping the hearer decide whether or not to accept ϕ .

Note that (31) considers not only good reasoners, but also bad ones, although the weights associated with bad reasoners will be small, hence their effect on the resulting probability will also be small. However, in a context in which there are many bad reasoners, their combined judgments may be significant. Consider the following attested example:

(32) It is clear that a strictly monotonically increasing function is one-to-one.

In the context of the mathematics text from which it is taken,¹⁵ (32) is true: the readers of the book are assumed to be versed in mathematical reasoning, and for them this mathematical statement is, presumably, clear. However, in most other contexts, (32) would be false, because reasoners who are able to affirm the statement are few and far between; although each one of them would be associated with a relatively high weight, the nonmathematically oriented reasoners, although associated with low weights, will win by sheer numbers.

8.6.1.2 *Clear*

Using $P_{\text{justification}}$, we can now represent the assertion of (2) as follows:

$$(33) \quad \mathbf{A}(P_{\text{justification}}(\mathbf{doctor}(\mathbf{a})) \geq \mathbf{d}(\mathbf{clear})), \\ P(P_{\text{justification}}(\mathbf{doctor}(\mathbf{a})) \geq \mathbf{d}(\mathbf{clear})) \geq \mathbf{HIGH}$$

This formula means that it is justified to conclude that Abby is a doctor (i.e. “good” reasoners would conclude that Abby is a doctor), and the speaker believes to a high degree that it is justified to conclude that Abby is a doctor.

As discussed above, in addition there is a presupposition that the belief is not defeated: there is no proposition ϕ' that is incompatible with $\mathbf{doctor}(\mathbf{a})$ and $P_{\text{justification}}(\phi') \geq P_{\text{justification}}(\mathbf{doctor}(\mathbf{a}))$.

Note that, as desired, (33) does not entail that Abby is a doctor: the good reasoners, however good they are, may be wrong. However, even though there is no entailment, there is a strong feeling that if (2) is true, then Abby is a doctor.

Indeed, as we have seen above, the speaker would not normally express clarity if she knew that the proposition in the scope of clarity is false, or may be false.

In this, clarity behaves in a way that is similar to factives. This is not to say that the predicate *clear* is a factive – in fact, Barker argues that it is not. He discusses the following examples:

- [(34)] a. We know that p , though p might be false.
 b. We believe that p , though p might be false.
 c. It is clear that p , though p might be false.

Thus for a factive such as *know*, [(34a)] is internally inconsistent. In contrast, in [(34b)] we can allow that our belief might be mistaken without any inconsistency. Clarity behaves more like belief in this respect than like knowledge, as expected on the belief theory of clarity. (p. 255)¹⁶

However, although we agree with Barker that (34c) is not contradictory, it is certainly not a common thing to say. And (35), where p is denied, is downright odd:

- (35) #It is clear that p , though p is false.

We suggest that this is so because if it is justified to conclude something, i.e. if the experts conclude something, it is very likely (though not necessarily) true.

In other ways, clarity actually does pattern with factive verbs.¹⁷ For example, it is well established that factives like *know* can take interrogative complements, while nonfactives, like *believe*, cannot:

- (36) a. I know who the murderer is.
 b. I know whether the murderer is John.
 (37) a. *I believe who the murderer is.
 b. *I believe whether the murderer is John.

Note that *clear* behaves like a factive, in that it, too, allows an interrogative complement:

- (38) a. It's clear who the murderer is.
 b. It's clear whether the murderer is John.

Thus, while *clear* is not a factive, it is, in many ways, similar to factives. In fact, Égré (2008) defines a category of predicates he calls *veridicals*,¹⁸ which subsumes both factives and *clear*.

8.6.1.3 Publicly available evidence

Sentences in the form of (35) become even worse if they rely on evidence that is not publicly available:

- (39) ##It was clear that I was in jail, when in fact I had broken out of jail.

Although it is quite possible that the best minds thought that the speaker was still in jail, this was because the relevant evidence was kept secret, and was not publicly available.

Recall from section 8.2 that Barker argues that clarity must refer to publicly available evidence, and we have proposed that this requirement is a presupposition rather than part of the truth conditions of clarity. We can now say more about the origins of this presupposition.

Since we use a mixture model of a number of reasoners, they cannot be assumed to rely on private evidence, which would be accessible, at most, to only one of them. Consequently, we propose that clarity carries a presupposition to the effect that the relevant evidence is publicly available. This presupposition explains why (39) is completely out.

Note that Barker's theory makes the wrong predictions concerning (39): it follows from it that the sentence is acceptable and true. This is because for Barker, the requirement of publicly available evidence is part of the truth conditions of clarity, rather than a presupposition: (39) is predicted to say that in all normal worlds compatible with publicly available evidence, the speaker is in jail. But, if people do not realize that the speaker escaped from jail, then in all worlds (hence in all normal worlds) consistent with publicly available evidence, the speaker is in jail, and (39) ought to be unproblematically true.

8.6.1.4 *Clearly*

Following Piñon's (2006) account of modal adverbs, we suggest that *clearly* modifies the strength, rather than content, of the assertion.¹⁹ Hence, the assertion of (40a) can be represented as (40b):

- (40) a. Abby is clearly a doctor.
 b. $A(\text{doctor}(\mathbf{a}), P_{\text{justification}}(\text{doctor}(\mathbf{a})) \geq d(\text{clear}))$

This formula means that Abby is a doctor, and it is justified to conclude that Abby is a doctor, i.e. good reasoners would conclude that she is a doctor.

Following Barker, we require that (40) *not* entail that Abby is a doctor. But we must proceed with care, since *clearly* modifies the entire utterance, rather than just its propositional content. What does it mean for an *utterance*, rather than a proposition, to entail something? What would it mean for the utterance in (40) to entail the utterance in (23), whose representation is (24)?

- (23) Abby is a doctor.
 (24) $A(\text{doctor}(\mathbf{a}), P(\text{doctor}(\mathbf{a})) \geq \text{HIGH})$

Note that since in our system assertions can be made at varying degrees of strength, an entailment relation between the propositional contents of two utterances is not sufficient for an entailment relation between the assertions. Consider the examples in (26), repeated below.

- (26) a. Abby is certainly a doctor.
 b. Abby is a doctor.
 c. Abby is presumably a doctor.

Although the propositional contents of all three assertions are the same, intuitively (26c) does not entail (26b), which does not entail (26a).

A useful notion in this context is *illocutionary entailment*, as defined by Searle and Vanderveken (1985): an illocutionary act A_1 illocutionarily entails A_2 if it is impossible to perform A_1 without thereby performing A_2 . For example, a directive illocutionarily entails a suggestion, because it has the same content, but a higher degree of strength. But (26c) has a lower degree of strength than (26b), and, consequently, the former does not entail the latter.

Now our question becomes: does (40) illocutionarily entail (23)? The answer to this question is no: (40) does not necessarily have a higher degree of strength than (24). If good reasoners believe that Abby is a doctor to a degree higher than **d(clear)**, the speaker may believe this too, but does not have to. She certainly does not have to believe this to a degree greater than **HIGH**. Therefore, it is possible for the speaker to perform the act in (40), without thereby performing the act in (24).

8.6.2 Personal clarity

Let us now move to personal clarity:

(41) It is clear to X that Abby is a doctor.

It is tempting to say that (41) expresses nothing more than X's subjective belief. We could say that X believes to a degree of at least **d(clear)** that Abby is a doctor, and the speaker believes to a high degree that X believes to a degree of at least **d(clear)** that Abby is a doctor:

(42) $A(P_X(\mathbf{doctor}(a)) \geq \mathbf{d(clear)})$,
 $P(P_X(\mathbf{doctor}(a)) \geq \mathbf{d(clear)}) \geq \mathbf{HIGH}$

Such an approach, however, would be incorrect. It turns out that not only simple clarity, but even personal clarity, must be "objectivized."

Just like with simple clarity, the evidence for a statement of personal clarity must also be publicly available. Suppose I have secret knowledge that Mother Teresa, whom everybody admires, is actually corrupt. It would still be strange to say (43a); in contrast, (43b) would be perfectly fine.

(43) a. #It is (secretly) clear to me that Mother Teresa is a crook.
 b. It is (secretly) known to me that Mother Teresa is a crook.

Moreover, just like simple clarity, we do not normally use personal clarity if we know that the proposition in the scope of clarity is false. Thus, even if Mary is convinced that the Earth is flat, it would be strange (other than by way of irony) to utter:

(44) #It is clear to Mary that the Earth is flat.²⁰

Indeed, just like simple clarity, personal clarity patterns with factives in that it allows interrogative complements:

- (45) a. It is clear to the police who the murderer is.
 b. It is clear to the police whether John is the murderer.

Consequently, it makes sense to represent personal clarity, just like simple clarity, as a mixture model. The *to*-PP has the role of a speech-act modifier.²¹ It modifies the sincerity conditions of the utterance, indicating that the belief that competent reasoners conclude that Abby is a doctor is ascribed to X, rather than to the speaker. Thus, (41) receives the following representation:

(46) $A(P_{\text{justification}}(\mathbf{doctor(a)}) \geq d(\mathbf{clear}))$,
 $P_X(P_{\text{justification}}(\mathbf{doctor(a)}) \geq d(\mathbf{clear})) \geq \mathbf{HIGH}$

Note that (46) is very similar to (33), the representation of simple clarity: it says that it is justified to conclude that Abby is a doctor, and that this is believed to a high degree; the difference is that in (33) this belief is ascribed to the speaker, whereas in (46) it is ascribed to X. This representation of personal clarity explains its objective feel, and why (43a) and (44) are odd.

In expressions of personal clarity, X is also the one who assigns the weights to the reasoners indicating their competence. Thus, when Alex and Bill disagree in (10), repeated below, their disagreement is over the assignment of these weights: Alex thinks that good reasoners should concentrate on Abby's appearance, while Bill thinks that they ought to reach their conclusions based on her behavior.

- (10) Alex: It is clear to me that Abby is a doctor (because she is wearing a stethoscope).
 Bill: It is clear to me that she's not (because she has an unhealthy habit).²²

Our theory can explain why personal clarity is possible with *clear* but not with *clearly*: as we have seen, the adverb *clearly* is already, in itself, a modifier of sincerity conditions. We have also seen that one such modifier cannot be inside the scope of another: if the *to*-PP is a speech-act modifier, it cannot co-occur with *clearly*.²³

Our view of personal clarity is reminiscent of Anand's (2008) intuition concerning predicates of personal taste; he points out that faultless disagreement is possible even when an overt experiencer is indicated:

- (47) a. X: This is tasty for me.
 Y: But you're wrong. It's disgusting.
 b. X: This is boring for me.
 Y: But you're wrong. It's not boring.

Anand contrasts these examples with constructions that unambiguously express a perceptual experience, where faultless disagreement is impossible:

- (48) a. X: I like this.
 Y: #But you're wrong. It's disgusting.
 b. X: This bores me.
 Y: #But you're wrong. It's not boring.

On the basis of such examples, Anand proposes that predicates of personal taste, both with and without overt experiencers, express the perception of "normal" perceivers. We speculate that using a mixture model may be a fruitful way of formalizing this notion, but will not pursue the matter further here.

8.7 The origins of the mixture model

In this chapter we propose an account of clarity, both adjectival and adverbial, both simple and personal.

Adjectival clarity modifies the propositional content of the utterance, by stating that the proposition is believed by "good" reasoners. Adverbial clarity does not modify the propositional content of the utterance, but rather its sincerity conditions: it indicates that the speaker believes that "good" reasoners would reach the conclusion expressed by the propositional content of the assertion. Both simple and personal clarity are expressions of "objectivized" belief. The former ascribes this belief to the speaker, whereas the latter ascribes it to the person overtly indicated.

Thus, the theory we present in this chapter is an analysis of one particular vague property, clarity, in terms of collective judgments of good reasoners. But why does this property have such a meaning? Why is its meaning dependent on the judgment of competent agents?

We have seen in the previous section that there are grounds to believe that another class of vague predicate, namely predicates of personal taste, can be accounted for in similar ways. But this only pushes the question one step further: why do the meanings of these predicates involve a mixture model?

We do not know the answer to this question, but there is one speculation that appears particularly promising. As discussed above, both clarity and predicates of personal taste are vague. Perhaps, then, vagueness itself requires the use of mixture models: perhaps any vague property ought to be handled in this way.

There is at least one established theory of vagueness that makes a suggestion along these lines.

Wright (1987) says: "For an object to be (definitely) red is for it to be the case that the opinion of each of a sufficient number of competent and attentive subjects ... would be that it was red" (p. 244). This idea can be naturally formalized in terms of a mixture model of the judgments of competent and attentive subjects: if the value of the mixture passes a certain threshold, the object is red, otherwise it is not.

If Wright's theory is on the right track, there is nothing so special about clarity, and its meaning comes simply from its vagueness: for an object to be red is for competent observers to judge that it is red, and for a proposition to be clear is for competent reasoners to judge that it is true.

Notes

* We would like to thank Chris Barker, Manfred Krifka, Hubert Truckenbrodt, and an anonymous referee for helpful comments and suggestions. Our thanks go especially to Paul Égré and Nathan Klinedinst for very detailed comments on an earlier draft.

1. Since clarity is vague, so is the standard for clarity. Analogously, one is tall if one's height is greater than some standard, but, as can be seen by the sorites paradox, defining this standard precisely is problematic.
2. Degrees of normality are formalized relative to a stereotypical ordering source.
3. Barker claims that epistemic *must* is not vague.
4. From Keller (2007: 31).
5. Sentence (17b) is, of course, much better in a context in which it echoes (and questions) a previous utterance. But note that no such context is necessary to make (17a) good.
6. From <http://www.casinobooksonline.com/>.

7. Of course, the sentence also has an implicature that if you do use these formulas, you will not lose. We will not comment on the truth of this implicature ...
8. See Cohen (1999: 20–4) for more on this and other arguments against a normality-based account of generics.
9. In fact, our judgment is that the sentence is inappropriate rather than simply false, providing further evidence that it constitutes a case of presupposition failure.
10. <http://www.thedvdforums.com/forums/showthread.php?t=451805&page=22>
11. If this classification of knowledge sources brings evidentials to the mind of the reader, this is not a coincidence. We believe that evidentials may indeed be accounted for in terms of our view of context update, but will say no more about this here.
12. See McLachlan and Peel (2000) for a thorough exposition.
13. From Oremland and Windholz (1971).
14. As Paul Égré (p.c.) points out to us, this discrepancy between the speaker's belief and the beliefs of good reasoners is probably indicated by the word *reasonably*, for without it (4) would become considerably worse.
15. Krantz (2009).
16. Since he argues against the belief theory of clarity, Barker then accounts for these facts in a different way: the standard for clarity can be low enough so as to make the sentence noncontradictory.
17. We are indebted to Hubert Truckenbrodt (p.c.) for this observation.
18. Égré defines veridicals to be predicates that entail their complements in positive contexts; thus, according to him, (2) does entail that Abby is a doctor.
19. The fact that *clearly* cannot occur in the protasis of the conditional, does not entail that it can *never* be embedded: (i), suggested to us by an anonymous reviewer, appears perfectly fine:

(i) I think that Abby is clearly a doctor.

However, this is not an argument against *clearly* being a speech-act modifier, since, as Krifka (2001) and others have shown, such modifiers *can* be embedded in certain contexts.

20. In contrast, (ii) is perfectly fine; but of course if ϕ seems clear, it means that ϕ is *not* clear.

(ii) It seems clear to Mary that the Earth is flat.

21. We are thankful to Manfred Krifka for suggesting this idea to us.
22. One implication of this view is that (2) and (iii) have the same meaning.

(iii) It is clear to me that Abby is a doctor.

We believe that this is, in fact, the case, and the difference is that (iii) is used to indicate contrast, much like the use of an overt pronoun in a pro-drop language (Enç 1986).

23. Of course, that it not to say that the *idea* that clarity is judged by somebody else cannot be expressed. As an anonymous reviewer points out, all the sentences in (iv) are perfectly felicitous, and appear synonymous with (41).
- (iv) a. From X's point of view, Abby is clearly a doctor.
 b. As far as X is concerned, Abby is clearly a doctor.
 c. To X, Abby is clearly a doctor.

However, the synonymy is only apparent: these sentences actually talk about X's subjective belief, which is not "objectivized." For example, the sentences in (iv) would be quite acceptable even if Abby is known not to be a doctor.

References

- Anand, P. (2008). Predicates of taste: context-shift, assessment-shift, or something else? Presented at the 2nd International Conference on Quotation and Meaning, ZAS, Berlin.
- Barker, C. (2009). Clarity and the grammar of skepticism. *Mind and Language*, 24 (3): 253–73.
- Barker, C. and Taranto, G. (2003). The paradox of asserting clarity. In Koskinen, P. (ed.), *Proceedings of the 14th Western Conference on Linguistics (WECOL)*, pp. 10–21. Vancouver, BC.
- Bruce, K. and Farkas, D. F. (2007) Context structure for dialogues, MS, UC Santa Cruz.
- Cohen, A. (1999). *Think Generic!* CSLI, Stanford.
- Douven, I. (2002). A new solution to the paradoxes of rational acceptability. *British Journal for the Philosophy of Science*, 53: 391–410.
- Égré, P. (2008). Question-embedding and factivity. *Grazer Philosophische Studien*, 77: 85–125.
- Enç, M. (1986). Topic switching and pronominal subjects in Turkish. In Slobin, D. I. and Zimmer, K. (eds), *Studies in Turkish Linguistics*, pp. 195–208. Benjamins, Amsterdam.
- Ginzburg, J. (1996). Dynamics and semantics of dialogue. In Seligman, J. (ed.), *Language, Logic and Computation*, Vol. 1. CSLI Lecture Notes, CSLI, Stanford.
- Halpern, J.Y. (1990). An analysis of first-order logics of probability. *Artificial Intelligence*, 46(3): 311–50.
- Keller, J. A. (2007). *Problems of Evil and the Power of God*. Ashgate, Aldershot.
- Kölbl, M. (2003). Faultless disagreement. *Proceedings of the Aristotelian Society*, 104: 53–73.
- Krantz, S. G. (2009) *A Guide to Real Variables*. Mathematical Association of America.
- Kratzer, A. (1991). Modality/conditionals. In von Stechow, A. and Wunderlich, D. (eds), *Semantik. Ein internationales Handbuch der zeitgenössischen Forschung*, pp. 639–59. Walter de Gruyter, Berlin.
- Krifka, M. (2001). Quantification into question acts. *Natural Language Semantics*, 9: 1–40.
- Krifka, M. (2007). More on the difference between *more than two* and *at least three*. Presented at the University of California, Santa Cruz.

- Krifka, M., Pelletier, F. J., Carlson, G., ter Meulen, A., Link, G., and Chierchia, G. (1995). Genericity: an introduction. In Carlson, G. and Pelletier, J. F. (eds), *The Generic Book*. University of Chicago Press, Chicago.
- Kyburg, H. (1961). *Probability and the Logic of Rational Belief*. Wesleyan University Press, Middletown.
- Kyburg, H. (1976). Chance. *Journal of Philosophical Logic*, 5(3): 355–93.
- Lasersohn, P. (2005). Context dependence, disagreement, and predicates of personal taste. *Linguistics and Philosophy*, 28(6): 643–86.
- Lasersohn, P. (2008). Quantification and perspective in relativist semantics. *Philosophical Perspectives*, 22(1): 305–37.
- McLachlan, G. and Peel, D. (2000). *Finite Mixture Models*. Wiley, New York.
- Oremland, J.D. and Windholz, E. (1971). Some specific transference, countertransference and supervisory problems in the analysis of a narcissistic personality. *International Journal of Psycho-Analysis*, 52: 267–75.
- Piñon, C. (2006). Modal adverbs again. Presented at the Workshop on Syntax, Lexicon and Event Structure, The Hebrew University of Jerusalem.
- Pollock, J. (1995). *Cognitive Carpentry*. MIT Press, Cambridge, Mass.
- Ryan, S. (1996). The epistemic virtues of consistency. *Synthese*, 109: 121–41.
- Searle, J. R. and Vanderveken, D. (1985). *Foundations of Illocutionary Logic*. Cambridge University Press, Cambridge.
- Stalnaker, R. (1978). Assertion. In Cole, P. (ed.), *Syntax and Semantics 9: Pragmatics*. Academic Press, New York.
- Vanderveken, D. (1990–91). *Principles of Language Use*. Cambridge University Press.
- Wolf, L. and Cohen, A. (2009). Modal adverbs as negotiation chips. *Sprache und Datenverarbeitung*, 33(1–2).
- Wright, C. (1987). Further reflections on the sorites paradox. *Philosophical Topics*, 15: 227–90.

9

Commentary on Wolf and Cohen: Reasoning about Public Evidence

Chris Barker
New York University

9.1 The main issues

Barker and Taranto (2003) and Taranto (2006) give an account of clarity based on shared belief. On the shared-belief account, *It is clear that p* will be true just in case both the speaker and the listener mutually believe that *p*. In Barker (2009), I argue that belief is neither necessary nor sufficient to guarantee clarity. After all, if the evidence is weak, *It is clear that God exists* is false no matter how strongly or how mutually we believe in God. Likewise, admitting that something is painfully clear does not guarantee that the discourse participants believe it, at least, not as a matter of entailment. Barker (2009) proposes instead a modal account based on justification: a proposition *p* is clear just in case all sufficiently normal worlds consistent with the common ground (i.e. consistent with publicly available evidence) are *p*-worlds. Then *It is clear that p* is true just in case the publicly available evidence justifies concluding that *p*.

Wolf and Cohen (W&C) agree that shared belief is inadequate, but do not endorse a justification account. Instead, they propose to rescue the belief approach by relativizing belief to a set of idealized judges: “people with sound judgment, people who are knowledgeable about the matter, who can reason rationally, and whose judgment is to be trusted.” This is objectivized belief, since it does not depend (directly) on the subjective beliefs of the discourse participants. However many judges there are – W&C are never specific – we sum the degree to which each of the judges would believe the object proposition, weighted by the reliability of each judge, in what W&C call a mixture model. If the weighted belief sum is sufficiently high, then it is clear that *p*. On this view, then, clarity depends not on belief, but on believability.

One major difference between the approaches concerns the connection with vagueness. On the justificational approach, the entire point of asserting clarity is to constrain a vague standard. This is because whether or not a proposition qualifies as clear crucially depends on how normal a world has to be in order to count as sufficiently normal. The more normal a world has to be, the less skeptical the vague standard of clarity (because we have excluded from consideration more worlds in which unusual things happen). Since the evidence justifying the proposition is already public, that is, already part of the common ground, asserting clarity does not provide any new information about the world. Rather, we learn only something about the vague standard of skepticism. Thus on the justification view, assertions about clarity are assertions directly about vague standards.

On the objectivized belief account, asserting clarity entails facts about what an expert would believe. Since what an expert would believe can vary independently of the standards of any vague predicates, assertions of clarity would potentially provide new information about the world, specifically, information about how experts would behave. On the objectivized belief model, then, assertions of clarity have no special status with respect to vagueness.

There is much to recommend the objectivized belief approach, not least of all its naturalness: certainly anyone who would assent to the clarity of p would be strongly inclined to agree that any objective reasonable person would believe that p , and this fact follows immediately from W&C's analysis. I will point out some additional virtues of their analysis, notably the possibility of deriving the requirement that the evidence supporting clarity must be publicly available. First, however, I will comment on W&C's main empirical arguments in favor of objectivized belief over justification, which concern personal clarity and adverbial clarity.

The theme that will run throughout my commentary is the role of public versus private knowledge. I will suggest that W&C's move from shared belief to objectivized belief should be understood as a move from private belief to public belief, and that this move is at the very heart of clarity.

9.2 Personal clarity

W&C argue that personal clarity (e.g. *It is clear to me that p*) is best analyzed as objectivized belief, so simple clarity is too. (That personal clarity and simple clarity must have the same analysis is just a hunch on their part, though a reasonable hunch.)

Their main piece of evidence that personal clarity depends on belief is the status of (1) in a particular discourse situation.

(1) It is clear to me that Abby is a doctor.

Imagine John and Bill are studying a photograph of Abby in which she is wearing a stethoscope and smoking a cigarette. John asserts (1), considering the stethoscope as decisive. Bill denies (1), reasoning that no doctor would do something as unhealthy as smoke a cigarette.

As W&C point out, this situation is unexpected on the justification account. The reason is that the justification account only allows for two factors to influence the truth of an assertion of clarity: the public facts available, and the prevailing standard of skepticism (i.e. the degree to which the truth of a proposition must be justified before it counts as clear). But the dispute over the truth of (1) is neither a dispute about what evidence is available (both discourse participants are looking at the same photograph), nor is it a dispute about the justification threshold beyond which a proposition counts as clear.

W&C explain this situation by suggesting that there can be a third factor, namely beliefs about probabilities, especially about the likelihood that doctors smoke. If John believes that some doctors smoke, he will judge (1) to be true; if Bill believes that few doctors smoke, he will judge (1) to be false.

I am only inclined to agree that John and Bill may differ with respect to (1) if they are not aware of the deciding factor. If John explicitly says "Only doctors wear stethoscopes," Bill will deny it. Likewise, if Bill declares that "No doctor smokes cigarettes," John will deny that. Once these preliminary disagreements are overtly made, then it seems to me that neither John nor Bill can assert clarity, whether simple clarity or personal clarity, until their other differences of opinion have been resolved.

Thus on the justification view, the dispute in (1) is not about what is clear, but about what is normal. For Bill, who thinks doctors do not smoke, the most normal worlds in which Abby smokes are worlds in which she is somehow not a doctor, despite other evidence to the contrary. If so, then for Bill, asserting clarity is not justified on the basis of the available evidence.

Let us dig deeper. In first-person clarity ("It is clear to me that *p*"), the conclusion must be based on evidence known to the speaker. As a result, first-person clarity can be based on private knowledge ("Based on my secret observations, it is clear to me that ..."). But this is not

quite the same thing as saying that it depends only on the speaker's beliefs. To tease apart private knowledge from private belief, we must assume that it is possible to have beliefs that are not fully justified by any consciously accessible evidence. Perhaps your subconscious mind has come to a conclusion based on evidence that your conscious mind is not aware of – you have an intuition, you have faith in something.

For instance, imagine that you wake up suddenly, convinced that there is an unwelcome presence in the house. You are dimly aware that a sound woke you, but you cannot be sure what it was. You wake your partner, and begin a discussion:

- (2) You: I believe that there is someone downstairs.
 Partner: It's not clear to me that there is anyone down there.
 You: #Well, it is clear to me that there is someone downstairs.

If belief were sufficient to justify asserting clarity, then the final exchange of the dialog in (2) should be OK. And in fact, this ought to be exactly the kind of situation in which asserting personal clarity (as opposed to simple clarity) is useful, since the interlocutor has just declared that they reject clarity.

Yet the assertion of first-person clarity in (2) is not quite right. Even if your belief is quite strong, you cannot assert clarity without concrete evidence, even though the only person who needs to have access to that evidence is yourself. Apparently, if you assert first-person clarity, you are implying that you could exhibit your private evidence if challenged.

On the justification approach, then, the difference between personal clarity and simple clarity is that personal clarity explicitly indicates who must have (conscious!) access to the relevant evidence.

Since personal belief is not sufficient to justify asserting (2), the objectivized belief theory appears to have a problem. However, I would like to offer a natural refinement of the objectivized belief theory that makes the right predictions. In the scenario as given, all the speaker has to go on is a feeling. This would not normally be enough to persuade an idealized reasoner that there is someone downstairs. Even for first-person clarity, we must hypothesize an idealized reasoner, an idealized version of the speaker.

Therefore even on the objectivized belief approach, it is possible for the speaker to believe that *p* at the same time that the speaker is not able to assert "It is clear to me that *p*"! The reason the objectivized belief account can get this example right is precisely because of the move from personal belief to objectivized belief: because the objectivized reasoner does

not have conscious access to whatever evidence caused the unreasoning belief in the presence of an intruder, clarity fails.

I have argued that the justification approach can account for personal clarity at least as well as the objectivized belief account. Furthermore, I have suggested that on the objectivized belief approach even first-person clarity must be objectivized. If so, then the connection W&C would like to make between belief and clarity is just as indirect for personal clarity as it is for simple clarity.

9.3 Adverbial clarity

W&C's second main empirical argument comes from adverbial clarity (*Clearly Abby is a doctor*). W&C take pains to demonstrate that adverbial clarity behaves differently than adjectival clarity. For instance, only adjectival clarity can occur in the antecedent of a conditional (*??If Abby clearly is a doctor ...*).¹

The reason that W&C decide that adverbial clarity is based on objectivized belief rather than justification is that adverbial clarity behaves similarly to a class of adverbials that Piñon (2006) analyzes as expressing the speaker's degree of commitment. W&C construe degree of commitment as degree of belief, and endorse Piñon's analysis (without further argument).

I agree that adverbial clarity is used only when the speaker's commitment to the truth of the object of clarity is high. However, it does not follow that the meaning of adverbial clarity must be based on objectivized belief rather than on justification. After all, if adverbial clarity expresses a high degree of justification, then a rational speaker will normally have a correspondingly high degree of belief, a point discussed in some detail by W&C. But just as for adjectival clarity, when belief and justification come apart, adverbial clarity cleaves to justification, not to belief.

- (3) a. Although the evidence is weak, I believe there is a God.
 b. Therefore God clearly exists.

If adverbial clarity were based on degree of belief, then the reasoning in (3) should be sound, but in fact it is just as shaky as with adjectival clarity. That is, it is perfectly possible for a speaker to strongly believe (3a) without assenting to (3b).

- (4) a. Mars has clearly always been barren of life.
 b. Yet I still believe that there was once life on Mars.

Conversely, if adverbial clarity were based on degrees of belief, (4) should be a logical contradiction. There is something odd about it, to be sure. On the justification account, this is because the speaker of (4) is explicitly admitting that they hold an irrational belief, one that is not compatible with the available evidence. Yet because the evidence, though overwhelming, is incomplete, stubborn speakers can hold out hope for life on Mars without deciding that (4a) must be false.

These examples suggest that belief certainly is not sufficient and may possibly not be necessary for assertions of adverbial clarity, just as for adjectival clarity.

Apparently, a simple belief theory for adverbial clarity faces the same problems it faces for adjectival clarity. If so, then whether objectivized belief is a good theory for adverbial clarity depends on the same factors that are relevant for adjectival clarity, and adverbial clarity does not count as a separate argument in favor of a belief-based approach.

9.4 The lottery example favors justification over objectivized belief

So far in these comments I have argued that personal clarity and adverbial clarity do not necessarily favor either objectivized belief or a justification account. But I think that one argument remains that favors justification, a variant on Kyburg's (1961) lottery scenario. If you own one out of a large number of lottery tickets, the probability that you will win can be made arbitrarily small. The objectivized belief analysis predicts that I should then be able to say *It is clear that your ticket will not be the winning ticket*. But for at least some native speakers, the assertion of clarity is false, no matter what the odds: although it is highly likely that your ticket will not be the winning ticket, it is not clear.

On the implementation of the justification approach in Barker (2009), this judgment is predicted as follows: *p* is clear just in case all maximally normal worlds consistent with the common ground (i.e. consistent with publicly available evidence) are *p*-worlds. Since the unlikely world in which your ticket wins is just as normal as each of the many worlds in which some other ticket wins, the justification account predicts that the proposition that your ticket will not be the winning ticket does not qualify as clear.

W&C's view, following Levi, is that judgments of belief must be relative to some background question. If we focus on whether the individual ticket is likely to win, we can appropriately believe that it will not. If we focus instead on whether the individual ticket is any less likely to win than any other ticket, we will judge otherwise. I agree that under the

right circumstances, we can believe that the ticket will lose, confidently make plans on the basis of that belief, and so on. However, my judgment is that it remains false to assert clarity.

In fact, W&C's attested example from an advertisement for a casino betting strategy – [If you do not buy our product] *It is clear that you will lose* – strikes me as something the copy writer would like to intimidate the reader into accepting, rather than something that is in fact true.

I am not sure what to make of the discrepancy between my judgments concerning the lottery scenario and those of W&C. Some speakers report that their judgments hinge on whether first-person clarity is used or not (*It is clear (to me) that your ticket will lose*): true for first-person clarity, false for simple clarity. If judgments differ systematically on lottery examples for different speakers, it is possible that there are two competing semantics for clarity. Perhaps for some people clarity involves justification, and for other people it involves objectivized belief.

9.5 Deriving the public evidence requirement

One aspect of W&C's hypothesis that I find particularly promising is the possibility that it can derive one of the main features of clarity, namely that clarity depends on publicly available information. That is, based on my private knowledge, I can assert that Abby must be a doctor, but not that it is clear that she is a doctor.

According to W&C, asserting clarity presupposes that all relevant evidence is public.

(5) #It was clear that I was still in jail, even though I had escaped.

They judge this sentence to be infelicitous by virtue of presupposition violation, since a relevant piece of information (that the speaker had escaped) was not public.

W&C correctly point out that the justification approach – at least, the particular technical implementation of the justification approach in Barker (2009) – cannot explain the status of (5). In that analysis, based on Kratzer's theory of modality, there is a realistic modal base reflecting information in the common ground. By assumption, the most normal worlds in that modal base will be ones in which the speaker was still in jail, and that is enough to predict that (5) should be both felicitous and true.

Nor could the analysis be rescued by simply assuming instead that the modal base reflects the complete epistemic state of the speaker, since

that would enable clarity assertions to be based on private knowledge known only the speaker. Thus asserting clarity amounts to a declaration that you have laid all of your cards on the table.

But is this forthrightness requirement a presupposition? If all relevant evidence were public, then surely the existence of a degree from some accredited medical school would be highly relevant to the issue of whether Abby is a doctor. Yet we can appropriately discuss whether it is clear whether someone is a doctor without knowing whether such a document exists. Furthermore, if we presupposed that all relevant evidence were already in the public record, no fact could ever become clear as a result of new evidence coming to light, when in fact it is easily possible for something to become clearer bit by bit over time. So if the public requirement is a presupposition, the presupposition can cover at most evidence already known by one of the discourse participants.

It does seem possible for a speaker to withhold evidence without creating a presupposition failure.

- (6) a. Abby must be a doctor.
- b. It is not clear to you that she is a doctor.
- c. But in fact, it's clear that she is a doctor.

I can assert (6a) based on evidence known only to me (I saw her degree in her home office one day). Assume (6b) is true because you lack that relevant piece of information. If I immediately go on to assert clarity, as in (6c), I say something false, since it is not equally clear to you that she is a doctor. If clarity presupposes that all relevant evidence is public, and (6a) and (6b) in the given context entail that I possess evidence you do not, then (6c) ought to give rise to a presupposition violation; yet it strikes me as being perfectly felicitous (though false).

However, the objectivized belief analysis may not need to stipulate the presence of a presupposition. It may be perfectly adequate to simply assume that the public requirement is part of the truth conditions of clarity, though care is needed when formulating this requirement for an objectivized belief account. Perhaps something like this would be appropriate: *It is clear that p* is true just in case an idealized set of experts would believe that p, where their belief arises entirely through consideration of publicly available evidence.

This may sound a bit complicated, but I suspect that most or all of its complexities can be derived from W&C's basic hypothesis. Here is how the derivation might go: if clarity depends on the opinions of an idealized collection of reasonable people, and if these reasonable people

are not distinguished from one another in any particular way, then we can assume that they are all effectively in the same epistemic state as one another. This is the same thing as saying that they all have the same information available.

Furthermore, simple Gricean reasoning suggests that the information must be publicly available. To see why, imagine that the speaker's idealized reasoners conclude that p is based in part on some information known to the speaker but not to the addressee. Then the speaker should expect the addressee to reject an assertion of clarity, since the addressee may not assume that her reasonable people have access to this nonpublic information. So the addressee's group of idealized reasoners must have access to enough of the information that convinced the speaker's hypothetical reasoners. The easiest way to guarantee this is if the publicly available information is enough to justify the desired conclusion, since that information will manifestly be available to the addressee's idealized group as well.

Thus the publicness of the information that justifies asserting clarity can be derived from W&C's basic hypothesis in combination with some ordinary reasoning about when a speaker will expect an utterance to be accepted.

9.6 Who are these idealized evaluators?

Who are these reasonable people that the objectivized belief account depends on? For instance, can they be significantly smarter than either the speaker or the addressee?

No. Imagine that we have already established that either $P = NP$ or $P \neq NP$, though we do not know which (in particular, we have rejected the possibility that neither proposition can be proved). Imagine further that we believe that a sufficiently smart person, given sufficient time, should be able to figure out which of the two alternatives obtains. Then if the abstract clarity reasoners can be arbitrarily smart, then we should be able to say

- (7) Either it is clear that $P = NP$, or it is clear that $P \neq NP$.
- (8) It might be clear that $P = NP$.

But it seems to me that neither of these claims is true.

Along similar lines, one of us could make the following claim:

- (9) It is clear that $P = NP$.

We would expect this assertion to be true just in case a sufficiently smart person would arrive at a proof of $P = NP$. But in fact, an alert addressee should object to (9). According to the objectivized belief theory, the objection should be “You’re not smart enough to know what a sufficiently smart person would conclude”; but it seems to me that the objection would actually be “No, it’s not clear at all.”

(I am reminded of a fictional detective who sarcastically explained how he was able to untangle a mystery: “I asked myself what a person much smarter than myself would do, then I did that.”)

These arguments suggest that the reasonable people cannot be significantly smarter than the speaker and the addressee.

Well then, can the set of reasonable people be significantly less smart than the speaker and the listener?

Again, no. If the two best chess players in the world are having a private conversation about a game being played in front of them, one might say

(10) It is clear that white will win in nine moves.

She does not mean “clear to a set of ordinary rational observers,” since it is possible that only a very small set of the best chess players could perceive the accuracy of the claim. To see this, imagine that a mediocre chess player joins the conversation. When this happens, the stronger player can no longer justifiably assert (10).

So we conclude that the reasonable people cannot be significantly stupider than the speaker or the addressee.

We have seen that the abilities of the reasoners who are required to evaluate the evidence in favor of clarity must closely match the speaker and the addressee. On the objectivized belief approach, this might be interpreted as a natural constraint on the selection of the group of idealized experts. On the justification view, this follows immediately from assuming that the assessors default to the discourse participants themselves.

By the way, one thing that the chess scenario shows is that clarity cannot be asserted on the basis of a mere belief that a suitable rationale exists. That is, it is not enough for the mediocre player to trust in the accuracy of the grandmasters. Rather, the entire chain of reasoning must be present in the minds of all of the participants, or must at least be closely accessible (as when we have just completed a laborious series of lambda conversions too complex to hold in our minds at one time, at the end of which we say “Now it is clear that the variable remains unbound.”) The inference steps appear to have a status similar to evidence with respect to clarity, in that they must be publicly available.

9.7 Becoming clear

As mentioned above, one of the main differences between the justification theory and the objectivized belief theory concerns the way in which clarity depends on vague thresholds. We can explore this issue by considering a case in which the clarity of a proposition increases over time:

(11) It's becoming clear that John is tired.

One of the ways that something can become clear is that additional pieces of evidence can become public that support the clarity claim. For instance, if John is computing simple sums, each occasion on which he makes an arithmetical mistake might count as a new piece of evidence that he is getting tired. The more frequent the mistakes, the more support for the clarity claim.

On the justification account, there is another way that clarity can conceivably incrementally increase: if the standard of skepticism slowly lowers. Imagine that it is well-known that you and I both get progressively more credible the drunker we get. Imagine further that we have some evidence that John is tired, but not enough to be conclusive. After a couple of beers, I might try to assert (11) as a way of commenting that we are slowly getting drunk. But in fact, (11) cannot be true by virtue of our increasing degree of inebriation. Apparently, the only way for something to become clear is for new evidence in favor of the object proposition to come to light.

On the objectivized belief account, there should be yet a third way for a proposition to become clear: if one's belief in the proposition increases. Imagine that we are under the impression that the longer a person performs mindless arithmetic, the more likely they are to become tired. We are watching John summing numbers without a break. He begins fresh after a good night's sleep, and he performs flawlessly. Nevertheless, after several hours, we both begin to suspect that John must be getting tired, just by virtue of the length of time that he has been working, though we see no signs of fatigue. The likelihood we assign to the proposition that John is tired increases, and we should then be able to assert (11). But once again, this is not possible. There must be some overt, external, concrete manifestations of John's tiredness, some piece of evidence to justify concluding that he is tired. A mere increase in subjective probability unsupported by evidence is not enough to justify an assertion of clarity.

On the justification account, this is easy to explain: only public evidence can support an assertion of clarity. On the objectivized belief

account, this is a mystery: it should be true that the clarity of John's tiredness is increasing, even without any overt evidence.

9.8 Discourse update effects

One very nice aspect of the objectivized belief approach is the predictions it makes concerning the distinctive discourse update effect of clarity assertions. W&C discuss these issues insightfully, though briefly, so more can be said.

One often observed property of clarity assertions is that they are less open to rejection than an ordinary assertion. As Taranto (2006) notes, this side effect of clarity – putting the truth of a proposition beyond negotiation – is so conventional that clarity can be asserted at the very beginning of a conversation, as when a lecture begins with *It is clear that the Roman Empire was one of the greatest of all human civilizations*, presupposing familiarity with the evidence that supports this claim.

This is not surprising on a justification story, since if the object of clarity is sufficiently supported by public evidence, it is reasonable to assume that the proposition in question can be entered into the common ground without fear of objection.

Yet the objectivized belief approach explains this fact about clarity in a particularly direct and satisfying way. The objectivized belief analysis can be paraphrased as saying “a reasonable assessor would believe this sentence is true (based on public evidence).” It immediately follows that a reasonable person would accept an assertion of the clarity object. On this approach, a clarity assertion does not claim that the addressee necessarily believes the object proposition, only that the addressee should believe it, upon pain of being deemed unreasonable. Thus asserting clarity helps signal which parts of a discourse are expected to be unproblematic, versus those other parts that may turn out to be more contentious.

9.9 Outlook

Important empirical issues remain that bear on whether clarity is best analyzed as justification or as objectivized belief, or perhaps as sometimes one, sometimes the other. These issues include the following: What range of judgments are systematically possible for lottery examples, and under what conditions? Is the deviance of the jailbreak examples presupposition failure, or some other kind of infelicity?

Whichever way these questions are answered, work remains to be done for the objectivized belief account exploring connections with the

semantics of other normative judgments, work that W&C have begun with their discussions of predicates of personal taste.

Importantly, both approaches agree that clarity provides unique insights into discourse update, though in rather different ways. On the justification view, clarity asserts claims about the state of the discourse, more specifically, where a vague threshold lies. On the objectivized belief view, W&C suggest that clarity asserts claims about believability. Because their believability corresponds to likelihood assessed by a rational agent based on evidence, it is tantamount to assertability. If so, then the meaning of *It is clear that p* is an assertion that *p* is assertable.

On either view, then, to study the semantics of clarity is in effect to study the grammar of discourse update.

Note

1. I am not sure that W&C's claim that adverbial clarity never occurs in questions is correct; I find *Isn't Abby clearly the best candidate?* fully grammatical.

References

- Barker, C. (2009). Clarity and the grammar of skepticism. *Mind and Language*, 24(3): 253–73.
- Barker, C. and Taranto, G. (2003). The paradox of asserting clarity. In Koskinen, P. (ed.), *Proceedings of the Western Conference on Linguistics (WECOL) 2002*, Vol. 14, pp. 10–21. Department of Linguistics, California State University, Fresno.
- Kyburg, H. (1961). *Probability and the Logic of Rational Belief*. Wesleyan University Press, Middletown.
- Piñon, C. (2006). Modal adverbs again. Presented at the Workshop on Syntax, Lexicon and Event Structure, the Hebrew University of Jerusalem.
- Taranto, G. (2006). *Discourse Adjectives*. Routledge.

This page intentionally left blank

Part III

The Sorites Paradox

This page intentionally left blank

10

Supervaluationism and Fara's Argument Concerning Higher-Order Vagueness

Pablo Cobreros

There is an intuitive appeal to truth-value gaps in the case of vagueness. The idea is that the facts that determine the meaning of vague expressions (facts about use, most likely) left unsettled for a range of cases whether the expression applies. This sort of *semantic unsettledness* has been taken for a long time as a proper motivation for truth-value gaps; since it is unsettled whether the expression applies the corresponding sentence is neither true nor false.

The issue might be formulated in terms of a sort of symmetry in our dispositions to assent to a borderline sentence or to assent to its negation. Suppose Tim is a borderline case of the predicate "thin." The semantic unsettledness is supposed to lead to a situation in which competent speakers refuse to classify Tim as thin and refuse to classify Tim as not thin. It looks that any consistent truth-value assignment would arbitrarily break the symmetry of our dispositions, and one way to respect the symmetry is by saying that neither sentence is true.¹

Another motivation for truth-value gaps is that many have found incredible the idea that we are ignorant in borderline cases. If there is a unique correct bivalent valuation of the language, then there must be a cutoff in a sorites series for, say, the predicate "tall." That is, there must be an item in the series which is the last tall member of the series, immediately followed by a nontall member. Since it is pretty clear that we do not know where this cutoff lies, then we are ignorant about where exactly it does.² Truth-value gap theories claim that we do not have any ignorance of this sort for the simple reason that there is no such cutoff point.

Though the truth-gap view on vagueness has been popular for a time, a number of problems have made many philosophers depart from this position. I take it that Fara's argument concerning higher-order vagueness is one of them. Fara argues that, in order to explain the seeming

absence of sharp transitions in sorites series, the truth-gap theorist is committed to an endless hierarchy of *gap principles*. But then she shows that, given her/his particular commitments on logical consequence, the truth-gap theorist cannot consistently endorse the truth of all these gap principles for any (finite) sorites series.

The aim of this chapter is to provide a way out of this problem for the defender of truth-value gaps (in particular, the supervaluationist). The basic idea is taking a third way between two unsatisfactory options: global and local validity. Local validity is not a satisfactory option for a defender of truth-value gaps because the notion of truth preserved by it does not allow for failures of bivalence. In this respect, global validity looks like a natural option. But the notion of truth preserved by global validity adopts an *external* perspective that makes impossible the accommodation of higher-order vagueness. The proposal considers these two problems and provides a notion of consequence in which the notion of truth necessarily preserved allows for truth-value gaps, but is formulated from an *internal* perspective. The relevant notion of consequence, which I call *regional validity*, lies strictly between global and local validity.

The chapter is divided into three sections. The first briefly presents the supervaluationist theory and Fara's argument concerning higher-order vagueness. The second section presents an argument for the supervaluationist commitment to regional validity and considers briefly the logic provided by this notion of consequence. The third section shows how we might endorse all the gap principles appealed to in Fara's argument given the regional notion of consequence.

10.1 Supervaluationism and Fara's argument

10.1.1 Supervaluationism

The thought underlying truth-value gap theories of vagueness is that vagueness is a matter of underdetermination of meaning. The facts that determine the meaning of a vague predicate do not determine, for a range of cases, whether the predicate applies. The supervaluationist theory provides a picture for understanding vagueness as a form of underdetermination of meaning. The idea is that a vague predicate such as "tall" can be made precise in several ways consistent with the use we make of it. If Peter is a borderline case of tallness, the sentence "Peter is tall" will be true in some ways of making "tall" precise and false in some others. Since each way of making it precise is consistent with the use we make of the expression, our use does not *decide* between these various ways and so the truth value of "Peter is tall" is left unsettled.

The previous picture assumes that the (intuitive or philosophical) notion of truth allows for failures of bivalence. This is often conveyed by the supervaluationist slogan that “truth is supertruth”: a sentence is true just in case it is true in every way of making precise the vague expressions contained in it. We might extend our language with a *definitely* operator (“ \mathcal{D} ” henceforth) which mirrors the notion of supertruth in the object language. The extended language is amenable to a semantics analogous to the possible-worlds semantics for a simple modal language. An interpretation for a language with “ \mathcal{D} ” is an ordered triple $\langle W, R, \nu \rangle$ where W is a nonempty set of *admissible precisifications* (intuitively, admissible ways of making precise all the expressions of the language), R is an admissibility relation between precisifications (wRw' will be read as w' is admitted [deemed admissible] by w) and ν a function assigning truth values to sentences at precisifications.³ Classical operators have their usual meaning (relative to precisifications); the definition of \mathcal{D} is analogous to the definition of the modal operator for necessity:

- $\varphi \rightarrow \psi$ takes value 1 at w just in case at w : either φ takes value 0 or ψ takes value 1.
- $\neg\varphi$ takes value 1 at w just in case φ takes value 0 at w .
- $\mathcal{D}\varphi$ takes value 1 at w just in case φ takes value 1 at every w -admitted precisification.

Constraints on R will depend on the informal reading of the semantics and on questions concerning higher-order vagueness, but it is almost universally admitted in the literature that R should at least be reflexive.

Given this sort of semantics the next question concerns how to define logical consequence. Since logical consequence is a matter of necessary preservation of truth, the commitment to a particular notion of consequence will follow from the commitment to a particular notion of truth. We consider in the first place *local validity* which is the standard way to define logical consequence in modal semantics:⁴

Definition 1 (Local validity)

A sentence φ is a local consequence of a set of sentences Γ , written $\Gamma \vDash_l \varphi$, iff for every interpretation and any point w in that interpretation: if all the $\gamma \in \Gamma$ take value 1 in w then φ takes value 1 in w .

Local validity preserves *local truth*, where a sentence φ is locally true at w just in case φ takes value 1 at that point. Local validity is a natural

way to interpret logical consequence in modal semantics since under this reading of the semantics, being true (in the intuitive or philosophical sense) is being locally true. But local truth does not allow for failures of bivalence in the sense that given an interpretation and a point w , any sentence of the language will be either locally true or locally false at w (and so, there are no interpretations with points at which some sentence is neither locally true nor locally false).⁵ For this reason, given that validity is a matter of necessary preservation of truth, local validity cannot be adequate for a defender of truth-value gaps. It is usually assumed in the literature that the supervaluationist is committed to something like *global validity*.⁶

Definition 2 (Global validity)

A sentence φ is a global consequence of a set of sentences Γ , written $\Gamma \vDash_g \varphi$, iff for every interpretation: if all the $\gamma \in \Gamma$ take value 1 in every point then φ takes value 1 in every point.

Global validity preserves *global truth*, where a sentence φ is globally true at w just in case it takes value 1 at every point. Global truth allows for failures of bivalence since if a sentence φ takes value 1 at a point and value 0 at some other point in the same interpretation, φ is neither globally true nor globally false in any point in that interpretation. In this sense, global truth looks like the natural option for the supervaluationist. As we shall see later, however, global truth is not fully adequate in a different sense.

A particular distinctive feature of global validity is the validity of the following rule:

$$\mathcal{D}\text{-introduction } \Gamma \vdash \varphi \implies \Gamma \vdash \mathcal{D}\varphi$$

which has as a particular case the inference from φ to $\mathcal{D}\varphi$. While the inference is globally valid (if φ takes value 1 at every w so does $\mathcal{D}\varphi$), it is not locally valid (since φ might take value 1 at w whereas $\mathcal{D}\varphi$ does not). The inference plays a key role in Fara's argument below.

10.1.2 Fara's argument concerning higher-order vagueness

Take a long series of men. The first man in the series is 2.5 meters tall and the last is 1.5. Each man in the series differs from his successor by less than a millimeter. It seems that there is no sharp transition from the members of the series that are tall to those that are not tall. Truth-value gap theories explain this fact by appealing to the presence of borderline

cases in between. There is actually no sharp transition from the members of the series of which it is true that they are tall to those of which it is true that they are not tall (= falsely tall), because there are some members in between that are neither truly tall nor truly not tall (falsely tall). Thus, what explains the seeming absence of a sharp transition is the fact that there is no x in the series such that x is truly tall but its successor in the series is truly not tall (falsely tall). Or equivalently, for any x in the series, if x is truly tall, it is not the case that its successor in the series is truly not tall (falsely tall). This *gap principle* for the predicate “tall” is what explains, according to the truth-value gap theorist, the seeming absence of sharp transitions from the members of the series that are tall to those that are not. Given that, for the truth-value gap theorist, “ \mathcal{D} ” is an object language expression of the theory’s notion of truth, we might express this gap principle in the object language:

$$(\text{GP for “}T\text{”}) \mathcal{D}T(x) \rightarrow \neg\mathcal{D}\neg T(x'),$$

where “ T ” stands for “tall” and x' is the successor of x in the series.

However, the seeming absence of a sharp transition in the series cuts deeper than that. We are just as unable to find a sharp transition from the definite cases of “tall” to its nondefinite cases, as we were with respect to the simple positive and negative cases of the predicate. Since the phenomenon seems to be no different in kind, the truth-value gap theorist is compelled to explain this second seeming absence of a sharp transition in analogous terms. Thus, what explains the seeming absence of a sharp transition from the members of the series that are definitely tall to those that are not definitely tall is the truth of a gap principle, this time for “definitely tall”:

$$(\text{GP for “}DT\text{”}) \mathcal{D}DT(x) \rightarrow \neg\mathcal{D}\neg DT(x')$$

There seems to be no reason to claim that higher-order vagueness stops at some finite level. Once one accepts an account of vagueness based on borderline cases, one should be prepared to endorse the hierarchy of borderline cases all the way up. If the theory cannot treat nonterminating higher-order vagueness at least as a logical possibility, one might well start doubting the whole picture. If the truth-value gap theorist endorses the last claim then there might be no sharp transitions in suitably long sorites series, and explaining this fact requires endorsing the truth of all the gap principles of this form:

$$(\text{GP for “}\mathcal{D}^n T\text{”}) \mathcal{D}^n T(x) \rightarrow \neg\mathcal{D}\neg \mathcal{D}^n T(x')$$

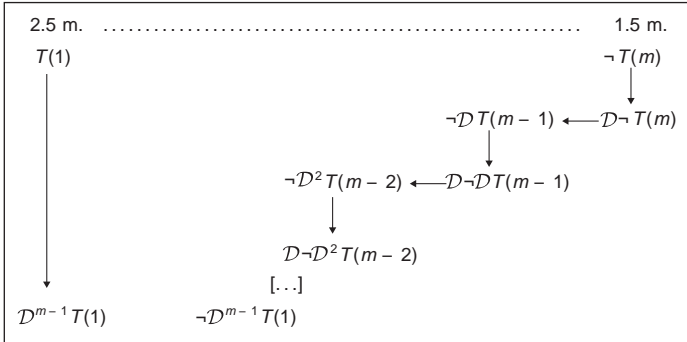


Figure 10.1 Fara’s argument

In her 2003 paper, Delia Fara argues that truth-value gap theorists cannot appeal to the truth of all these gap principles to explain the seeming absence of sharp transitions in sorites series. Fara points out that when we read \mathcal{D} as a truth operator, then one seems to be committed to the rule of \mathcal{D} -introduction (and in particular to the inference from φ to $\mathcal{D}\varphi$) (Fara 2003: 199–200). But Fara shows that, given \mathcal{D} -intro, the truth of all the gap principles is inconsistent for finite sorites series. The proof can be represented with Figure 10.1.⁷

Each move downwards in Figure 10.1 represents an application of \mathcal{D} -intro, while each move leftwards represents an application of the relevant instance of the relevant gap principle (in contrapositive form). As Figure 10.1 shows, the relevant instance of the relevant gap principle works in tandem with \mathcal{D} -intro, forcing us to move leftwards to pick out a new element in the series. A number $m - 1$ of moves leftward suffices to reach the conclusion that $\neg\mathcal{D}^{m-1}T(1)$. But this conclusion contradicts the assumption that the first member of the series is tall, since we might infer that $\mathcal{D}^{m-1}T(1)$ from it, by $m - 1$ applications of \mathcal{D} -intro.

The reasoning employed by Fara in her argument is globally valid (see Cobreros (2010) for discussion on adequate systems for global validity). This means that if one is committed to global validity, one might try to disguise the result, but cannot consistently endorse the truth of all the gap principles. This looks to me as a pretty bad result for global consequence. Still, the supervaluationist might argue that the notion of global validity is not adequate to model the supervaluationist notion of consequence.

10.2 Regional validity

10.2.1 From higher-order vagueness to regional validity

As mentioned above, local validity does not constitute a satisfactory notion of consequence for truth-value gap theories in general and for supervaluationism in particular. The reason is that validity is a matter of necessary preservation of truth, but the notion of truth preserved by local validity does not allow for truth-value gaps. In this respect global validity looks like the right choice, since a sentence might be neither globally true nor globally false in a point in an interpretation. Still, global validity is not adequate in a different sense, connected to the problem of higher-order vagueness.

Higher-order vagueness, from a truth-value gap theory's perspective, is linked with vagueness in the notion of truth. Take a vague predicate " F " and a suitably long sorites series for the predicate. If the truth predicate at work is precise there might be cases in the series of which it is neither true nor false that they are " F ", but there will be a sharp transition from the cases of which the predicate is true to those of which it is not. Likewise, there will be a sharp transition from cases in the series of which the predicate is false to those of which it is not. Now the notion of *global truth* is precise since in any interpretation any sentence will be either globally true in every point or not globally true in every point and so there cannot be borderline cases for "globally true."

Following Williamson (1999), the point might be stated in a somewhat more precise way. Second-order vagueness in φ , from a truth-value gap theory's perspective, might be defined as vagueness in a classification according to which either "It is true that φ " holds or "It is true that $\neg\varphi$ " holds or "It is neither true nor false that φ " holds. The classification is precise if and only if each member is precise and vague otherwise. A sentence φ is precise in an interpretation just in case it takes value 1 in every point or value 0 in every point in the interpretation (Williamson 1999: 131). Now a sentence φ is globally true at a point in an interpretation if and only if it is globally true in every point in that interpretation. Likewise φ is globally false if and only if it is globally false in every point in that interpretation and neither globally true nor globally false if and only if neither globally true nor globally false in every point in that interpretation. This means that global truth makes the above classification precise for any interpretation. Thus, global truth does not allow even for second-order vagueness.

The problem with global validity is that the notion of truth preserved by it takes an *external perspective* in the sense that relativity to points

plays no role in its definition. Saying that something is globally true at a point w in an interpretation is equivalent to simply saying that it is globally true in that interpretation, since, as pointed out before, φ is globally true at a point if and only if it is globally true at every point. This fact makes impossible the existence of borderline cases for *global truth*, since this would require something being globally true at a point and not globally true at some other point. The natural way to address this question is to consider a notion of truth defined from an *internal perspective*, that is, a notion of truth in which relativity to points plays a substantial role. Say that φ is *regionally true* at w in an interpretation just in case φ takes value 1 at every w -admitted point in that interpretation. The corresponding notion of consequence might be called *regional validity*.⁸

Definition 3 (Regional validity)

A sentence φ is a regional consequence of a set of sentences Γ , written $\Gamma \vDash_r \varphi$, iff for every interpretation and any point w in that interpretation: if all the $\gamma \in \Gamma$ take value 1 in every w -admitted point then φ takes value 1 in every w -admitted point.

A sentence φ might take value 1 at a w -admitted point and value 0 at a different w -admitted point, in which case φ will be neither regionally true nor regionally false at w . Unlike local truth, regional truth allows for failures of bivalence and in this sense regional validity is acceptable for the supervaluationist. But unlike global truth, regional truth does not rule out the possibility of second- or higher-order vagueness, since a sentence might be regionally true at a point and not regionally true at a different point in the same interpretation.

In a sense, the notion of regional truth provides a natural strategy for the supervaluationist to accommodate higher-order vagueness. We said before that for the supervaluationist higher-order vagueness is interpreted as vagueness in the notion of truth. Now the idea behind regional truth is that this notion might vary from one precisification to another and, thus, that the relevant notion of truth for the theory can be made precise in several ways (which is what is characteristic of vague predicates according to supervaluationism).

The supervaluationist might look at the situation as follows. If supervaluationism is really committed to global validity, then the discussion concerning higher-order vagueness stops here: the theory cannot accommodate even second-order vagueness. If we give her/him the chance to endorse regional validity instead, then the discussion goes on and we will see whether she/he can address Fara's argument.

10.2.2 Logic with regional validity

The question about which logic flows from the regional notion of consequence depends on restrictions on the admissibility relation. It is usually assumed that factivity ($\mathcal{D}\varphi \rightarrow \varphi$) is a minimal requirement⁹ and where admissibility is reflexive regional consequence is stronger than local consequence. For assume that $\Gamma \not\vdash_r \varphi$; in that case there is an interpretation and a precisification w such that every member of Γ takes value 1 at every w -admitted precisification and φ value 0 at some. The precisification at which φ takes value 0 shows that $\Gamma \not\vdash_l \varphi$. On the other hand, where admissibility is reflexive, not every regionally valid argument is locally valid. In particular $\{\varphi, \neg\mathcal{D}\varphi\} \vDash_r \perp$ (given reflexivity, these two sentences cannot be both regionally true at a precisification w since this would require φ taking value 0 at some w -accessible and φ value 1 at every w -accessible) but $\{\varphi, \neg\mathcal{D}\varphi\} \not\vdash_l \perp$.

Without any other constraint on admissibility, regional validity is weaker than global validity. In particular, the inference from φ to $\mathcal{D}\varphi$ no longer holds, since a sentence φ might be regionally true at w_0 while $\mathcal{D}\varphi$ is not. This is graphically explained in Figure 10.2, which shows that things change if admissibility is required to be transitive; the inference from φ to $\mathcal{D}\varphi$ is regionally valid for transitive interpretations.¹⁰ But there is a reason to reject that admissibility is transitive based on the vagueness of *supertruth* that is quite natural for the supervaluationist. The idea is that in the same way in which the supervaluationist rejects the validity of the schema “ $Fa \rightarrow DFa$ ” based on the vagueness of “ F ,” the supervaluationist rejects the validity of the schema “ $DFa \rightarrow DDFa$ ” based on the vagueness of “supertruth” (recall that \mathcal{D} is an object language expression of *supertruth*).¹¹ Since the transitivity of admissibility guarantees the validity of the schema, the rejection of the schema entails the rejection of the transitivity of admissibility.

The failure of the inference from φ to $\mathcal{D}\varphi$ might look surprising given the supervaluationist reading of “ \mathcal{D} .” As Fara points out, for a truth-value gap theorist, “ \mathcal{D} ” means something like “it is true that,” and in that case “it seems impossible for a sentence S to be true while another sentence, ‘it is true that ‘ S ’, that says (in effect) that it’s true is not true”




Figure 10.2 Failure of \mathcal{D} -intro

(Fara 2003: 199–200). But the supervaluationist might defend the reasonableness of the failure of that inference in the presence of higher-order vagueness. On the one hand, the two facts that lead to that failure (the relativity to worlds involved in the notion of regional truth and the failure of transitivity) are well motivated by the problem of higher-order vagueness. It actually seems to me that the source of the intuition appealed to by Fara is the same as the one stating the validity of the schema “ $Fa \rightarrow \mathcal{D}Fa$,” that the supervaluationist rejects. On the other hand, the failure of the inference should not be confused with the idea that the sentences “ φ ” and “ $\neg\mathcal{D}\varphi$ ” (the sentence stating φ that is not true) might both be true; the last is not accepted, as it is shown by the fact that $\{\varphi, \neg\mathcal{D}\varphi\} \vDash_r \perp$.

A further motivation for the failure of the inference from φ to $\mathcal{D}\varphi$ might be found in the way the supervaluationist addresses Fara’s argument concerning higher-order vagueness, as will be discussed below.

10.3 Gap principles and regional validity

In Figure 10.1, each move leftward represents the use of an instance of a gap principle. The diagram shows that to reach the contradiction for a sorites series of m elements, we need to move $m - 1$ times leftwards and, thus, there are $m - 1$ relevant instances of a gap principle:

$$\begin{aligned} & \text{(GP. } m - 1) \mathcal{D}\neg T(m) \rightarrow \neg\mathcal{D}T(m - 1) \\ & \text{(GP. } m - 2) \mathcal{D}\neg\mathcal{D}T(m - 1) \rightarrow \neg\mathcal{D}^2T(m - 2) \\ & \vdots \\ & \text{(GP. 3) } \mathcal{D}\neg\mathcal{D}^{m-4}T(4) \rightarrow \neg\mathcal{D}^{m-3}T(3) \\ & \text{(GP. 2) } \mathcal{D}\neg\mathcal{D}^{m-3}T(3) \rightarrow \neg\mathcal{D}^{m-2}T(2) \\ & \text{(GP. 1) } \mathcal{D}\neg\mathcal{D}^{m-2}T(2) \rightarrow \neg\mathcal{D}^{m-1}T(1) \end{aligned}$$

In order to show that gap principles are consistent under the regional notion of consequence we have to show that there is an interpretation such that the assumptions of the sorites argument (the first element is tall, the last is not tall) and the relevant instances of the relevant gap principles are regionally true. In other words, we have to show that there is an interpretation with a precisification w in it such that the assumptions of the sorites argument plus all the relevant instances of the relevant gap principles take value 1 in every w -admitted precisification. We will give such an interpretation, but before that I would like to consider a remark concerning the existence of absolutely definite cases.

10.3.1 Absolute definiteness

Suppose that Peter has no hair on his head. Then Peter is bald and, in fact, definitely bald. But is he definitely definitely bald?

Definition 4 (Absolute definiteness)

A sentence φ is *absolutely definite* in a precisification w in an interpretation, just in case each member of $\{\mathcal{D}^n\varphi \mid n \in \omega\}$ holds in w . An object a is an absolutely definite positive case of “ F ” just in case “ a is F ” is absolutely definite; similarly, a is an absolutely definite negative case of “ F ” just in case “ a is not F ” is absolutely definite.

It seems that, at least for some vague predicates, there are absolutely definite positive and/or negative cases of the predicate. This claim, of course, will depend to some extent on the informal reading of “definite.” But the claim seems to be fully justified under the supervaluationist reading. In the case of “bald,” for example, it seems that the use we make of the expression forbids us to count as not bald any person without hair on his head. Thus, according to the supervaluationist theory, there is no precisification in which someone without hair on his head is counted as not bald. This means that if Peter has no hair on his head, then the sentence “Peter is bald” is absolutely definite: you will never reach a precisification through any number of admissibility steps where this sentence takes value 0. Other examples of predicates that seem to have absolutely definite (positive and/or negative) cases seem to be “flat,” “empty” or “open.”

If the previous remark is correct we should impose a further constraint for an acceptable supervaluationist solution. An interpretation showing that gap principles for a vague predicate “ F ” are regionally consistent cannot presuppose that the first element of the series is not an absolutely definite positive case of “ F ” or that the last is not an absolutely definite negative case of “ F ” (since for the supervaluationist there are cases of this sort, this *solution* would be looking at the wrong place).

10.3.2 Regional consistency of gap principles

Given regional validity, the minimum number of elements required to show the consistency of gap principles with absolute definite positive and negative cases is four. Consider the following pocket-sized example. Take the assumptions,

- (1) $\{\mathcal{D}^n T(1) \mid n \in \omega\}$
- (2) $\{\mathcal{D}^n \neg T(4) \mid n \in \omega\}$

- (GP. 1) $DDDT(1) \rightarrow \neg D\neg DDT(2)$
- (GP. 2) $DDT(2) \rightarrow \neg D\neg DT(3)$
- (GP. 3) $DT(3) \rightarrow \neg D\neg T(4)$

An interpretation showing the regional consistency of all the assumptions might look like as shown in Figure 10.3. We assume here (not explicit in the figure) that $T(1)$ and $\neg T(4)$ hold in every world in the model (this guarantees that (1) and (2) hold in w_1) and also that every world accesses itself (to guarantee the reflexivity of R). (GP. 1) is regionally true at w_1 because the consequent takes value 1 at w_1 and w_2 . (GP. 3) is regionally true at w_1 because the antecedent takes value 0 in both w_1 and w_2 . Finally (GP. 2) is regionally true at w_1 because the antecedent takes value 0 in w_1 and the consequent value 1 at w_2 .

Though the model in Figure 10.3 is just a pocket-sized example, it suffices to show that gap principles can be consistently endorsed under the regional notion of consequence: if we can endorse them for sorites series of just four elements, we should be able to do it for larger series.¹²

The following question concerns what is the intuition, if any, of why we can accommodate gap principles given regional validity (and not given global validity). Fara (2003: 201–3) gives the following intuitive explanation of what is going on in her argument (when global validity is at play):

Grasp the first member of a length- m sorites series for “tall” in your left hand; grasp the last member in your right hand. To illustrate that there’s no “sharp” boundary between the tall and the not-tall, you want to move your right hand leftward to grasp a different object that is a borderline case of the predicate “tall” that’s true of the object in your left hand but false of the object in your right hand. After one move leftward of your right hand you still have the object in your left hand that is tall, hence definitely tall, and a new object in your right hand that’s a borderline case of “tall”, hence not definitely tall. Now to illustrate that there is no sharp boundary between the definitely tall and the not definitely tall, you want to move your right

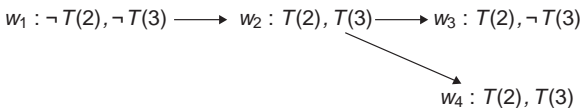


Figure 10.3 Regional consistency of gap principles

hand leftward again, to grasp an object that's a borderline case of the predicate "definitely tall" that's true of the object in your left hand but false of the object in your right hand. Each time you do this, you find you have an object in your left hand of which a predicate of the form " $\mathcal{D}^n T(x)$ " is true, and an object in your right hand of which that predicate is false. The collection of $m - 1$ gap principles appealed to in my argument entail that you can do this at least $m - 1$ times. But you cannot do this as many as $m - 1$ times; there were only $m - 2$ objects between your hands at the start.

Why does this intuitive explanation no longer work if we assume regional rather than global validity? The reason is (of course) connected to the failure of the inference from φ to $\mathcal{D}\varphi$ in the regional case. In Fara's informal explanation, we grasp in our right hand an object a that is a borderline case of the predicate "definitely ^{n} tall" to illustrate that there is no sharp boundary between the definitely ^{n} tall and the not definitely ^{n} tall members of the series. Now it does not follow that this object is definitely not definitely ^{n} tall (that is, it does not follow that the predicate "definitely ^{n} tall" is *false* of this object) and so we are not forced to move leftwards to pick out a new object to illustrate the failure of a sharp transition for the predicate "definitely ^{$n+1$} tall." The same object might illustrate both that there is no sharp transition for the predicate "definitely ^{n} tall" and for the predicate "definitely ^{$n+1$} tall" (this would not have been possible if the object were definitely not definitely ^{n} tall). The general idea is that the borderline cases we take to illustrate the absence of sharp transitions need not be definitely borderline and so the same object might be a borderline case of "definitely ^{n} tall" and a borderline case of "definitely ^{$n+1$} tall" (that is, the same object might illustrate both the absence of a sharp transition for "definitely ^{n} tall" and for "definitely ^{$n+1$} tall").

The way in which regional consequence allows us to address Fara's argument looks to me intuitively appealing. In order to accommodate higher-order vagueness, the truth-value gap theorist is committed to the existence of an endless hierarchy of borderline cases arranged in a particular order in suitable long sorites series. Now given regional validity, one might endorse a commitment of this sort without endorsing the idea that each borderline case in the series is definite. On this view, not everything that smacks of being a borderline case is treated as a clear borderline case¹³ and this looks to me intuitively appealing since the borderline status of an object is, very often, something elusive. When talking about vagueness we say things like "assume a is a clear borderline

case of *F*-ness"; but the truth is that pinning down a particular case that is a clear borderline case of a predicate is usually trickier than what might seem at first sight. The presence of a borderline case is often better witnessed by our disagreement on whether the predicate applies than by our agreement on its being a borderline case. Regional validity has the virtue of allowing us to endorse that something is a borderline case without the commitment of it being a definite borderline case.

Notes

1. Some discussion on this might be found in Williamson (1994) followed by Burgess (2001) and Weatherson (2003).
2. A general point might be made using bivalence, unknowability (if it is borderline whether φ then we do not know whether φ) and a definition of "ignorance."
3. I shall often talk about *points* instead of *precisifications* to remain neutral about the informal reading of the semantics.
4. See, for example, Blackburn et al. (2001: 31).
5. I equate *falsity* with *truth of the negation*.
6. Fine implicitly assumes the global notion of consequence in his classical paper (Fine 1975: 290). Keefe also assumes the global notion (2000b: 176, and 2000a).
7. A picture of this sort might be found in Fara (2003). The argument is based on the same idea as an argument of Wright's (1987, 1992). However, the reasoning used in Wright's argument is too strong since it is not globally valid (this observation on Wright's argument is to be found in Heck 1993). Further discussion of the argument can be found in Edgington (1993) and Sainsbury (1991).
8. As far as I know the first author to consider this sort of approach is Williamson (1994: 159, fn. 32). The term *regional validity* is introduced in Cobreros (2008).
9. See for example Williamson (1999: 130) or the discussion between McGee and McLaughlin (1998: 227) and Williamson (2004: 119).
10. The two notions actually collapse for reflexive and transitive interpretations (Cobreros 2008: 305).
11. An argument of this sort is used in McGee and McLaughlin (1998: 224).
12. In fact, a stronger result can be shown. Fara's argument can be run *in the other direction*, this time taking the relevant instances of the schema:

$$(\text{GP for } "D^n \neg T") \quad \mathcal{D}\mathcal{D}^n \neg T(x) \rightarrow \neg \mathcal{D} \neg \mathcal{D}^n \neg T(x')$$

The previous schema was designed to show that there are no sharp boundaries between the positive cases and the others; this schema is intended to show that there are no sharp boundaries between the negative cases and the others. It can be shown that, given regional validity, the instances of both kinds of schemas are consistent for finite sorites series. The minimum number of elements to show this is five.

13. See Fine (1975: 297).

References

- Blackburn, P., Rijke, M., and Venema, Y. (2001). *Modal Logic*. Cambridge University Press (reprinted with corrections: 2004).
- Burgess, J. A. (2001). Vagueness, epistemicism and response-dependence. *Australasian Journal of Philosophy*, 79(4): 507–24.
- Cobreros, P. (2008). Supervaluationism and logical consequence: a third way. *Studia Logica*, 90(3): 291–312.
- Cobreros, P. (2010). Varzi on supervaluationism and logical consequence. *Mind*, forthcoming. DOI: 10.1093/mind/fzq069.
- Edgington, D. (1993). Wright and Sainsbury on higher-order vagueness. *Analysis*, 53(4): 193–200.
- Fara, D. G. (2003). Gap principles, penumbral consequence, and infinitely higher-order vagueness. In *Liars and Heaps: New Essays on Paradox*, pp. 195–221. Originally published under the name “Delia Graff.” Oxford University Press.
- Fine, K. (1975). Vagueness, truth and logic. *Synthese*, 30: 265–300.
- Heck, R. (1993). A note on the logic of (higher-order) vagueness. *Analysis*, 53(4): 201–8.
- Keefe, R. (2000a). Supervaluationism and validity. *Philosophical Topics*, 28: 93–106.
- Keefe, R. (2000b). *Theories of Vagueness*. Cambridge University Press.
- McGee, V. and McLaughlin, B. (1998). Review of *Vagueness*. *Linguistics and Philosophy*, 21: 221–35.
- Sainsbury, M. (1991). Is there higher-order vagueness? *Philosophical Quarterly*, 41(163): 167–82.
- Weatherson, B. (2003). Epistemicism, parasites and vague names. *Australasian Journal of Philosophy*, 81(2): 276–9.
- Williamson, T. (1994). *Vagueness*. Routledge.
- Williamson, T. (1999). On the structure of higher-order vagueness. *Mind*, 108: 127–43.
- Williamson, T. (2004). Reply to McGee and McLaughlin. *Linguistics and Philosophy*, 27: 113–22.
- Wright, C. (1987). Further reflexions on the sorites paradox. *Philosophical Topics*, 15: 227–90.
- Wright, C. (1992). Is higher-order vagueness coherent? *Analysis*, 52: 129–39.

11

Truth in a Region

Delia Graff Fara
Princeton University

Vague predicates have borderline cases. A certain shortish, thin and scraggly fir tree might only be a borderline case of “suitable as a Christmas tree” for example. A number of theorists of vagueness would regard this as equivalent to a claim that they would express using a “definitely” operator: the tree is neither definitely suitable nor definitely unsuitable. Standard versions of supervaluation semantics for vague predicates yield the surprising result that every sentence entails its definitization: if it is true that Firby the fir tree is suitable, then it is true that Firby is definitely suitable.¹ The truth of the claim that Firby is suitable *guarantees* the truth of the claim that Firby is definitely suitable. Let us call the entailment principle in question **D-introduction**.

D-intro is problematic for at least two reasons. First, prefixing a claim with “definitely” seems to yield a stronger assertion than making the claim by itself: if I say that Firby is suitable, it does not feel like I have committed myself to the seemingly stronger claim that Firby is *definitely* suitable. It seems harder, so to speak, to be definitely suitable than it is to be just suitable. A natural thought is that when I say that Firby is suitable, I am expressing just my own judgment, whereas when I say that Firby is definitely suitable (or *clearly* suitable) I am making a claim about how others would or should judge – that Firby’s suitability is not something I would expect others to disagree about if they are reasonable and competent.

Second, **D-intro** embroils the supervaluationist in a certain paradox – which we will present shortly – if she wants to adopt what I call “gap principles” in order to resolve problems that arise from higher-order vagueness.²

Pablo Cobreros (Chapter 10, this volume) has devised an ingenious version of supervaluation semantics – call it *region-valuation semantics* – that

is designed to invalidate **D-intro** in order to save the gap principles from paradox. The central thesis of standard supervaluationism is that a sentence is true just in case it would be true according to classical bivalent semantics no matter how vague predicates were admissibly “precisified” in ways that are (in some sense) compatible with their meaning. Cobreros’s regional semantics does not admit of such a straightforward gloss, but it will suffice to say, for now, that region-valuation semantics renders a sentence true just in case it would be true, according to bivalent semantics, everywhere in a certain restricted “region” of admissible precisifications, where the region need not include all of the admissible precisifications there are. Call this conception of truth “region-truth.”

In order to demonstrate that regionalist semantics renders gap principles coherent, Cobreros devises a model in which the gap principles are region-true and which, therefore, demonstrates that **D-intro** is not regionally valid – that is, that it does not preserve region-truth.

There are two overarching themes to this chapter. One is that the supervaluationist has an independent conception of what it is to be a borderline case; if she takes that independent conception as primitive, then her views about the connection between borderline status and definite status commit her to a conception of definiteness that validates **D-intro**.

The second theme is that we should question the relevance of Cobreros’s model. If we are using a formal model theory to tell us anything about which inferences are good inferences, then we should restrict our attention only to models that are plausible representations of the way things could in fact be. I argue that Cobreros’s model is not such a model in a number of respects. Consequently, the burden is shifted back to the region-valuationist to provide us with a more plausible model that shows that the gap principles can all be true.

My plan is to begin by presenting an instance of the sorites paradox (section 11.1), then to continue by explaining why a supervaluationist might want to adopt gap principles, then to proceed from there to present the gap-principles paradox (section 11.2). After that I provide my own presentation of supervaluation semantics (section 11.3). In that section I argue that the supervaluationist’s conception of borderline cases as *things that could have been classified either way* straightforwardly leads to a semantics for the definitely operator, *D*, that validates **D-intro**. In the following section (11.4) I present a version of supervaluationism that revises its semantics for *D* in order to shed certain serious problems with simple supervaluationism, but which nevertheless validates **D-intro**. I will present Cobreros’s regional semantics and his counter-model

to **D-intro** along with my criticisms of it (sections 11.5–8). Finally (section 11.9), I will present an incontrovertible argument for why supervaluationists and region-valuationists, as well as any theorist who proposes that borderline cases of a predicate are to be identified with the things of which the predicate is neither true nor false, are all committed to **D-intro**. I conclude by explaining why my arguments, taken together, constitute an argument against any truth-value gap approach to vagueness.

11.1 A ‘suitable’ version of the sorites paradox

It can be vague whether a tree is *suitable as a Christmas tree*. Some trees are definitely suitable – for example, any fir or spruce that is full and cone-shaped and between 1 and 2 meters tall. Other trees are definitely unsuitable – for example, deciduous ones or hedge-shaped ones. But there are further trees that fall in between: that are neither definitely suitable nor definitely unsuitable. These trees fall into a *gap* between the definitely suitable and the definitely unsuitable. Accounting for this gap – saying what kind of gap it is (e.g. a gap between truth and falsity or a gap between areas where there would be no reasonable disagreement about suitability) – is “the problem of borderline cases.”

Also, given enough firs (and some shears), we can construct a *sorites series* of trees ranging, in very small increments of change in size, shape, and fullness, from a tall, conical and full one on the left (definitely suitable) to a puny, hedge-shaped and scraggly one on the right (definitely unsuitable). Problematically, there appears to be no sharp transition from the suitable trees to the unsuitable ones. There does not appear to be a bivalent yes–no switch from suitability to unsuitability. Flat-out denying, however, that there is a switch from a simple “yes – suitable” to anything else requires denying that at the end of the series there has been a switch to “no – not suitable.” Resolving this difficult predicament is “the problem of the sorites paradox.”

Many theorists have let their account of borderline cases serve double duty by using it not only to solve the problem of borderline cases but also to resolve the sorites predicament. They say that it is the intervention of borderline-suitable Christmas trees that makes it correct to deny a sudden, bivalent, yes–no switch from suitability to unsuitability. A sudden switch from “yes” to “no” is incorrect, they say, because there are “neither-yes-nor-no” cases in between. I will call these theorists “gap theorists.” (Really, that is short for “gap theorist of the sorites paradox.”)

11.2 The gap-principles paradox

These are the difficulties with vagueness that we have mentioned so far: the problem of borderline cases; the problem of the sorites paradox; and a second paradox – the gap-principles paradox – that is problematic for gap theorists in particular. “Gap theorists,” remember, are those who use their account of borderline cases to explain the appearance of a lack of a sharp boundary on a sorites series – a boundary that directly divides the cases where a predicate applies and the cases where it does not apply.

Let us turn now to that second paradox. Gap theorists promote the following *gap principle*:

Gap principle: For any object x and its successor x' on our series, if x is definitely suitable then x' is not definitely unsuitable.

The principle is that definitely suitable trees and definitely unsuitable trees can never lie right next to each other on a series of trees, each of which differs only in minuscule ways from its neighbors. In symbols we can write it like this:

$$\forall x \forall y ((DSx \wedge D\neg Sy) \rightarrow y \neq x').$$

Or, just as well, we will write it like this:

$$G1 : \forall x (DSx \rightarrow \neg D\neg Sx').$$

Gap theorists' use of borderline cases to provide the explanation of a seeming lack of sharp boundary have their back up against a hard wall. Just as there appears to be no sharp boundary between the suitable trees and the unsuitable trees, there appears to be no sharp boundary between the definitely suitable trees and the borderline cases. If – as gap theorists propose – it is the intervention of borderline suitable trees that explains the appearance (which they deem true) of no sharp boundary between the suitable trees and the unsuitable ones, then it should be the intervention of borderline-definitely suitable trees that explains the appearance of no sharp boundary between the trees that are definitely suitable and those that are not. This leads to the following second-order gap principle: the definitely definitely suitable trees and the definitely not definitely suitable trees can never lie right next to each other on a series of trees, each of which differs only in minuscule ways from its

successor. In symbols:

$$\mathbf{G2} : \forall x(DDSx \rightarrow \neg D\neg DSx').$$

But again, there is no sharp boundary between the trees that are definitely definitely suitable and those that are not. This, the story must continue, is because there are some borderline-definitely definitely suitable trees between the ones that are *definitely* definitely definitely suitable and those that are *definitely not* definitely definitely suitable:

$$\mathbf{G3} : \forall x(DDDSx \rightarrow \neg D\neg DDSx').$$

Ad nauseum:

$$\mathbf{G4} : \forall x(DDDDSx \rightarrow \neg D\neg DDDSx'),$$

$$\mathbf{G5} : \forall x(DDDDDSx \rightarrow \neg D\neg DDDDSx'),$$

$$\mathbf{G6} : \forall x(DDDDDSx \rightarrow \neg D\neg DDDDDx'),$$

⋮

$$\mathbf{G(n + 1)} : \forall x(DD^n Sx \rightarrow \neg D\neg D^n Sx').$$

But this is paradoxical. The last tree of our 1000-term sorites series is definitely unsuitable, so its predecessor cannot be definitely suitable, in accordance with the first-order gap principle, G1. But then since the second to last tree is not definitely suitable, and since that entails that it is definitely not definitely suitable (again, by **D-intro**), its predecessor cannot be definitely definitely suitable, in accordance with the second-order gap principle, G2. Repeat the reasoning backwards along the series and we conclude that the first, beautiful, full, green, and conical 2-meter tall tree is not definitely definitely definitely ... definitely suitable (that is 999 total “definitely’s”). But it is.

Laid out in a column of reasoning:

$D\neg S(1000)$	
$\neg DS(999)$	Gap principle for $S(x)$
$D\neg DS(999)$	D-intro
$\neg D^2 S(998)$	Gap principle for $DS(x)$
$D\neg D^2 S(998)$	D-intro
$\neg D^3 S(997)$	Gap principle for $D^2 S(x)$
⋮	⋮
$\neg D^{999} S(1)$	Gap principle for $D^{998} S(x)$

A further argument beginning with $T(1)$ yields $D^{999}T(1)$, after 999 applications of **D-intro**. Contradiction. Let us call this the gap-principles paradox.

A supervaluationist who wants to avoid the gap-principles paradox while retaining her gap principles might deny that **D-intro** is valid. But in canonical supervaluation semantics, it is valid. Let us quickly review that canonical semantics.³

11.3 Super-truth and bi-truth

The supervaluationist says that on our sorites series for “suitable,” the predicate will be true of some initial segment, false of some final segment, and neither true nor false of some intervening segment. Let us, though, draw a single yes–no line between the things of which the predicate is true and the things of which it is false (its extension and anti-extension, respectively), so that everything falls into one or the other category. Let us also draw such a line for every predicate in our language. We will use the term *precisification* to refer to any group of bivalent line drawings when the group contains exactly one such line drawing for each predicate.

I will use the term *model* to refer to any collection of precisifications. A supervaluationist uses two notions of truth evaluation to present her theory. One is the notion of *supervaluation*, which is not bivalent – there will be sentences that are neither super-true nor super-false. The second is a notion from classical semantics, *bi-valuation*, according to which every sentence is either bi-true or bi-false. The supervaluationist then says that a sentence is “super-true” in a model just in case it is bi-true “at” (as we will say) every precisification in the model, super-false if bi-false at every precisification in the model. Bi-truth and bi-falsity are values that a sentence has relative to a precisification in a model, while super-truth and super-falsity are values that a sentence has relative to a model.

Bi-truth and bi-falsity: A sentence is bi-true at a precisification just in case it is not bi-false at that precisification. Bi-truth and bi-falsity conditions for logically complex sentences at a precisification are the same as the classical truth and falsity conditions for those sentences.

Super-truth: A sentence is super-true in a model just in case it is bi-true at every precisification in that model.

Super-falsity: A sentence is super-false in a model just in case it is bi-false at every precisification in that model.

Some precisifications are “admissible” while others are not. Here are two criteria for admissibility (I leave it to supervaluationists to explicitly

articulate any others): (i) the precisification respects all comparative relations – for example, by putting something in the extension of “tall” only if everything taller than it is also in there, or by putting a frequency of events in the extension of “often” only if every event of that type that occurs with greater frequency is also in there;⁴ and (ii) it respects analytic relations among different words – for example, by putting something in the extension of “tiny” only if it is in the extension of “small,” or by putting a pair into the extension of “to the left of” if, and only if, it puts the reverse of that pair into the extension of “to the right of.” We will say that a model is admissible when all of its precisifications are admissible.

11.3.1 The argument for D-intro from ‘borderline’ bi-truth

But how does the supervaluationist evaluate claims involving “definitely” or “borderline” in a model? What are the conditions for the super-truth of the claim that a given tree is *borderline* suitable? Like other sentences, the supervaluationist will deem these claims to be super-true in a model if they are bi-true at every precisification in the model and super-false if bi-false at every precisification in the model. We have indicated the bi-truth conditions for truth-functionally or quantificationally complex sentences at precisifications – they are the classical ones. But the “definitely” and “borderline case” operators are not operators of standard languages for classical logic; they do not come along with their own antecedent bivalent truth conditions for evaluation of their bi-truth at a precisification. Let me begin with the question of what bi-truth conditions the supervaluationist should give for the borderline-case operator *B*. I begin with the question about *B* since my concern is to argue that the intuitive understanding of *B* leads to a bi-truth condition for *it* that yields the validity of **D-intro**. The canonical answer is that a “borderline”-prefixed claim is bi-true at a precisification in a model when the claim itself is bi-true at some precisifications in the model but bi-false at others.

‘Borderline’ bi-truth at a precisification: $B\Phi$ is bi-true at a precisification in a model just in case Φ is bi-true at some precisifications in the model while bi-false at others.

This canonical answer is the right answer for the supervaluationist to give. One key component of the supervaluationist’s understanding of borderline cases is this: she takes an object’s *having borderline status* to mean that things *could have gone either way* for that object depending on how sharp lines for predicates could reasonably have been drawn. But whether things could have gone either way among the different

precisifications is independent of which way things did go at any one precisification. So the bi-value of a “borderline”-prefixed claim will not vary from precisification to precisification in a model and will be bi-true at a precisification in a model just in case the bare claim, without the prefix, has different bi-values at different precisifications.⁵

But once this bi-truth condition for B is admitted – as it has to be, given the above considerations – the canonical bi-truth condition for D just drops out:

‘Definite’ bi-truth at a precisification: $D\Phi$ is bi-true at a precisification in a model just in case Φ is bi-true at every precisification in the model.

This bi-truth condition for D follows from that for B given only two minimal assumptions. Given the standard definition of $B\Phi$ as $\neg D\Phi \wedge \neg D\neg\Phi$ (first assumption), $D\Phi$ is bi-true at a precisification in a model only if $B\Phi$ is bi-false at that precisification in that model. Given the above bi-truth condition for B , this means that $D\Phi$ can be bi-true at a precisification in a model only if Φ is either bi-true at every precisification in the model or bi-false at every precisification in the model. But then, given that the analog of the T schema for modal logic is valid: $\models D\Phi \rightarrow \Phi$ (second assumption),⁶ $D\Phi$ is bi-true at a precisification in a model only if Φ is bi-true at that precisification, and hence only if Φ is bi-true at every precisification in the model.

Conversely, suppose that Φ is bi-true at every precisification in a model. Then at any given precisification, $B\Phi$ will be bi-false, given the bi-truth condition for B . The definition of $B\Phi$ as $\neg D\Phi \wedge \neg D\neg\Phi$ then requires that, at any given precisification, $D\Phi \vee D\neg\Phi$ be bi-true. This in turn requires that at each precisification either $D\Phi$ or $D\neg\Phi$ is bi-true. If it were $D\neg\Phi$ at some precisification, then since Φ is true at that precisification (by hypothesis) $D\neg\Phi \rightarrow \neg\Phi$ would be false at that precisification, in violation of the validity of the T-schema.

So now we have the standard truth clauses for supervaluation semantics on the table. On this standard picture, **D-introduction** is valid; any sentence Φ entails its definitization, $D\Phi$. For if Φ is super-true in a model, then it is bi-true at every precisification in the model; that is the condition for $D\Phi$ to be bi-true at any precisification, so it is bi-true at every precisification in the model, and so $D\Phi$ is super-true in the model.

Whether our condition for “definite” bi-truth at a precisification is the standard supervaluational one is irrelevant. We just gave an *argument* that supervaluationists must accept it. The ideas that supported the argument, and which make it incontrovertible for the supervaluationist, are

(i) that something has borderline status when it could reasonably be classified either way; (ii) that truth is super-truth; (iii) that $B\Phi$ is definable as $\neg D\Phi \wedge \neg D\neg\Phi$; and (iv) that $D\Phi \rightarrow \Phi$ (whatever is definitely suitable is suitable) is true at every precisification in every model. The point to underline here is one that we have mentioned before, namely that the semantics for “definitely” follows from independent considerations about the semantics for “borderline.”

This standard version of supervaluationism has a number of serious problems. One is the problem we have been focusing on, namely its validation of **D-intro**. This is a problem because it leads to the gap-principles paradox for the supervaluationist who wants to maintain gap principles in order to deal with higher-order vagueness in the way that she deals with first-order vagueness.

Another problem, related to the first, is that it validates the following principle of iterability of D :

Iteration: $D\Phi \rightarrow DD\Phi$.⁷

Iteration entails that any borderline definite case of a vague predicate is not a definite case: iteration entails $BDF(x) \rightarrow \neg DF(x)$. This is not good. In general we could not have reason to assert both $B\Phi$ and $\neg\Phi$. For example, if we are in a position to deny that a given tree is suitable, then we should not be in a position to assert that it is a borderline case. If we were justified in asserting that something was *borderline* Φ , we should not be justified in asserting that it *is not* Φ . If an attribution of borderline status is justified when things *could have gone either way*, then surely it cannot simultaneously be justified to say that things *have not gone that way*. But this is precisely what iteration dictates. In our case, Φ is the predicate “definitely F .”

Borderline definite but not definite: $\models BD\Phi \rightarrow \neg D\Phi$.⁸

Perhaps a more serious problem is the validation of a weak Euclidean principle, $B\Phi \rightarrow DB\Phi$, the principle that all borderline cases are definitely borderline cases – that there are no borderline borderline cases.

Weak Euclidean principle: $\models B\Phi \rightarrow DB\Phi$.⁹

This is problematic for the gap theorist because it rules out his idea that the seeming absence of a sharp boundary between those trees that are borderline suitable and those that are not is due to an intervening segment of trees that are neither definitely borderline suitable nor definitely

not borderline suitable. If the weak Euclidean principle is super-true in every model, there cannot be such trees.

11.4 Revised supervaluationism

In order to avoid the gap-principles paradox while retaining her gap principles, the supervaluationist needs to disavow her commitment to **D-intro**. One way to do this would be first, to revise the condition of bi-truth-at-a-precisification for D . A second option would be to reject the idea that *truth is super-truth*, and, correspondingly, that validity is preservation of super-truth in a model. Pursuing the first option without the second, though, would not achieve the desired result.

One salient way to pursue the first option, that is, to revise the bi-truth condition of the “definitely” operator, would build on two observations: (i) that supervaluational models are analogous to standard models for modal operators in that they involve collections of “points” at which (something like) truth is bivalent and in that the truth of certain sentences at a “point” in a model is determined by the truth of certain other sentences at other “points” in the model – in the modal case, the “points” are possible worlds rather than precisifications; and (ii) the intended meaning of the definitely operator is like that of the necessity operator in that both can be construed as having, each in its own way, the meaning “true no matter what” – in the modal case, this would be *true no matter how things had been*; in the supervaluationist case, this would be, *true no matter how our vague predicates had been (admissibly) precisified*.¹⁰ In the basic semantics for the necessity and definitely operators, the truth of $\Box\Phi$ and $D\Phi$ at a point in a model depends on the truth of Φ at every point in that model. One standard revision of the semantics for the necessity operator involves the introduction of an “accessibility” relation between worlds: then $\Box\Phi$ is defined to be true at a world just in case Φ is true at every world that is *accessible* to it. Similarly, one might introduce an accessibility relation among precisifications in a supervaluational model and define $D\Phi$ to be bi-true at a precisification just in case Φ is bi-true at every precisification that is accessible to it.

Revised ‘definite’ bi-truth at a precisification: $D\Phi$ is bi-true at a precisification p in a model when Φ is bi-true in the model at every precisification that is accessible to p .

A revision of the bi-truth condition for D requires a corresponding revision of the bi-truth condition for B . This is required in order to preserve

the equivalence between being borderline suitable (for example) and being neither definitely suitable nor definitely unsuitable. The required revision is this:

Revised “borderline” bi-truth at a precisification: $B\Phi$ is bi-true at a precisification p in a model when Φ is bi-true in the model at some precisifications that are accessible to p while bi-false at others.

The upshot of the revised definitions is that “definitely” or “borderline” claims are no longer super-bivalent (always either super-true or super-false) – there will be models containing precisifications at which such claims will be neither super-true nor super-false.

If “Firby is suitable” is bi-true at every precisification that is accessible to p , then “Firby is definitely suitable” will be true at p . But this condition might obtain even if “Firby is suitable” is bi-false at some precisification p' that is accessible to a precisification p' that is accessible to p . For p' might not be accessible to p . In this case, “Firby is definitely suitable” will not be bi-true at p' , in which case “Firby is definitely suitable” will be neither super-true nor super-false in this model; it will be bi-true at some precisifications (e.g. p) while bi-false at others (e.g. p').

One consequence of the new definition for the bi-truth of D that is considered to be advantageous – this is in fact its *raison d'être* – is that the principle of iteration, “ $D\Phi \rightarrow DD\Phi$,” and the weak Euclidean principle, $B\Phi \rightarrow DB\Phi$, are no longer super-valid (super-true in every model). They were both super-valid given the original condition for bi-truth at a precisification but neither is given the revised one.

But the validity of **D-intro** did not turn on the super-bivalence of “definitely” claims. The claim of **D-intro**, is that if Φ is super-true in a model, then $D\Phi$ is super-true in that model. However much the bi-truth condition of D is weakened, though, the super-truth of Φ in a model must surely suffice to render $D\Phi$ bi-true at each precisification p : everywhere p “looks,” it will find only precisifications at which Φ is bi-true. That must surely suffice for $D\Phi$ to be true at p , and so on for every other precisification in the model. So the new semantics for D gives no shelter from the validity of **D-intro**. Nor will any other change in the semantics for D as long as the bi-truth of Φ at every precisification ensures the bi-truth of $D\Phi$ at each precisification – as it should.

Our argument (pp. 228–9) for the canonical bi-truth condition for D did, however, turn on the identification of truth with super-truth. We had emphasized that the formalization of the supervaluationist’s intuitive understanding of what it is to be a borderline case by itself had the

canonical semantics for D as a consequence. The intuitive understanding of what it is to be a borderline case was this: *things could have gone either way*, depending on how our predicates had been admissibly precisified. We in turn took it that things possibly having gone either way was to be formally represented as bi-truth at some but not all precisifications in a model. But if *truth* does not depend on bi-truth at *every* precisification in a model, then that formal representation of borderline status can be rejected.

So perhaps the gap-theorist supervaluationist can find refuge by giving up *the* core idea of her theory, namely, that truth is super-truth – truth no matter how precise lines for our predicates might reasonably have been drawn. For this would yield an opening for a corresponding revision of what *truth* is. If *validity* is preservation of *truth* but truth is not super-truth, then although **D-intro** does preserve super-truth, it might nonetheless not be valid. Giving up the identification of truth with super-truth is precisely what Cobreros does. In order to appreciate the motivation for his revision, it helps to contrast it with an alternative revision.

Instead of characterizing validity as preservation of super-truth in a model, one might characterize it instead as preservation of bi-truth at any precisification in any model. Timothy Williamson gives the names “global validity” to the first characterization and “local validity” to the second.¹¹ **D-intro** is globally valid, as we have seen, but it is not locally valid. Even without the introduction of an accessibility relation among precisifications in a model, and the concomitant weakening of the bi-truth condition for D , Φ might be bi-true at a precisification p in some model without $D\Phi$ being bi-true at p since Φ might be bi-false at other precisifications in that model, in which case $D\Phi$ will be bi-false at p .

As Williamson effectively argues, the supervaluationist has no right to the localist characterization of validity, since she cannot identify it with necessary preservation of *truth* (Williamson 1994: 147–9). For her, truth cannot be bi-truth, since for her, truth is not bivalent while bi-truth at a precisification is.

11.5 Region-valuationism

Cobreros takes this problematic as his starting point. He wants a characterization of validity that, like local validity, does not validate **D-intro** but which nevertheless can be rightly understood by the supervaluationist as necessary preservation of *truth*. His self-imposed constraint is that,

unlike local validity, validity proper must be characterized as necessary preservation of a truth-like property that is not bivalent.

Cobrerros's region-valuation semantics is like the supervaluationist's in that his models consist of collections of *precisifications*, in the sense already described; and, like the supervaluationist's, Cobrerros's semantics involves two notions of truth, in his case: bi-truth at a precisification and region-truth at a precisification (rather than super-truth in a model). Models also come along with an "accessibility" relation between precisifications, which is required to be reflexive. A sentence is said to be *regionally true* (I have called that "region-true") at a precisification when it is bi-true at every precisification that is accessible to that precisification; *region-false* at a precisification when it is bi-false at every accessible precisification.

Bi-truth at a precisification is defined for the standard logical operators in the usual classical way (a conjunction is bi-true just in case both of its conjuncts are bi-true, a disjunction is bi-true just in case one of its disjuncts is bi-true, *et cetera*). For the "definitely" operator, bi-truth at a precisification is defined as bi-truth at every accessible precisification. Given this condition, a sentence turns out to be region-true at a precisification just in case its definitization is bi-true at that precisification, region-false at a precisification just in case the definitization of its negation is bi-true at that precisification.

Region-valuation at a precisification: A sentence is region-true at a precisification when it is bi-true at every accessible precisification; region-false at a precisification when bi-false at every accessible precisification.

'Definite' bi-truth at a precisification: $D\Phi$ is bi-true at a precisification when Φ is bi-true at every accessible precisification.

Immediate consequences: A sentence is region-true at a precisification just in case its definitization is bi-true at the precisification, region-false at a precisification just in case the definitization of its negation is bi-true at that precisification.

The regionalist then characterizes validity as preservation of region-truth at a precisification.

Regional validity: An inference from Φ to Ψ is *regionally valid* when Ψ is region-true at every precisification at which Φ is region-true; equivalently, when the bi-truth of Φ at every precisification accessible to a precisification p in a model guarantees the bi-truth of Ψ at every precisification accessible to p .


Region-valuation semantics is a species of supervaluation semantics in that its primary notion of truth – region-truth – is identified with bi-truth at each member of a certain collection of precisifications.

In standard supervaluation semantics, bi-truth and bi-falsity are the bivalent truth values that sentences have at a precisification, and super-truth and super-falsity are the gappy truth values that sentences take in a model. Unlike standard supervaluation semantics, Cobreros’s regionalism does not assign any sort of truth value to sentences *in a model*. A sentence has a bi-value at each precisification (Cobreros calls the bi-values “1” and “0”), and it has its region value (if it has one) *at a precisification*. There is no sort of truth value that a sentence has relative to a model itself. This, I think, is one of the main failings of region-valuationism.

Region-valuation semantics has two crucial features. First, region-truth – the preservation of which is required for validity – is not bivalent at a precisification: a sentence Φ might be bi-true at some precisifications that are accessible to p while bi-false at others. Second, **D-intro** is not regionally valid: we might have Φ but not $D\Phi$ be region-true at a precisification p in a model when Φ is bi-true at every precisification that is accessible to p but $D\Phi$ is not bi-true at every precisification that is accessible to p because one of *those* precisifications has a precisification that is accessible to it at which Φ is not bi-true. Note that accessibility must not be transitive in this model. The failure of **D-intro** requires that accessibility not be transitive in general.

11.6 Regional avoidance of the gap-principles paradox

To complete his defense of gap principles within the supervaluationist framework, Cobreros presents us with a model in which there is a precisification at which the gap principles for “suitable” (let us say) are all region-true, compatibly with there being a sorites series that has something that is absolutely definitely suitable at one end and something that is absolutely definitely unsuitable at the other.



Here, $\llbracket S \rrbracket_{bi}$ is the bivalent extension of S at a precisification – these are the things of which S is bi-true. Reflexivity of accessibility is assumed, but not depicted.

At the precisification p_1 it is region-true that the first object, 1, is absolutely definitely suitable and that the last object, 4, is absolutely definitely unsuitable. For a length-four sorites series there are three gap principles that would come into play in generating a gap-principles paradox for the series. At precisification p_1 , all three of these gap principles are region-true.

Regional truths at p_1 :

- 1. $D^n S(1)$ (for any n)
- 2. $D^n \neg S(4)$ (for any n)
- 3. **G1:** $DS(3) \rightarrow \neg D\neg S(4)$.
- 4. **G2:** $DDS(2) \rightarrow \neg D\neg DS(3)$.
- 5. **G3:** $DDDS(1) \rightarrow \neg D\neg DDS(2)$.

11.7 Two more notable features of Cobrerós’s model

It is worth laying out at least two other notable features of the model. First, let us write down the other precisifications in the model at which the above sentences are region-true/region-false.


Region-verifiers/region-falsifiers:

- 1. $D^n S(1)$ ($p_1, p_2, p_3, p_4/-$)
- 2. $D^n \neg S(4)$ ($p_1, p_2, p_3, p_4/-$)
- 3. **G1:** $DS(3) \rightarrow \neg D\neg S(4)$. ($p_1, p_3/p_4$)
- 4. **G2:** $DDS(2) \rightarrow \neg D\neg DS(3)$. ($p_1, p_4/p_3$)
- 5. **G3:** $DDDS(1) \rightarrow \neg D\neg DDS(2)$. ($p_1, p_2, p_3, p_4/-$)

The thing to notice here is that none of the precisifications other than p_1 region-verifies all of the gap principles. At p_2 , neither **G1** nor **G2** is region-true. At p_3 , **G2** is region-false. At p_4 , **G1** is region-false.

Next, let us depict the regional extension gaps of S – these are the things of which S is neither region-true nor region-false.

“Suitable” extension gaps in Cobrerós’s model:



Here, $\{S\}_r$ is the regional extension gap of S at a precisification – the things of which S is neither region-true nor region-false. As before, reflexivity of accessibility is assumed, but not depicted.

If n is in the extension gap of S at a precisification, then $BS(n)$ is bi-true at that precisification. So $BS(x)$ is bi-true of 2 and 3 at p_1 , of 3 at p_2 , and of nothing else at any other precisification.

11.8 Good models and bad models

If we are told that a certain form of inference is not in general valid on the grounds that there are models in which the premises of the inference form are true but in which its conclusion is not true, then we can sensibly and reasonably ask whether the invalidating models are ones that we should care about. For instance, if I am developing a semantic theory of English, I might well ignore models that did not verify the sentence “everything that is tiny is small.” If one thing I am interested in doing is capturing the fact that an inference from “it is tiny” to “it is small” is a good inference, then I should not care about models in which the extension of “tiny” is not a subset of the extension of “is small.”

The models that we care about as they relate to the questions of validity of inferences involving vague terms we will call the “good models.” The rest are the “bad models.”

11.8.1 Transitivity

If our theoretical understanding of the accessibility relation in region-valuation semantics requires that relation to be transitive, then it should be irrelevant to us that there are models that invalidate a given form of inference but that have a nontransitive accessibility relation. What matters is whether there are invalidating models that conform to our theoretical understanding of the accessibility relation between precisifications as it relates to our understanding of the phenomena of vagueness. But if accessibility is required to be transitive, then **D-intro** is regionally valid – it preserves region-truth at any precisification. If Φ is region-true at p , then $D\Phi$ must be region-true at p . For suppose it were not, then there would be some precisification p' accessible to p at which $D\Phi$ was bi-false. But then there would have to be some precisification p'' accessible to p' at which Φ were bi-false. Transitivity ensures that p'' is accessible to p , so p'' is a precisification that is accessible to p at which Φ is bi-false. But then Φ cannot be region-true at p , as we supposed it to be.

We have yet to hear any explanation of what might make one precisification but not another accessible to any given precisification. We have yet to hear any theoretical grounding of the accessibility relation. The success of Cobreros’s project of defending the gap principles depends on accessibility’s not being a transitive relation. What understanding,

then, could ground an accessibility relation that was not necessarily transitive? Why should we accept that nontransitive models are good models?

One idea, in keeping with the understanding of precisification as *narrowing* a truth-gap, if not closing it all the way, would be that one precisification is accessible to another just in case its region-truths are a superset of those of the first and its region-falsehoods are a superset of those of the first; in other words, that its region-value gaps were a subset of those of the first. Since *being a superset of* is a transitive relation, accessibility should be transitive as well.

Transitive *narrowing* accessibility: One precisification is accessible to another just in case it narrows the other's region gap.

The most salient alternative understanding of accessibility that would allow for it to be nontransitive would be that of narrowing a gap, but not by more than some given amount. Cobreros's model meets this condition if we take the limit of narrowing to be one member. The extension gap of S at p_2 contains one less element than the extension gap of S at p_1 does, but the gaps at p_3 and p_4 contain two fewer elements, and so are not by these lights accessible to p_1 , although they are to p_2 .

Nontransitive *slight-narrowing* accessibility: One precisification is accessible to another just in case it narrows the other's region gap, but not by much.

The regionalist can talk this way of course, but that talk should not convince us that *being a slight gap-narrowing* is an understanding of an accessibility relation that should matter for vagueness. Analogously in modal logic: we can talk as if one possibility is alternative to another, though not to a third, but this by itself should not convince us that what might have been the case depends on what in fact is the case. There might be other arguments to the effect that this is so, but mere structural coherence does not by itself constitute such an argument.

11.8.2 Monotonicity validates iteration

There is a straightforward argument against the *slight-narrowing* conception of accessibility: it validates a certain restricted iteration principle. Iteration, remember, was the following principle:

Iteration: $D\Phi \rightarrow DD\Phi$.

When Φ is required to be either an atomic predication, a negation of an atomic predication, or any definitizationⁿ of such, iteration is valid given the slight-narrowing conception of accessibility.¹² The reason for this is that the nontransitive conception of accessibility requires monotonicity of narrowing in the sense that the extension gap at any precisification will be a subset of that of any of its ancestors. This requires that extensions and anti-extensions grow monotonically from precisification to precisification along an accessibility chain.

Monotonicity: If p “sees” p' , then the extensions and anti-extensions of atomic predicates at p are subsets of those at p' .

We had identified a serious problem with *iteration*:

Borderline definite but not definite: $\models BD\Phi \rightarrow \neg D\Phi$.

If we assume that justified assertability is transferred from antecedent to consequent of a valid conditional, then we could be justified in asserting that a tree was borderline definitely suitable only if we were justified in asserting that it was *not* definitely suitable. We found this problematic because we should not simultaneously be justified in asserting both that things *could have gone either way* and also that they *did not go that way*.

11.8.3 Truth at a precisification and truth in a model

The regionalist uses two different truth-like properties in his semantics: *bi-truth* and *region-truth*. Each of these kinds of truth is relativized to a precisification in a model. We have not been presented with a regionalist notion of *truth in a model*. Let us again make an analogy with the standard semantics for modal logic. In modal logic, models should come along with a designation of one world as the actual world. Truth in a model would then be identified with truth at the actual world of that model. If validity is preservation of truth in a model, then on the actual-world semantics for modal logic, an inference from Φ to Ψ counts as valid when any model in which Φ is true at its actual world is one in which Ψ is true at that actual world. But for the purposes of characterizing validity, the need for a designation of an actual world in each model can be dispensed with. For every model and every world in that model, there is some model that differs from the first only by designating that other world as actual. Given this fact, an inference from Φ to Ψ will be valid on the first definition just in case it is valid according to the following definition: an

inference from Φ to Ψ is valid just in case in every model, any world at which Φ is true is one at which Ψ is true. This is local validity. The new definition – equivalent to the original one – makes no use of a notion of truth in a model, only of truth at a world in a model. The equivalence of the two definitions should not, however, let us lose sight of the fact that we do need a notion of truth in a model (truth at its actual world) if what we want to give a characterization of is not merely validity but also truth as a matter of fact.

Some modal claims are true as a matter of fact – actually true, true in reality – while others are false as a matter of fact – actually false, false in reality. When we make a modal inference, we should be concerned about whether our inference preserves truth from premises to conclusion no matter what things are actually like, no matter how things are in reality. Since the real truth of modal claims depends on what things might in fact have been like, the real model has to be accurate with respect to the whole space of its possible worlds. This is to be the space of all and only those worlds that accurately represent a way things might have been. Without a notion of truth in a model that is identified with truth at the actual world of that model, we have lost sight of the substance of our main concern to make only valid inferences.

Similarly, the supervaluationist should take one of his admissible models to be the *real* model. For the standard supervaluationist, this will be one in which a claim is super-true just in case it is *really* true – true in reality, true as a matter of fact; and in which a claim is super-false just in case it is false in reality, false as a matter of fact. Since it is true as a matter of fact that man *X* is a tall man given that he is 2 meters tall, it will be super-true in the real model that *X* is a tall man: in the real model, *X* will be in the extension of “tall man” in every precisification.

And since it is false as a matter of fact that man *Y* is a tall man given that he is 1 meter tall, it will be super-false in the real model that *Y* is a tall man: in the real model, *Y* will be in the anti-extension of “tall man” in every precisification.

And since, as a matter of fact, it is neither true nor false (we will suppose) that man *Z* is a tall man given that he is 175 centimeters tall, it will be neither super-true nor super-false in the real model that *Z* is a tall man: in the real model, *Z* will be in the extension of “tall man” in some precisifications but not in others.

For the standard supervaluationist, the real model is the one that contains all and only those precisifications that are admissible given the actual facts (for example that Barack Obama is 187 centimeters tall),

and given the actual extensions of our vague predicates (for example, that Barack Obama is in the extension of “tall”). For the standard supervaluationist, the claims that are true as a matter of fact are the ones that are true in the real model. Those are the ones that are super-true in the real model, bi-true at every precisification in the real model.

But what about the region-valuationist? Some claims involving vague predicates are true, others are false. How does the regionalist represent this? We are right to demand from him an answer to the question how the notion of regional-truth at a precisification in a model aligns with the desired notion of *truth*. Without an answer to this question, we can have no understanding of why the joint verification at a precisification of some sentences would demonstrate the actual compatibility of those sentences. In our case the sentences in question are the gap principles and it is the significance of their joint verification at p_1 that we want to understand.

The obvious way for the regionalist to answer our question is to say that there is some model and some precisification in that model which is such that all and only the *truths* are region-true at that precisification, and all and only the *falsehoods* are region-false at that precisification. This makes salient the possibility that the regionalist include in each model a designated precisification and then define truth in a model as regional-truth at the designated precisification.

This development of region-valuation semantics brings a curiosity to the fore that should have been there all along: what is the significance of bi-truth at a precisification? In particular, how are we to understand the significance of bi-truth at the designated precisification of a model? For the supervaluationist, there is only one notion of truth at a precisification in a model: bi-truth. And there is no precisification in any model at which this kind of truth accurately reflects real truth, truth in reality. Real truth is not bivalent for the supervaluationist, but bi-truth at a precisification is. For the supervaluationist, it is a model rather than any precisification in a model that purports to represent the way things are. In the case of the real model, it represents the way things are by including all of the precisifications that admissibly precisify our predicates by closing their actual extension gaps in admissible ways.

But for the regionalist, there *is* a notion of truth at a precisification in a model that does purport to represent real truth, namely regional-truth. This representation is accurate at the designated precisification of the real model. But if in some model there is some designated precisification that represents how things are in reality, why would there be any particular

bi-valuation associated with that precisification? What would make one bi-valuation for our vague predicates rather than some other bi-valuation be the one that represents reality? Just as we need some theoretical basis for the accessibility relation before we honor nontransitive ones, we need some theoretical basis for bi-valuation that makes sense of there being some special, accurate bi-valuation for our vague predicates that reflects how things really are.

11.9 From Neg-B-intro to D-intro

The point of region-valuation semantics was to save gap principles from paradox. The success of this endeavor required that **D-intro** not in general be valid. We gave an argument that supervaluationists are committed to a semantics for *D* that validates **D-intro** (pp. 228–9). That argument turned on the idea that supervaluationists want to identify *being a borderline case of a predicate* as falling into the extension gap of that predicate. But it also depended on the identification of *truth* with super-truth. The regionalist rejects the latter, but should accept the former. As long as he does, though, we can serve him an independent argument for why supervaluationists (and regionalists along with them) are in fact committed to **D-intro** – given their antecedent and independent understanding of what it is to be a borderline case. If my argument is a good one, then since region-valuation semantics does not validate **D-intro**, it is not compatible with its own theoretical foundations.

11.9.1 The argument

Supervaluationists, and their regionalist cousins, say that when we apply a vague predicate to one of its borderline cases, our claim lacks a truth value – it is neither true nor false. Correlatively, they must say that if we truly apply a vague predicate to an object, then our claim does not lack a truth value – so the object is not a borderline case. If it is true that Juan is clever, then it is false, they must say, that John is a borderline case of clever. In other words, they are committed, although I have never seen it spelled out, to the following entailment principle:

– **B-intro**: $\Phi \models \neg B\Phi$.

Supervaluationists (among many others, including those who reject truth-value gaps) regard borderline-case claims as definitionally

equivalent to claims about what it is for an object to be a definite case of a given vague predicate:

B-def: $B\Phi \leftrightarrow_{def} \neg D\Phi \wedge \neg D\neg\Phi.$

But given some uncontroversial principles, there is an equivalent definition of definiteness in terms of being a borderline case:¹³

D-def: $D\Phi \leftrightarrow_{def} \Phi \wedge \neg B\Phi.$

These definitions have their pros and cons. The first definition gives a more intuitive expression of the connection between being a borderline case and being a definite case of a vague predicate: to be a borderline case of “suitable,” for example, is to be neither definitely suitable nor definitely unsuitable.

But here is a consideration in favor of the second definition. Since we have an understanding of what it is to be a borderline case that is independent of our understanding of what it is to be a definite case (or a *clear case*), we can clarify or add to our understanding of what it is to be a definite case by seeing how it is to be defined when “borderline” is taken as the primitive notion. That we do have an independent understanding of “borderline” is evidenced by the strangeness many feel when first confronted with **D-def**, despite its equivalence to **B-def**.

The second definition does not, however, give a very intuitive expression of the connection between definiteness and borderline status. In fact, it gives a counter-intuitive expression of that connection. If I were to say, “that is suitable, and it is not a borderline case,” I would sound and feel as if I had repeated myself. Why? – because I would not take myself to have justification for the first conjunct unless I took myself to have justification for the second. If I did not take myself to be justified in saying that Picasso was not borderline talented, then I would not take myself to be justified in saying that he was in fact talented. The supervaluationist can offer a different reason in terms of truth rather than justification: I would sound and feel as if I had repeated myself in saying “that is suitable, and it is not a borderline case” because, they must say, the second disjunct is *entailed* by the first. On their view, Φ entails $\neg B\Phi$, at least in the case where Φ is an atomic predication.

Given that Φ entails $\neg B\Phi$ (\neg **B-intro**), the definition displayed in **D-def** straightforwardly yields the following entailment:

D-intro $\Phi \models D\Phi.$

This follows straightforwardly from **¬B-intro** and **D-def** because the validity of the biconditional in **D-def** ensures – given $D\Psi \rightarrow \Psi$ – that (i) $\Phi \wedge \neg B\Phi \models D\Phi$. Given also our initial observation that $\Phi \models \neg B\Phi$, it follows that (ii) $\Phi \models \Phi \wedge \neg B\Phi$. Then since entailment is transitive, we have from (i) and (ii) that $\Phi \models D\Phi$.

We can also lay out the argument in a way that does not appeal directly to **D-def** but rather to **B-def**, the more familiar biconditional. But we should be long-winded about this. When we make an argument to supervaluationists, we have to be long-winded since they reject all sorts of common principles of logical reasoning – for example, they do not think that validity is preserved under substitution of equivalents, since they think that Φ is always equivalent to $D\Phi$, and they agree that $\Phi \rightarrow \Phi$ is valid, but they do not think that $\Phi \rightarrow D\Phi$ is valid. So our arguments to them cannot compress any of the usual steps. Being perspicuous and long-winded gives the supervaluationist the benefit of having more opportunities to balk at our argument. Let us therefore perspicuously lay out our argument as follows:

(1) $\Phi \models \neg B\Phi$ (**¬B-intro**).

Truth-value gap theorists should commit to this because borderline cases, on their view, require truth-value gaps. So the lack of a truth-value gap – in this case, the truth of Φ – precludes being a borderline case.

(2) $\Phi \models \Phi \wedge \neg B\Phi$.

This follows from (1).

(3) $\models B\Phi \leftrightarrow \neg D\Phi \wedge \neg D\neg\Phi$. (**B-def**)

This states the validity of the standard definition of what it is to be a borderline case.

(4) $\models D\neg\Phi \rightarrow \neg\Phi$. (**T_D**)

T_D is an instance of the principle about definiteness, “ $DP \rightarrow P$,” that corresponds to the **T**-schema of modal logic.

(5) $\models (B\Phi \leftrightarrow (\neg D\Phi \wedge \neg D\neg\Phi)) \rightarrow ((D\neg\Phi \rightarrow \neg\Phi) \rightarrow ((\Phi \wedge \neg B\Phi) \rightarrow D\Phi))$.

This says, in effect, that if the definition of B in terms of D holds, then if **T_D** also holds then the right-to-left direction of our definition of D in terms of B also holds. It is a tautology of the following form: $(P \leftrightarrow (\neg Q \wedge \neg R)) \rightarrow ((R \rightarrow \neg S) \rightarrow ((S \wedge \neg P) \rightarrow Q))$.

(6) $\models (D\neg\Phi \rightarrow \neg\Phi) \rightarrow (\Phi \wedge \neg B\Phi \rightarrow D\Phi)$.

This follows from (3) and (5).

(7) $\models (\Phi \wedge \neg B\Phi) \rightarrow D\Phi$.

This is the right–left direction of **D-def**. It follows from (6) and (4).

(8) $\Phi \models D\Phi$

This follows from (2) and (7).

Given a commitment to the validity of all tautologies, the supervaluationist can discharge the commitment to **D-intro** only by rejecting either (i) \neg **B-intro**, (ii) **B-def**, or (iii) **T_D**, the incontrovertible conditional corresponding to the T-schema of modal logic: " $D\Phi \rightarrow \Phi$." We tacitly appealed only to meager classical principles of reasoning: (a) reflexivity ($A \models A$); (b) Modus Ponens ($P, P \rightarrow Q \models Q$); (c) the closure of entailment under conjunction (if $P \models Q$ and $P \models R$, then $P \models Q \wedge R$); and (d) generalized transitivity (if $\Delta \models \gamma$ for every $\gamma \in \Gamma$, and $\Gamma \models P$, then $\Delta \models P$). Presumably, these are not to be given up.

Given that the above argument from \neg **B-intro** to **D-intro** appealed only to principles that the regionalist accepts, it should not be surprising that region-valuationism invalidates \neg **B-intro**.¹⁴ But if someone does not think that the truth of a predication precludes the subject of that predication from being a borderline case of the predicate in question, then it is unclear why that person would have a truth-value gap theory of borderline cases at all. If the point of supervaluationism in its standard, revised, or regionalist incarnations is to uphold a truth-value-gap theory of borderline cases, then the supervaluationist and her kin should accept \neg **B-intro**, and therefore **D-intro** along with it, and should therefore give up her gap principles as providing her with the solution to the problem of higher-order vagueness.

11.10 Conclusion

Cobrerros developed the regionalist version of supervaluationism in order to protect gap principles from paradox. The reason for accepting the gap principles was to enable an account of higher-order vagueness that paralleled the truth-value-gap account of first-order vagueness. At the first order, the lack of a sharp-appearing boundary between the suitable trees and the unsuitable ones was to be explained by the intervention of some borderline suitable trees between those that were definitely suitable and those that were definitely unsuitable. At the second order, the lack of a sharp-appearing boundary between the definitely suitable trees and the borderline suitable trees was to be explained by the intervention of some trees that were borderline definitely suitable. In general, at the n th order, the lack of a sharp-appearing boundary between the trees that

were definitelyⁿ suitable and those that were not was to be explained by the intervention of some trees that were borderline definitelyⁿ suitable.

The (section 11.2) argument for the inconsistency of gap principles with mundane facts turned on the validity of any inference from Φ to $D\Phi$. We argued (section 11.3) that standard supervaluationists must accept **D-intro**. But this argument turned on the identification of *truth* with super-truth – something that the regionalist rejects. We gave some general, theoretical rather than technical, reasons for rejecting Cobreros's argument for the compatibility of gap principles with mundane facts (section 11.8). These arguments were directed at two central tenets of regionalism: (i) that some precisification might be accessible to one precisification but not to another – that precisification is precisification-relative – and that this accessibility is not transitive and (ii) sentences have their region-truth values relative to precisifications rather than to models. We then gave a general argument (section 11.9) that truth-value gap theorists of borderline cases are in general committed to the validity of **D-intro**. Like the corresponding argument of section 11.3, this argument depended on the validity of **–B-intro**.

The truth-value-gap theorist of borderline cases has no right to reject **–B-intro**, the principle that the truth of a vague predication precludes the subject of that predication from being a borderline case of the predicate in question. Commitment to that principle, however, commits one to the validity of **D-intro**, the principle that the truth of a vague predication requires that its subject be a definite case of the predication in question. This principle, however, renders gap principles incoherent. The upshot is that the truth-value-gap theorist of borderline cases should give up his gap principles. But gap principles state that the explanation of the lack of sharp-appearing boundaries at higher orders of vagueness be the same as that at the first order. If the explanation is not coherent at higher orders of vagueness then it must not be accepted at the first order. The recommendation is that that no theorist accept a truth-value-gap account of borderline cases at even the first order of vagueness.¹⁵

Notes

1. I take the canonical supervaluation semantics to be that described in Kit Fine's (1975) formative presentation and also that defended in Rosanna Keefe's (2000) more recent book.
2. For presentation of the gap-principles paradox, see Fara (2003).
3. By canonical supervaluation semantics, I mean that presented in Fine (1975) and Keefe (2000).

4. The comparative relations in question will not be only those associated with gradable adjectives and gradable adverbs but also those that can be associated with nouns, such as “is closer to being undeniably a heap than.”
5. If we liked, we could spell this out more explicitly: consider a “borderline” claim $BS(x)$ that ascribes borderline suitability to a given tree. Our question is what the bi-value of $BS(x)$ is at an arbitrary precisification p in some arbitrary model. Suppose $BS(x)$ is bi-true at p . Suppose also that the bare claim $S(x)$, which ascribes suitability to the tree, is bi-true at p . Then there are two cases: (i) $S(x)$ is bi-true at every precisification in the model or (ii) true at some precisifications but false at others. Only in case (ii) could things have gone either way for x . So since cases (i) and (ii) exhaust all possibilities, $BS(x)$ is bi-true at p in this case, when $S(x)$ is bi-true at some precisifications but bi-false at others. A similar argument can be given for the converse of this. And then a parallel argument can be run in the case where $S(x)$ is bi-false at p .
6. This states the factivity of D . Just as nothing can be known without being true and nothing can be necessary without being true, nothing can be definite without being true.
7. Here is why. If $D\Phi$ is either super-true or super-false, then if it is bi-true at any precisification, it is bi-true at every precisification. But this is the condition for the bi-truth of $D\Phi$ at any precisification, and hence at every precisification, since D -prefixed formulas are true at all precisifications or at none. So whenever $D\Phi$ is true at a precisification, $DD\Phi$ is true at that precisification.
8. Here is why. Given iteration, $\models D\Phi \rightarrow DD\Phi$, but $\models BD\Phi \rightarrow \neg DD\Phi$, given the definition of $B\Phi$ as $\neg D\Phi \wedge \neg D\neg\Phi$. So $BD\Phi \wedge D\Phi$ is inconsistent given iteration. But given the law of excluded middle, $\models D\Phi \vee \neg D\Phi$, so $\models BD\Phi \rightarrow \neg D\Phi$.
9. The Euclidean principle of modal logic, $\diamond\Phi \rightarrow \Box\diamond\Phi$ can be thought of as saying that if I can see a p -world, then anyone I can see can see a p -world. (Here “ x sees y ” is our shorthand for “ y is accessible to x .”) The *weak Euclidean principle* says that if I can see both a p -world and a not- p world, then anyone I can see can see both a p -world and a not- p world.
10. Both Timothy Williamson (1994: Ch. 5) and Keefe (2000: Ch. 8) explain revised supervaluationism by analogy to accessibility-relation semantics for modal operators.
11. Consult Ch. 5 of Williamson (1994).
12. A definitization” of Φ is the sentence consisting of Φ preceded by n occurrences of D .
13. The equivalence can be derived from “ $D\Phi \rightarrow \Phi$ ” and “ $B\Phi \leftrightarrow B\neg\Phi$,” along with classical logic.
14. For suppose we have Φ being bi-true everywhere in the region of some precisification p . Let there be a precisification p' in the region of p . Although we have ensured that Φ is bi-true everywhere in the region of p , there might nonetheless be a precisification in the region of p' at which Φ is bi-false. This suffices to render $B\Phi$ bi-true at p' , which in turn suffices to render $\neg B\Phi$ not region-true at p .
15. I am grateful to Paul Égré and Nathan Klinedinst for detailed comments on this chapter and also to Pablo Cobreros for writing such a stimulating and ingenious essay which was thoroughly enjoyable to engage with and which deepens our understanding of the supervaluational approach to vagueness.

References

- Cobrerros, P. (2011). Supervaluationism and Fara's argument concerning higher-order vagueness. In Égré, P. and Klinedinst, N. (eds), *Vagueness and Language Use*. Palgrave Macmillan.
- Fara, D. G. (2003). Gap principles, penumbral consequence and infinitely higher-order vagueness. In Beall, J. C. (ed.), *Liars and Heaps: New Essays on Paradox*, pp. 195–221. Oxford University Press. Published under the name "Delia Graff."
- Fine, K. (1975). Vagueness, truth and logic. *Synthese*, 30: 265–300.
- Keefe, R. (2000). *Theories of Vagueness*. Cambridge University Press, Cambridge, England.
- Williamson, T. (1994) *Vagueness*. Routledge, London.

12

Vagueness and Practical Interest*

Paula Sweeney and Elia Zardini

12.1 Introduction and overview

Take the vague expression “tall.” Two outstanding *phenomena* of its vagueness seem to be:

SORITES SUSCEPTIBILITY. One is inclined to accept the *soritical principle*:
(S₀) For every x and y , if x is tall and y is 1 inch shorter than x , then y is tall,

and

IGNORANCE OF CUTOFFS. For every x and y , if y is 1 inch shorter than x , one does not know that x is tall and y not tall.

It is a very widespread assumption in the vagueness debate that a successful theory of vagueness must account for both SORITES SUSCEPTIBILITY and IGNORANCE OF CUTOFFS.

A leading project in that debate is to account for either or both of these phenomena by appealing to a peculiar *shiftiness* exhibited by vague expressions (see e.g. Kamp 1981, Raffman 1994, Soames 1999, Fara 2000, Gaifman 2002, Shapiro 2006). We can usefully label all such theories “*contextualist theories*.” The starting point of contextualist theories is the observation that vague expressions are typically *context-dependent*. Given the crucial role played in contextualist theories by the notion of context-dependence, we start with a brief summary of the nature and varieties of context-dependence. In general, the context-dependence of an expression ε can be neutrally characterized as:

(CD) An expression ε is context-dependent iff, for some contexts c_0 , c_1 , c_2 and c_3 , the *extension*¹ of ε as uttered in c_0 is correctly

assessed in c_1 to be X while the extension of ε as uttered in c_2 is correctly assessed in c_3 not to be X .

Notice that (CD) does not imply that the extension of ε varies in virtue of the *content* expressed by ε varying from c_0 to c_2 : for all (CD) says, ε as uttered in c_2 could express the same content as ε as uttered in c_0 does and yet vary in extension. More generally, (CD) does not imply that the extension of ε varies in virtue of *features of the contexts of utterances* c_0 and c_2 : for all (CD) says, c_2 could have all and only the features had by c_0 (indeed, could be numerically identical to c_0) and yet ε could vary in extension.

This neutrality of (CD) accords well with contemporary wisdom in the philosophy of language, which distinguishes at least four ways in which the extension of an expression ε can vary across contexts. Suppose that the right-hand side of (CD) holds. Then, for what we may call “the *standard contextualist*,” that fact holds in virtue of ε as uttered in c_0 expressing a content different from that expressed by ε as uttered in c_2 ; for what we may call “the *nonindexical contextualist*,” that fact holds rather in virtue of ε as uttered in c_0 and as uttered in c_2 expressing a single content that ε as uttered in c_0 brings to bear on a circumstance different from that on which ε as uttered in c_2 brings it to bear; for what we may call “the *truth relativist*,” that fact holds rather in virtue of ε as uttered in c_0 and as uttered in c_2 expressing a single content that is correctly assessed in c_1 to determine X as extension and correctly assessed in c_3 not to determine X as extension; finally, for what we may call “the *content relativist*,” that fact holds rather in virtue of ε as uttered in c_0 and as uttered in c_2 being correctly assessed in c_1 to express a single content different from that which they are correctly assessed in c_3 to express (this taxonomy is in many respects very rough, but will do well enough for our purposes; see Weatherson 2009: 339–42 for a similar one).

In this chapter, we wish mostly to focus on a particular type of contextualist theory, according to which the context-dependence that, at least partially, generates the phenomena of vagueness has its source in the variation of our *practical interests* (henceforth “interests”) – in other words, according to which the phenomena of the vagueness of an expression are, at least partially, due to the *interest relativity* of its correct application (we call such a type of theory “the *IR-theory*”). The IR-theory has been most influentially developed and defended by Delia Graff Fara in a series of papers (Fara 2000, 2008a, b). Accordingly, in this chapter we largely focus on Fara’s very specific version of the IR-theory – as we demonstrate, in the explanations given by contextualist theories, the devil is often in the details, and Fara is very usefully quite explicit about

many of these (so, for example, we will see in section 12.6 where Fara's version of the IR-theory exactly sits in terms of the four-way semantic divide sketched above).² We suspect however that something along the lines of the details we discuss is often implicit in the thoughts of those attracted by the IR-theory and by contextualist theories more generally, and so hope that our discussion will largely retain an interestingly wide scope. In effect our observations work at different levels of generality, some relevant only to the specifics of Fara's version of the IR-theory, others relevant to all contextualist theories of a certain type. To anticipate our findings, we believe that one encounters fundamental difficulties at all these levels, and it is the main purpose of this chapter to elaborate on what we think these various difficulties are.

In keeping with the spirit of this volume, we start with some general linguistic theses fairly common for contextualist theories, move to see how they are supposed to help with the (nonlinguistic) phenomena of vagueness and finally examine the specific way in which Fara justifies and understands the linguistic theses themselves. In detail, the rest of the chapter is organized as follows. Section 12.2 critically discusses some arguments given by Fara in favor of the claim, arguably required by many contextualist theories, that the semantic context-dependence of a certain class of adjectives cannot wholly be captured in terms of variation of comparison classes. Moving on to the theory of vagueness proper, section 12.3 argues against the independent plausibility of a salient-similarity constraint again assumed, in some form or other, by many contextualist theories. Section 12.4 ascends to an even higher level of generality and shows that, contrary to what seems to be presupposed by many such theories, an alleged consequence of that constraint still falls very much short of yielding satisfactory explanations of *SORITES SUSCEPTIBILITY* and *IGNORANCE OF CUTOFFS*. Section 12.5 descends to a much lower level of generality and criticizes Fara's own attempt at justifying the salient-similarity constraint on the basis of a certain semantic hypothesis about the relevant class of adjectives and of considerations pertaining to our interests. Section 12.6 situates Fara's version of the IR-theory in the four-way semantic divide sketched above and compares its pros and cons against standard-contextualist theories along two axes: how well they score in dealing with arguments from verb-phrase ellipsis and how well they manage to preserve our pre-theoretic conception of what tallness (and many other properties expressed by vague expressions) depends on. In the face of the overwhelmingly negative findings of the previous sections for a wide class of contextualist theories, section 12.7 concludes by sketching two alternative approaches to *SORITES SUSCEPTIBILITY* and *IGNORANCE OF CUTOFFS*, each favored by (exactly) one of the authors.

12.2 Comparison classes and kinds

Fara's version of the IR-theory is restricted to *gradable adjectives* (henceforth, "GAs") and requires some quite substantial claims about their semantics (presupposing but significantly going beyond the main tenets of Kennedy's (1999) theory).³ Before moving on in the remaining sections to see how Fara builds on these claims in her theory of vagueness, we pause in this section to expose and assess some of the more general points she makes about the semantic context-dependence of GAs – in particular, her claim that such context-dependence cannot wholly be captured in terms of variation of *comparison classes*.⁴ In arguing for this claim, Fara (2000: 55–7) offers two different arguments.

In the first argument, which proceeds by way of *example*, we are asked to imagine the following case. We are in the West End of London and there are two different auditions taking place. One of the auditions is to find an actor suitable for playing Mikhail Gorbachev, the other is to find an actor suitable for playing Yul Brynner. Fara claims, a bit dubiously in our view, that "bald" can be truly asserted of an actor in the context c_0 of the first audition and truly denied of the same actor in the context c_1 of the second audition. Furthermore, she claims that this is so even though the comparison class (the class of men) does not vary from c_0 to c_1 .⁵ We disagree with the latter claim. At best, the case seems only to show that the *comparison classer* (as we may call Fara's dialectical opponent in this section) should assume that the relevant comparison class in the Brynner audition is a very restricted one (roughly, one containing samples from men with no hairs on their scalp to men who do not have a much larger number of hairs on their scalp than Gorbachev). That assumption seems to us independently plausible: that is the kind of comparison class that makes salient the peculiar way in which Brynner is bald.

The second argument, which proceeds in a more *theoretical* fashion, is more challenging. It goes as follows:

- (i) The comparison class of *Fs* is not the set of things that happen to be *F*. It is rather, roughly, the set of things that are *typically F*;
- (ii) The notion of a typical *F* requires the *Fs* to form a *kind*;
- (iii) There are cases of context-dependence of GAs where no relevant kind is available;
- (iv) Therefore, there are cases of context-dependence of GAs that cannot be explained in terms of variation of comparison classes.

We may take it that the argument from (i)–(iii) to (iv) is valid. Fara motivates (i) with the following example. Imagine that, by some tragic accident, all but some very short basketball players are killed so that their average height drops drastically and Shorty is the tallest surviving one. Fara claims that, intuitively, a post-accident utterance of:

(SHORTY) Shorty is tall for a basketball player⁶

would still be false. While we agree that one *can* understand an utterance of (SHORTY) this way, we do not think that that does anything to support (i). Let us explain. The way we see it, in the peculiar situation just described there are (at least) two salient sets: roughly, the set of pre-accident basketball players and that of post-accident ones. These two sets are everything one needs in order to explain the data, and to do so in a more satisfactory way than a friend of (i). On the one hand, there arguably is an understanding of a post-accident utterance of (SHORTY) under which it is *true*. That understanding can be explained by the comparison classer by supposing that its truth condition is that Shorty is tall in comparison to the set of *post-accident* basketball players, whereas it is very unclear how it could be explained by a friend of (i). On the other hand, there admittedly is also an understanding of a post-accident utterance of (SHORTY) under which, in the original situation described by Fara, it is *false*. That understanding can be explained by the comparison classer by supposing that its truth condition is that Shorty is tall in comparison to the set of *pre-accident* basketball players, rather than in comparison to the set of typical basketball players, as the friend of (i) would have it. And such understanding can in turn be easily accommodated by the independently plausible assumption that NPs carry an implicit time variable, which in the case of (SHORTY) can be contextually assigned a post-accident time (thereby generating the first reading) or a pre-accident time (thereby generating the second reading).⁷

Setting aside this particular example, we can also note that there are many GAs for which it should be uncontroversial that the comparison class of *Fs* just is the set of things that happen to be *F*. Examples include: “rare,” “frequent,” “original,” etc. So, it should be uncontroversial that (i) is false under many substitutions for “*F*.” For all this last point shows, of course, a restricted version of the argument might still be sound. However, extra care would now be required to check that the cases witnessing (iii) do not involve one of the GAs for which (i) should uncontroversially fail. And it should always be borne in mind that the

conclusion (iv) could no longer be regarded as stating a general fact about GAs.

Having said all that about (i), let us now grant it for the sake of discussion. As far as we can tell, (ii) is assumed without much argument. Given the extreme looseness of the notion of a kind, it is hard to assess (ii) in abstraction from its use in motivating other claims, like (iii). It seems to us, however, that the notion of a kind in play should be one so relaxed as to be applicable, under appropriate circumstances, to just about any collection of things. Consider for example “big.” Take any (not all-inclusive) collection X of concrete objects. Take any concrete object x not in X . Consider then a rather impoverished context that has no antecedent information for the application of “big.” An utterance of “If we added x to X , we would have another big thing” can then be quite naturally interpreted as expressing the proposition that, if we added x to X , we would have another thing that is big for a *typical* thing in X – i.e. a proposition where the contextually selected comparison class is intensional in the “typical”-sense recommended by (i), so that the utterance would be false even if x were a normal ant and X only comprised all sorts of microscopical exemplars of cars, elephants, stars, etc. that happened to be much smaller than x .⁸

In this rather extreme case, the NP “thing in X ” does not directly provide any *conception of typicality*, which is gleaned instead from a conception of typicality for each of the things that happen to be in X .⁹ In less extreme cases, it is the NP itself that directly provides an at least inchoate conception of typicality. For example, Fara writes:

Comparison classes do not work just by contributing sets; for one, they need to form a kind. That is why it sounds strange to say that my computer is tall for a thing on my desk, even though it is in fact the tallest thing on my desk. Because the things on my desk don’t form a kind, we have no notion of what a *typical* height is for a thing on my desk. (Fara 2000: 56)

But, despite the fact that adjoining the PP “for a thing on one’s desk” to a GA results in a mild oddity, things on one’s desk can very well serve as comparison class, even in the intensional, “typical”-sense recommended by (i). For example, if you give a massive sequoia to a colleague as a gift and say that it is to be placed on her desk, she would certainly be within her rights in replying “But that is extremely tall,” thereby plainly expressing the proposition that it is extremely tall for a *typical* thing on her desk – i.e. a proposition where the contextually selected comparison

class is intensional in the “typical”-sense recommended by (i), so that her reply would be true even if all the things on her desk happened to be skyscrapers.

It thus seems that the notion of a kind should be appropriately relaxed, but doing this makes (iii) extremely unlikely to be true. Fara gives the example of a greyish-bluish book sometimes truly called “grey,” sometimes truly called “blue,” depending on which books surround it (whitish-bluish in the former situation, reddish-greyish in the latter). She thereby assumes that books in the surroundings do not form a kind. But again, that does not seem right. To repeat the style of example used above for “thing in X ,” suppose that the books in the surroundings are books published by presses $p_0, p_1, p_2 \dots p_n$, all famous for typically producing huge books. Suppose that we are considering whether to buy a certain standard Palgrave Macmillan paperback to place in the surroundings. We would certainly be within our rights in uttering “That would be very small,” thereby plainly expressing the proposition that it would be very small for a *typical* book in the surroundings – i.e. a proposition where the contextually selected comparison class is intensional in the “typical”-sense recommended by (i), so that our utterance would be true even if all the books in the surroundings happen to be miniature books exceptionally published by $p_0, p_1, p_2 \dots p_n$. This independently motivated relaxed notion of a kind allows the comparison classer to maintain that the relevant comparison class in Fara’s example just is the class of books in the surroundings. For books in the surroundings now do form a kind. Moreover, insofar as one is interested in an understanding of the above true color predications where the comparison class is taken in the intensional, “typical”-sense recommended by (i), one may assume that each of the books in the surroundings has its typical color. If so, the different prevailing colors in the two situations among the books in the surroundings can determine different standards for the application of “grey” and “blue.”^{10,11}

We conclude that we have not seen much compelling evidence against the comparison classer. Since, as we will see in section 12.5, Fara’s version of the IR-theory – along with many other contextualist theories – arguably requires that the context-dependence of GAs is not exhausted by variations in comparison classes, this raises an important issue about the independent plausibility of these approaches.¹²

12.3 Absolute cases and salient similarity

Let us move on to the theory of vagueness proper. Before examining which determinants of standards for the application of GAs Fara proposes

to add to comparison classes, we focus in this and the next section on some crucial *constraints* that she thinks such standards must respect. For a typical GA (e.g. “tall”), these are (modulo relabelling):¹³

ABSOLUTE CASES. Certain objects must be in the extension (anti-extension) in every context (e.g. “Yao Ming is tall”);

INTERNAL STRUCTURE. Extensions and anti-extensions must respect certain internal structural features in every context (e.g. “Everyone at least as tall as someone tall is tall”);

EXTERNAL STRUCTURE. Extensions and anti-extensions must respect certain other extensions and anti-extensions in every context (e.g. “No one is both tall and short”);

SALIENT SIMILARITY. Everything which in a context c is saliently similar to something in the extension (anti-extension) in c must also be in the extension (anti-extension) in c .

As we will see, it is SALIENT SIMILARITY that ends up doing much of the work in Fara’s version of the IR-theory. Yet, assuming classical logic, SALIENT SIMILARITY is in great tension with ABSOLUTE CASES (to preserve equity, we will see in section 12.5 that a crucial element of Fara’s own justification of SALIENT SIMILARITY is also in great tension with EXTERNAL STRUCTURE and in section 12.6 that Fara’s own understanding of SALIENT SIMILARITY is also in great tension with INTERNAL STRUCTURE).^{14,15}

Take a suitably fine-grained but discrete soritical series for “tall.” By classical logic, ABSOLUTE CASES entails there is a first case x_n to which “tall” is not truly applied in any context. SALIENT SIMILARITY then recklessly entails that there is *no* context in which [“tall” is truly applied to the preceding case x_{n-1} and the similarity¹⁶ between x_{n-1} and x_n is salient].¹⁷ But what is supposed to prevent the existence of such a context? Following Fara (2000: 67) we assume throughout an intuitive understanding of salient similarity according to which *active consideration* of the similarity between two objects *suffices* for their similarity to be salient.¹⁸ Now, by construction of n , there is a context – the one with the loosest possible standard for “tall” – in which “tall” is truly applied to the preceding case x_{n-1} . SALIENT SIMILARITY then entails that if, starting from that context, we do not do anything new but simply begin to consider the similarity between x_{n-1} and x_n , that will suffice to shift us to a new context in which *neither* x_{n-1} *nor* x_n are in the extension of “tall” (since it is not the case that they both can be and that their similarity is now salient). Although we discuss in section 12.4 cases where Fara accepts similarly induced context shifts, the shift that is claimed to take place here is particularly

problematic, for it seems to us that, on many plausible *metasemantic* views, there are sufficient conditions [for an object being in the extension of a GA in a context] whose holding is perfectly compatible with the relevant similarity being considered in that context. For example, on some views it might be sufficient for x_{n-1} to be in the extension of “tall” in a context c if the relevant subjects exhibit a complex pattern of dispositions to apply “tall” to x_{n-1} in c (see e.g. Williamson 1994: 231). Or, on some views it might be sufficient for x_{n-1} to be in the extension of “tall” in c if an assertion of “ x_{n-1} is tall” goes unchallenged in c (see e.g. Lewis 1979: 347). But we do not see why any suchlike sufficient condition cannot hold for x_{n-1} in c while x_{n-1} ’s similarity to x_n is being considered in c . And, if it can, SALIENT SIMILARITY is false. Notice that, arguably, this particular problem arises specifically for x_{n-1} , since, for every $i < n - 1$, the move is arguably available to say that, whenever any suchlike sufficient condition holds for x_i in c while x_i ’s similarity to x_{i+1} is being considered in c , *both* x_i and x_{i+1} are in the extension of “tall” in c . That move is however foreclosed in the case of x_{n-1} , since it would require that there is a context in which a_n is in the extension of “tall,” contrary to ABSOLUTE CASES.

Even if we bought into an odd metasemantic picture that would in effect allow SALIENT SIMILARITY to trump any positive determination whatsoever made by all the other determinants of standards, related problems would remain. Consider for example the series of natural numbers from 1 to 5 and a context c such that:

- (I) In c , 1 is a positive case for “small” whereas 5 is a negative case;
- (II) In c , 1 is saliently similar to 2;
- (III) In c , 2 is saliently similar to 3;
- (IV) In c , 3 is saliently similar to 4;
- (V) In c , 4 is saliently similar to 5.

We would have thought that there is such a c . But (I)–(V) are inconsistent if SALIENT SIMILARITY holds. This, we think, is a bad enough consequence of SALIENT SIMILARITY. For we assume that (I), (II), (V) and the result of stripping the salience gloss off (III) and (IV) uncontroversially hold, and we do not see how c could not also be such that one is considering both the similarity between 2 and 3 and that between 3 and 4. Think for example of a discussion about what counts as a small number of years on a postdoctorate (Weatherson 2010: 80–1): it seems to us obvious that there could be many discussions of this kind where 1 and 2 count while 5 and 4 do not, and where both the similarity between 2 and 3 and

that between 3 and 4 are considered. The latter is after all precisely one reason why it is so hard to tell whether 3 counts! But, even if the proponent of SALIENT SIMILARITY were to stick to her guns and deny that both (III) and (IV) hold, as long as she still granted that either can hold she would be left with the embarrassment of having by her own lights an infallible method of coming to know whether “small” applies or not to 3: for she knows that simply raising to salience the similarity of 2 with 3 will make “small” apply to 3, while simply raising to salience the similarity of 3 with 4 will make “small” not apply to 3. SALIENT SIMILARITY would only be preserved at the expense of IGNORANCE OF CUTOFFS.

The last point gestures at a more general problem with SALIENT SIMILARITY. If the constraint held, it would be utterly misguided of one to entertain the hypothesis that the cutoff is (one of the borderline cases) where one is looking (very much like it is utterly misguided of one to entertain the hypothesis that the current day is the previous day). And if one knew about the constraint, one could conclusively rule out that the cutoff is (one of the borderline cases) where one is looking (very much like one can conclusively rule out that the current day is the previous day). *Given our assumption that there is in effect a cutoff* (see note 14), both conditionals strike us as having a clearly false consequent: if there is indeed a cutoff, clearly it might well be any of the borderline cases, and the fact that one is looking at some of them would seem to have absolutely no bearing on the possibility that the cutoff is exactly one of those (where the possibility in question can be either broadly conceptual – as for the first consequent – or broadly epistemic – as for the second consequent). That is also evinced by the manifest oddity of thinking or saying “Although there is a cutoff and I don’t know where it is, since I am looking at x_n and x_{n+1} it cannot be x_n .” This point is easily missed by focusing exclusively on SORITES SUSCEPTIBILITY, since, as we will see in section 12.4, SALIENT SIMILARITY is alleged to imply [that, typically, one believes that there is no cutoff] and those who so believe will in effect be likely to think that the two consequents in question are true rather than false. But we are not claiming that it is that alleged implication of SALIENT SIMILARITY to be implausible. The point we are making is rather that, setting aside for the time being whether SALIENT SIMILARITY has implications that can satisfactorily explain why those who believe that there is no cutoff do so, that constraint (and knowledge of it) does have negative implications about the possibility of certain borderline cases being the cutoff – implications which, once it is assumed that there is in effect a cutoff, seem actually false.

12.4 Explaining the phenomena of vagueness?

If it is so problematic, why ever think that SALIENT SIMILARITY is true? Well, at least on Fara's version of the IR-theory, it is supposed to play a crucial role in the explanation of both SORITES SUSCEPTIBILITY and IGNORANCE OF CUTOFFS. By bivalence, SALIENT SIMILARITY is claimed to entail:¹⁹

TRUE INSTANCE WHEN CONSIDERED. Every instance of a soritical principle is true when one considers it²⁰

(See Fara 2000: 59.) We say that TRUE INSTANCE WHEN CONSIDERED is only claimed to be entailed by SALIENT SIMILARITY partly because there clearly are lots of instances of soritical principles that one can consider *without* the similarity between the two objects becoming salient to one. Appealing to some far-fetched scenario, this could be argued even for instances of (S_0). But the point can plainly be made for instances of soritical principles that do not bear on their sleeves the similarity of the two objects. Suppose for example that Andy, Bill, Charlie ..., Mark, Nick ..., Xandra, Yetta and Zac form a decreasing soritical series for "tall" but a subject, Hero, is not aware of the similarities of the series' neighboring objects. Then Hero might consider "If x is tall and y immediately follows x in the series, y is tall" under an assignment of Mark to " x " and Nick to " y ," but, given the absence of any consideration on her part of the similarity between Mark and Nick, it would seem that that similarity might nevertheless fail to be salient to Hero. Anticipating a bit, this point need not affect the explanation of SORITES SUSCEPTIBILITY, for one might still fully explain that phenomenon by using TRUE INSTANCE WHEN CONSIDERED suitably restricted as not to encompass subjects like Hero (after all, in her state of ignorance, Hero need not have any particular inclination to accept that conditional under that assignment). The point does, however, affect the explanation of IGNORANCE OF CUTOFFS: it seems to us out of question that, assuming the cutoff to be Mark, Hero is in no better position than us to find that out.²¹

Conversely, another reason to doubt the claimed entailment is that it manifestly assumes "if" as it occurs in a soritical principle to be some kind of *material implication* (in the sense that identity of truth values of antecedent and consequent suffices for its truth; Fara (2000: 79, note 19) seems to be aware of this glitch, but does not elaborate on it). Again anticipating a bit, this point need not affect the explanation of IGNORANCE OF CUTOFFS, for one might still fully explain that phenomenon by using TRUE INSTANCE WHEN CONSIDERED suitably restricted

to the material-implication reading of the relevant soritical principle. After all, a cutoff is a point at which the material-implication reading of the principle fails, so that, presumably, knowing of something that it is a cutoff point requires knowing that it is a point at which the material-implication reading of the principle fails. The point does, however, affect the explanation of SORITES SUSCEPTIBILITY: in most cases a soritical principle will be just as plausible under a reading of the conditional stronger than material implication.²²

Setting these worries aside, once TRUE INSTANCE WHEN CONSIDERED is in place SORITES SUSCEPTIBILITY is supposed to be explained as follows (see Fara 2000: 59). If one considered at a time t an instance of (S_0), then, by TRUE INSTANCE WHEN CONSIDERED, that instance would be true at t . *Assuming* the conditional fact about a sentence that, if one considered it at t , it would be true at t suffices for unconditionally inclining one to accept that sentence at t , one is inclined to accept any instance of (S_0) at t . *Assuming* that being inclined at a time to accept any instance of a universally quantified sentence implies being inclined at that time to accept the universally quantified sentence itself, one is inclined to accept (S_0) at t .

We would first like to stress the vast implausibility of the two assumptions emphasized in the previous explanation. As for the first assumption, there are many cases where the sheer conditional fact about a sentence that, if one considered it at t it would be true at t , plainly does not promote any unconditional inclination to accept that sentence at t nor, for that matter, any conditional inclination to accept that sentence at t provided that one considers it at t . For example, supposing that Goldbach's Conjecture is true, it would indeed be the case that, if one considered it at t , it would be true at t , but that sheer fact plainly does not promote any unconditional inclination to accept Goldbach's Conjecture at t nor, for that matter, any conditional inclination to accept Goldbach's Conjecture at t provided that one considers it at t . Since presumably the unconditional inclination to accept is supposed to be secured by the conditional inclination to accept, one should at least make sure not only that every instance of (S_0) is *true* when one considers it (which is given by TRUE INSTANCE WHEN CONSIDERED), but also that one is *inclined to accept* it when one considers it. We do not find in Fara's or other contextualists' work any argument to this effect. We suspect that some *transparency* principle is implicitly being assumed, according to which we are automatically inclined to accept certain kinds of truths concerning tallness (contrary, say, to truths concerning even numbers and sums of prime numbers). We will not investigate further here to what extent some such principle might be upheld.

We prefer to emphasize that, even if some correct transparency principle allowed us to move from TRUE INSTANCE WHEN CONSIDERED to a *conditional* inclination to accept, the further move to an *unconditional* inclination to accept would still be entirely unwarranted. To see this, consider for example:

(INST) For every x , if x exists, then an instance [of a universally quantified sentence] referring to x is considered by someone,

with “if” expressing material implication. Clearly, a reflective subject will have the conditional inclination to accept any instance of (INST) when she considers it (after all, reflecting on the fact that she considers it, she will easily see that its consequent is true). But, just as clearly, there may still be lots of instances of (INST) that a reflective subject will have no unconditional inclination to accept (say, because she is not aware of the existence of the relevant object or of the relevant instance). Notice that the problem cannot be fixed by saying that the move from a conditional inclination to accept to an unconditional inclination to accept becomes warranted under the further assumption that the subject is considering the relevant instance. For, since the move has to be warranted for each instance, this would require that the subject is considering each instance, a consequence which, together with SALIENT SIMILARITY and ABSOLUTE CASES, generates contradiction (see note 26).

As for the second assumption, being inclined to accept every instance of a universally quantified sentence certainly does not imply being inclined to accept the universally quantified sentence itself, for the simple reason that one might disbelieve or doubt that the instances one is inclined to accept are *all* the instances there are. It is arguable that this possibility is realized even in some cases where SORITES SUSCEPTIBILITY holds: for example, in cases where the soritical principle is known to have uncountably many instances. However, let us assume that that possibility is not realized. The main worry here is that this still falls very much short of implying that one is inclined to accept the universally quantified sentence itself. This is because of psychological versions of notorious *aggregation-failure* phenomena:²³ one might be inclined to accept of each participant of a fair lottery that she will lose, but one is typically not inclined to accept that everyone will lose (a psychological version of the “*lottery paradox*”, see Kyburg 1961: 197). Or, to give another example, one is certainly inclined to accept of each proposition about philosophy that one believes to be true that that proposition is true, but one might not be inclined to accept that every proposition about philosophy that one

believes to be true is true – i.e. that one always gets it right in philosophy (a psychological version of the “*preface paradox*”, see Makinson 1965).²⁴

Moreover, even assuming that the previous explanation works as far as it goes, we would like to mention some cases which would seem to exhibit exactly the same kind of psychological phenomenon as that highlighted in SORITES SUSCEPTIBILITY but where SALIENT SIMILARITY and TRUE INSTANCE WHEN CONSIDERED cannot kick in as desired. Consider for one the modification of “tall” as “tall in context *c*,” with *c* being a particular ordinary context. That phrase may mostly be used in philosophical English, but for all that it is a perfectly acceptable AP for competent speakers of that dialect. It is gradable. And one is inclined to accept the soritical principle that, for every *x* and *y*, if *x* is tall in context *c* and *y* is 1 inch shorter than *x*, then *y* is tall in context *c*. However, setting aside irrelevant issues of existence, there is little plausibility to the idea that there is *any* context-dependence in that phrase, given the usual fine-grained understanding of contexts as completely specific with regard to their agent, time, world and standards. But SALIENT SIMILARITY and TRUE INSTANCE WHEN CONSIDERED can only hold for context-dependent expressions, on pain of reinstating the sorites paradox in the metalanguage: for example, assuming to be a context-independent predicate, in a situation where, for every instance of a soritical principle for , someone considers it (possibly with different subjects considering different instances at different times), SALIENT SIMILARITY would in effect act as a soritical principle for the metalinguistic predicate “is in the extension of .”²⁵ (Indeed, as we explain in sections 12.5 and 12.6, on Fara’s version of the IR-theory the relevant kind of context-dependence that is supposed to be at play in SORITES SUSCEPTIBILITY and IGNORANCE OF CUTOFFS only requires dependence on *times*. Hence, we could just as well have considered the uncontroversially ordinary AP “tall at time *t*,” with *t* being a particular time.)

There are also cases of *quantifying* thought. Start with the soritical principle:

- (S₁) For every *x*, if *x* is small, then, for every *y* smaller than the second smallest natural number larger than *x*, *y* is small,

with “*x*” and “*y*” ranging over positive real numbers. Contrary to the consequent of an instance of (S₀), the consequent of an instance of (S₁) does not *refer* to any particular object – instead, it universally *quantifies* over an (uncountable) domain. Since that consequent quantifies rather than refers, it is not clear how SALIENT SIMILARITY is supposed to help with it – and in particular how SALIENT SIMILARITY could be used to derive TRUE

INSTANCE WHEN CONSIDERED – for it is not clear that, when one quantifies over a set X , one needs to consider any member of X (think for example of an ordinary citizen quantifying over the set of spies). And even granting that there is in some case no obstacle in going from quantification to consideration of the objects one is quantifying over, the specifics of (S_1) present the *additional* difficulty that, in order for SALIENT SIMILARITY to be at least in the ballpark for entailing TRUE INSTANCE WHEN CONSIDERED, the similarity of x with each and every real number smaller than the second smallest natural number larger than x should be considered. Setting aside the fact that it is highly controversial that normal human beings have so much as the *capacity to entertain thoughts about* many of the real numbers, it just strains credulity that *uncountably many* cases of similarity could be considered at the same time.²⁶ So we do not see how SALIENT SIMILARITY could be used to derive TRUE INSTANCE WHEN CONSIDERED for (S_1) , given that one cannot plausibly appeal to the relevant cases of similarity being all considered at the same time. Maybe there are ways of being *salient* that do not require being *considered*, and maybe for some of these ways what is important is that, roughly, the *span* within which the cases of similarity are located – rather than their *cardinality* – be small enough.²⁷ All this will depend on the relevant theory of salience, which as far as we know still needs to be developed in detail by contextualists appealing to the notion. But even assuming all that to be the case, one would still not be out of the woods: for even if the final theory of salience allows for the possibility of the relevant cases of similarity being all salient at the same time, SALIENT SIMILARITY could be used to derive TRUE INSTANCE WHEN CONSIDERED for (S_1) only if *consideration* of an instance of (S_1) suffices to realize that possibility. Given that the salience in question is now admittedly generated by something *other* than consideration, we find the required connection speculative at best.

Consider next the soritical principle:

- (S₂) For some F , for every x and y , if x is F and y is suitably similar to x , then y is F .

We are as much inclined to accept (S_2) as we are to accept (S_0) , but it is not clear that our inclination to accept (S_2) is wholly parasitic on our inclination to accept (S_0) (or some analogous principle in which “tall” is replaced by another GA).²⁸ Rather, it seems to us that one could be inclined to accept (S_2) without being aware of any witness to its truth, just as a *general* claim that properties – ways things can be – are not always sharply bounded. But such an inclination to accept could not be explained by SALIENT SIMILARITY and TRUE INSTANCE WHEN CONSIDERED.

We now turn to the explanation of IGNORANCE OF CUTOFFS. Once TRUE INSTANCE WHEN CONSIDERED is in place, IGNORANCE OF CUTOFFS is supposed to be explained as follows (see Fara 2000: 59). Suppose for *reductio* that one knows at a time t that [x is tall and y is not tall].²⁹ Then, by factivity of knowledge, at t [x is tall and y is not tall]. Moreover, at t one in effect knows the negation of an instance of (S_0). Assuming that knowing the negation of an instance at a time requires considering the instance itself at that time, one is considering the relevant instance of (S_0) at t . But then, by TRUE INSTANCE WHEN CONSIDERED, that instance is true at t . Since the truth of the instance at t implies that at t it is not the case that [x is tall and y is not tall], we get a contradiction.

Again, we would first like to stress the vast implausibility of the assumption emphasized in the previous explanation. On no intuitive understanding of “considering” does *knowing* at a time require *considering* (any constituent) at that time. One can know at 2.12 a.m. on 04/19/1770 [that, [if James Cook is a man, he is not a wombat] and $2+2 = 4$] without considering that, nor any of its constituents, at 2.12 a.m. on 04/19/1770.

Moreover, points analogous to those made above at the end of our discussion of the explanation of SORITES SUSCEPTIBILITY can be made here (save possibly for (S_2)): there are cases which would seem to exhibit exactly the same kind of epistemic phenomenon as that highlighted in IGNORANCE OF CUTOFFS but where SALIENT SIMILARITY and TRUE INSTANCE WHEN CONSIDERED cannot kick in as desired. We leave the details of this to the reader.

Finally, although this is not at the center of our interests in this chapter, it should be mentioned that almost every theorist of vagueness maintains a claim much stronger than IGNORANCE OF CUTOFFS, namely that *borderline cases in general* cannot be known by one.³⁰ It is likely that, for anyone adhering to that tenet, the *fundamental* explanation of IGNORANCE OF CUTOFFS will be the more general explanation of the unknowability of borderline cases. For such theorists then, an explanation of IGNORANCE OF CUTOFFS that could not be naturally extended to a more general explanation of the unknowability of borderline cases would crucially be prevented from limning the deepest features of the epistemic phenomenon surfacing in IGNORANCE OF CUTOFFS. Needless to say, any SALIENT SIMILARITY-based explanation is, alas, one such explanation.

12.5 Salient similarity and interests

Let us take stock. Section 12.3 contained our criticisms not only of Fara’s version of a SALIENT SIMILARITY-based IR-theory, but more generally of

any theory relying on SALIENT SIMILARITY. The previous section ascended to an even higher level of generality with our criticisms not only of Fara's version of a TRUE INSTANCE WHEN CONSIDERED-based IR-theory, but more generally of any theory that tries to explain SORITES SUSCEPTIBILITY and IGNORANCE OF CUTOFFS with TRUE INSTANCE WHEN CONSIDERED. This section descends to a much lower level of generality, investigating Fara's specific version of a SALIENT SIMILARITY-based IR-theory, and in particular the distinctive justification she gives for SALIENT SIMILARITY on the basis of a certain semantic hypothesis about GAs and of considerations pertaining to our interests. It is precisely at this level that we will see how interests are supposed to come into play in the generation of vagueness.

Fara thinks that SALIENT SIMILARITY is not a brute semantic constraint on GAs, but that it follows from their pure semantics together with some alleged facts about our interests. So let us first look at what Fara (2000: 71–5) understands their pure semantics to be. Taking “tall” as an example, something like the following truth condition is proposed:

(TALL) “ x is tall” is true (at circumstances of evaluation with world w and time t) in a context with interest bearers ii , norm n and comparison class c iff (at w and t) x has significantly (for the ii 's interests) more height than n as applied to c .

So, for example, if the interest bearers are Jodie and Caitlin, the norm a norm of typicality and the comparison class the class of basketball players, “ x is tall” is true (at w and t) iff (at w and t) x has significantly (for Jodie and Caitlin's interests) more height than what is typical for a basketball player. Note that Fara understands the relativization of truth to all the three contextual parameters in the standard-contextualist fashion: different utterances differing in these parameters will express different contents.

Let us record a couple of worries concerning (TALL). First, how can that style of analysis be *generalized* to other GAs? How can, say, “normal” be analysed in terms of the “significantly more than” construction? “ x is normal” is true iff x has a significant amount of what? Normality? And significantly more normality than what? Than what is the norm for strange things? It strains credibility to think that we are making implicit reference to such relatively exotic comparison classes in making common-or-garden judgments of normality. Second, for our purposes, Fara's analysis differs from Kennedy's (1999) by the addition of the “significantly (for the ii 's interests)”-clause. While, as we will see shortly, that is certainly required by Fara's explanatory strategy, it unfortunately

threatens to wreck havoc with the *exhaustivity* relations in which many GAs stand (such relations are in effect cases of EXTERNAL STRUCTURE and were first emphasized by Fine 1975: 270). So, for example, “red” and “orange” are arguably exhaustive over the range of the color spectrum that goes from red to orange: if something in that range is not red, it is orange (classically equivalently, everything in that range is either red or orange). However, given this, and leaving implicit interest bearers and comparison class, analogs of (TALL) would imply that, if something in that range does not have significantly more redness than the norm, then it has significantly more orangeness than the norm. If one considers that something that does not have significantly more redness than the norm might still have (not significantly) more redness than the norm, and that redness and orangeness are incompatible properties, that should strike one as a problematic implication.³¹

Having put that on the record, let us now look at how Fara (2000: 67–71) proceeds in deriving SALIENT SIMILARITY from (TALL) and its analogs together with some alleged facts about our interests. Here is her argument in a nutshell (taking “tall” as example and leaving implicit interest bearers and comparison class):

- (a) If two things are saliently similar (with respect to height), then they are the same for present purposes (with respect to height), i.e. for present purposes, it is fine to ignore the difference between them (with respect to height);
- (b) If it is fine to ignore the difference between two things (with respect to height), then the costs of discriminating between them (with respect to height) outweigh the benefits;
- (c) If the costs of discriminating between two things (with respect to height) outweigh the benefits, then one is significantly taller than the norm iff the other is;
- (d) Therefore, if two things are saliently similar (with respect to height), then one is in the extension of “tall” iff the other is.

At the end of this section we address the question of the validity of the argument from (a)–(c) and (TALL) to (d). Having already discussed (TALL), let us now focus on the premises (a)–(c). Frankly, we see little justification to believe any of (a)–(c) and indeed ample justification to disbelieve them.

Take (a). It may be that we want to organize two school basketball teams for different height leagues *A* and *B*. There are 10 schoolgirls, four of whom are equally very tall and four of which are equally very short.

The remaining two girls, Jodie and Caitlin, are of middle height, with one ever so slightly taller than the other. None of the very tall girls meets the shortness criteria of the *B* league and none of the very short girls meets the tallness criteria of the *A* league; Jodie and Caitlin meet both, and no player in the *A* league can be shorter than any player in the *B* league. Given our interest in organizing the two teams, clearly it is not the case that it is fine to ignore the difference in height between Jodie and Caitlin: to do so would forego the only way in which we can satisfy our desire to have the two teams. And all this holds, of course, even if the similarity (with respect to height) of Jodie and Caitlin is considered, and hence salient.³² Such counterexamples are legion. The underlying general point is that (a) gives to the thin *contemplative* fact of salient similarity the implausible power in the *practical* domain of trumping any deeper interest we may antecedently have in not ignoring the difference between two things.

As for (b), we simply remark that it might be fine to ignore the difference between two things because the *benefits* of discriminating between them *do not outweigh* the *costs* without it being the case that the *costs* of discriminating between them *outweigh* the *benefits*. Notice that, while one could simply weaken (b) accordingly, the consequent strengthening of (c) would be even less plausible than (c) itself.

As for (c), it is hard to see why it should be true in the absence of more of an explication of what “significantly” is supposed to mean. We can at least get a connection between its antecedent and consequent by interpreting the former as implying that the difference in height between two things is not significant. But it would be a fallacy in the *logic of significance* to think that that in turn implies the consequent of (c). In general, that the difference between *x* and *y* is not significant and that the difference between *x* and *z* is significant does not entail that the difference between *y* and *z* is significant.³³

Moreover, it is interesting to ask oneself why, if the (a)–(d) argument is any good, an analogous argument should not go through when *mere* similarity is substituted for *salient* similarity in (a), thus yielding:

- (a') If two things are similar (with respect to height), then they are the same for present purposes (with respect to height), i.e. for present purposes, it is fine to ignore the difference between them (with respect to height).

Such an analogous argument would of course be unacceptable, as it would reinstate the sorites paradox. Predictably enough, Fara seems to

want to block this new argument at (a'): the fact that two things are similar is supposed not to imply that, for present purposes, it is fine to ignore the difference between them. However, we do not see why mere salience should make a difference as to whether, for present purposes, it is fine to ignore the difference between them: as we have already mentioned in our discussion of (a), salience is too thin a contemplative fact to have such wide-ranging practical effects. In support of her line of thought here, Fara (2000: 68) offers the further claim that, if two things are not "live options," then there is no cost in discriminating between them. Again, we do not see why the thin contemplative fact of salience should have the substantial practical effect of making two things "live options": at least in the ordinary sense of "live option," not every option that happens to be considered is a live option (and not every live option happens to be considered). We also do not see why not being a "live option" should imply that there are no costs in making the relevant discrimination – if anything, we would have thought that it implies that there are rather no benefits (and some costs) in making the relevant discrimination! (Together with premise (c), that would of course reinstate the sorites paradox.)

More generally, it is unclear to us that it is not the case that, at least in a large variety of cases, [for present purposes, it is fine to ignore the difference between any two similar things and, for any two such things, the costs of discriminating between them outweigh the benefits]. Both these claims actually strike us as probably correct in a large variety of cases (although not in all cases, as we demonstrated in our discussion of (a)). And, given the arguable *nontransitivity* of the relations they employ, neither claim reinstates by itself the sorites paradox. Both do, however, in conjunction with (b), (c) and (TALL). Being wedded to the latter three, Fara (2000: 70–1) tries to parry this attack by observing that the costs of discriminating somewhere do not outweigh the benefits of doing so. We wholeheartedly agree and would go even further by saying that the benefits of discriminating somewhere outweigh the costs of doing so, but such a platitude does little to ward off the attack. For one need not make the discrimination between similar objects – one may rather do it only between nonsimilar objects (again, the arguable nontransitivity of nondiscrimination allows that, in general, one may perform what could be called a "rough" discrimination – i.e. discriminate between x and y without discriminating between any objects strictly between x and y).³⁴

Having said all this about the premises (a)–(c), we can briefly turn to the question of the validity of the argument from (a)–(c) and (TALL) to (d). So far, we have induced the reader not to think too hard about interest

bearers by saying that these have been left implicit in the (a)–(d) argument. However, the step from (a)–(c) and (TALL) to (d) is only valid if the interest bearers that are considering x and y and are (implicitly) referred to throughout (a)–(c) (let us call them “the *considering* interest bearers”) are the *same* as those that are semantically relevant and are referred to in the operative instance of (TALL) (let us call them “the *semantic* interest bearers”). Given how Fara’s explanation of the phenomena of vagueness is supposed to work (see section 12.4), obviously the group of considering interest bearers has at least to *contain the speaker*. But it is not even clear that the group of semantic interest bearers also has to do so: indeed, the assumption that, in using a GA, the value of the interest-bearers parameter necessarily contains the speaker should appear to be very problematic, given that similar assumptions about “speaker inclusion” are false (the domain to which “everyone” is contextually restricted need not contain the speaker, the body of knowledge relevant to “might” need not include the speaker’s, the standard of taste relevant to “tasty” need not take into account the speaker’s).

However, let us grant for the sake of argument that that assumption is correct. It remains very plausible [that the group of semantic interest bearers may also include people other than speaker], and the validity of the argument from (a)–(c) and (TALL) to (d) requires that group to be the same as the group of considering interest bearers. If so, the question must arise as to why, if the similarity of x and y is salient *to the speaker*, the discrimination between them could not be made *by some other* of the semantic interest bearers to whom such similarity is not salient, in such a way that, after all, the costs of discriminating between them would not outweigh the benefits. That possibility would sever the link between the similarity of x and y being salient to the speaker and x and y falling on the same side of the cutoff for “tall” as uttered by the speaker. Since at the relevant stage of Fara’s explanation of the phenomena of vagueness no more can be assumed than the similarity of x and y being salient to the speaker (rather than its being salient to the speaker and to everyone else who is a semantic interest bearer for “tall” as uttered by the speaker), for that possibility her explanation would not get off the ground. As far as we can see, the only way to rule that possibility out would be to admit as group of semantic interest bearers at a time only groups of people containing the speaker and in which the objects salient to the speaker at that time are also salient to everyone else at that time, but that restriction would strike us as hopelessly ad hoc (think of how implausible an analogous restriction would be in the case of candidate referents for “we”).³⁵

12.6 The IR-theory vs standard-contextualist theories

We finally turn to a comparison of Fara's version of the IR-theory with contextualist theories which try to explain the phenomena of vagueness by variously appealing to context-dependence understood along *standard-contextualist* lines (see e.g. Kamp 1981, Raffman 1994, Soames 1999, Shapiro 2006). Although, as has emerged in the previous section, Fara's analysis also postulates several elements of standard-contextualist context-dependence in the semantics of GAs (interest bearers, norm, comparison class), the previous section should also have made clear that it is not these elements that are exploited in her explanation of SALIENT SIMILARITY (which, to recall, is in turn supposed to yield TRUE INSTANCE WHEN CONSIDERED which is in turn supposed to yield SORITES SUSCEPTIBILITY and IGNORANCE OF CUTOFFS). For what her explanation exploits is rather the fact that salient similarity to the *iis* between x and y implies that the *ii*'s interests are in a certain way, a way that ensures that x is in the extension of a GA iff y is (again, recall the (a)–(d) argument of the previous section). Thus, the relevant changes that are supposed to ensure SALIENT SIMILARITY are changes neither of the interest bearers, nor of the norm, nor of the comparison class, but *changes over world and time in what the interest bearers' interests are*. These changes are due to what is saliently similar to the interest bearers at a world and time and concern what pairs of objects are in their interest to discriminate at a world and time. And they affect the semantics *only via the semantics' more general world and time dependence* (just as, for example, changes in weight affect the semantics of "heavy" only because an object can have different weights at different worlds and times and the semantics of "heavy" has the more general property of being world and time dependent). Fara further assumes a *temporalist* (and hence nonindexical contextualist) semantics, so that the relevant changes do not induce any change in the content expressed by a GA; on an *eternalist* (and hence standard contextualist) semantics, the only change in content would be that of the time involved in the content. For simplicity, we too will presuppose a temporalist semantics, but emphasize that the choice is irrelevant for the substance of our discussion in this section (see note 36).

It is this relative *invariance in content* that helps Fara's version of the IR-theory to avoid some of the arguments against standard-contextualist theories from the content-invariant interpretation of context-dependent expressions under *verb-phrase ellipsis* (see Stanley 2003). Assuming

(as we will mostly do) the content-invariant interpretation of context-dependent expressions under verb-phrase ellipsis, Stanley argued that standard-contextualist theories cannot explain analogs of SORITES SUSCEPTIBILITY and IGNORANCE OF CUTOFFS for the following sorites:

(VPES) This man is tall; and this₁ is too; and this₂ is too ... ; and this_n is too,

where “this_{*i*}” ($1 \leq i \leq n$) refers to a man just slightly shorter than the previous one and “this man” and “this_{*n*}” refer to a clearly tall man and a clearly short man respectively. This is so because the content of the elided verb phrase is invariant throughout (VPES), and so no variation in content is available, contrary to what would be required by a standard-contextualist theory. By contrast, Fara’s version of the IR-theory (with (TALL) understood temporalistically) has no problem with (VPES), since, even though the *content* of the elided verb phrase cannot change throughout (VPES), our *interests* can, and with them the *extension* of the elided verb phrase.³⁶

This is fine as far as it goes, but we do not think it goes very far. As Fara’s version of the IR-theory stands at this point, it cannot explain analogs of SORITES SUSCEPTIBILITY and IGNORANCE OF CUTOFFS for the following *modal* sorites:

(MS₀) For some maximally specific property, this man is tall and interests have not that property; but if they had, the man would still be even if he were just slightly shorter than that₁; and if they had, the man would still be even if he were just slightly shorter than that₂ ... ; and if they had, the man would still be even if he were just slightly shorter than that_{*n*},

with the same notational conventions as before plus the conventions that, under every assignment, [“that_{*i*}” ($1 \leq i \leq n$) refers to the height exemplified by the man in the previous counterfactual circumstance and “that_{*n*}” refers to a clearly short height]. This is so because the ellipsis on the bound verb phrase “have not that property” ensures that, in all the various counterfactual circumstances entertained throughout (MS₀), what on Fara’s version of the IR-theory are the semantically relevant interests cannot change, while the ellipsis on the verb phrase “is tall” ensures that the values of the standard-contextualist parameters of “tall” cannot change throughout (MS₀). Together, these two invariances force

not only the *content* expressed by “tall” throughout (MS₀) to remain constant, but also its *extension* relative to the various counterfactual circumstances entertained throughout (MS₀), contrary to what would be required by Fara’s version of the IR-theory.

Indeed, the trick of (MS₀) is really that of forcing, in the vanilla way afforded by counterfactuals, the evaluation of “tall” to be made relative to a circumstance where the interests are not guaranteed to be the actual ones of the speaker (see the end of section 12.5 for less vanilla ways of achieving similar results). It is no surprise then that a similar point could be made without exploiting verb-phrase ellipsis to keep the interests constant, as in the following *modal* sorites:

- (MS₁) This man is tall; and he would still be even if he were just slightly shorter than that₁ and interests were just slightly different from the way they now actually are;³⁷ and he would still be even if it were just slightly shorter than that₂ and interests were just slightly different from those₁ ... ; and he would still be even if he were just slightly shorter than that_{*n*} and interests were just slightly different from those_{*n-1*},

with the same notational conventions as before plus the conventions that “those_{*i*}” ($1 \leq i \leq n - 1$) refers to the interests obtaining in the previous counterfactual circumstance and “those_{*n-1*}” refers to clearly possible interests. Although interests are now allowed to change from one counterfactual circumstance to another, their change is predetermined and cannot be further influenced by salient-similarity facts, contrary to what would be required by Fara’s version of the IR-theory (for good measure, we still use verb-phrase ellipsis to ensure that such facts cannot even affect the values of the standard-contextualist parameters of “tall”).³⁸

Thus, we do not think that Fara’s version of the IR-theory has a substantial advantage over standard-contextualist theories as far as matters of verb-phrase-ellipsis arguments are concerned.³⁹ Indeed, similar considerations of modal embeddings lead us to observe a *clear disadvantage* of Fara’s version of the IR-theory over standard-contextualist theories. For suppose that the boundary between the positive and the negative cases of “tall” (as uttered by speaker *s* at world *w* at time *t*) falls between *x* and *y*. Then, on Fara’s version of the IR-theory, the conditionals “If I considered (consider) the similarity in height between *x* and *y*, either *x* would (will) cease to be tall or *y* would (will) start to be tall” are true (as uttered by *s* at *w* at *t*; we will assume here the solipsistic view mentioned in note 35 and

a material-implication reading of indicative conditionals). We find those conditionals extremely unpalatable and deem them to embody a severe misconception of *what tallness depends on*: not only – counterintuitively enough – as depending on someone’s considering things, but also as depending on *s* rather than anyone else considering things! Standard-contextualist theories, even if SALIENT SIMILARITY-based, do remarkably better on this score, since, at least on the usual way of developing them, they allow all suchlike conditionals to be false. What they do require to be true (as uttered by *s* at *w* at *t*) are rather more theoretical conditionals such as “If I considered (consider) the similarity in height between *x* and *y*, either *x* would (will) not be in the extension of ‘tall’ as uttered by me then or *y* would (will) be in the extension of ‘tall’ as uttered by me then.” The argument extends to many other conditional and non-conditional sentences that are extremely unpalatable, but true on Fara’s version of the IR-theory, like “It is possible that something becomes tall without changing height.”

Indeed, even INTERNAL STRUCTURE is now in jeopardy. Suppose that the cutoff for being tall in a certain soritical series comprising Andy (who is 6 feet and 2 inch tall) and Bill (who is 6 feet and 1 inch tall) is Andy, and that one is considering another series where the similarity between Charlie (who is also 6 feet and 2 inch tall) and Dave (who is also 6 feet and 1 inch tall) is salient. Then, on Fara’s version of the IR-theory, one could truly utter “Either Charlie is not tall and is at least as tall as someone who is tall or Dave is tall and someone [who is at least as tall as he is] is not tall,” which is straightforwardly inconsistent with our (and Fara’s 2000: 73) paradigmatic example of INTERNAL STRUCTURE. Fara, just as well as a standard contextualist relying on SALIENT SIMILARITY, might rejoin that SALIENT SIMILARITY should be charitably understood as pertaining to any pair of objects that agrees in the relevant supervenience base with the pair whose similarity is salient. We do not see a lot of independent plausibility for such a move, especially if SALIENT SIMILARITY is supposed to be explained in something like the praxis-oriented way attempted by Fara, but let that pass. Suppose that the cutoff lay between Andy and Bill at time t_0 and the similarity between Charlie and Dave becomes salient at a later time t_1 . Then, on Fara’s version of the IR-theory but not on a standard-contextualist theory, one could truly utter at t_1 “Either Charlie is not tall and is at least as tall as someone [who (at t_0) was tall] was (at t_0) or [Dave is tall and someone [who (at t_0) was at least as tall as he is] was not tall (at t_0)],” which is straightforwardly inconsistent with the platitudes that a man who is at least as tall as [a man who was tall was] is tall (first disjunct) and that a man who was

at least as tall as [a man who is tall is] was tall (second disjunct), platitudes which we take also to belong to INTERNAL STRUCTURE.

It is sometimes suggested that the last series of counterintuitive results can be avoided by the IR-theory if the semantics “*rigidifies*” the speaker’s actual and present interests (see e.g. Stanley 2003: 278–9), as in:

(TALL’) “*x* is tall” is true (at a circumstance of evaluation with world w_0 and time t_0) in a context with interest bearers *ii*, norm *n*, comparison class *c*, world w_1 and time t_1 iff (at w and t) *x* has significantly (for the *ii*’s interests as they are at w_1 and t_1) more height than *n* as applied to *c*.

We do not think however that, in this dialectical context, rigidification would be a wise move on behalf of the IR-theory. Firstly, as far as the conception of what tallness depends on is concerned, even a rigidified version of the IR-theory would be stuck for example with the truth of “If I’m considering the similarity in height between *x* and *y*, *x* is tall iff *y* is, even though as a matter of fact *x* is tall and *y* is not” (as uttered by *s* at w at t). The initial problem with INTERNAL STRUCTURE would also remain. Secondly, as far as arguments from verb-phrase ellipsis are concerned, a rigidified version of the IR-theory would have the same problem with (VPES) as standard-contextualist theories have, since, although our interests can change throughout (VPES), the invariant content of the elided verb phrase would always refer back to the interests we had at the time of utterance of the first sentence of (VPES).

We conclude that, as far as the issues discussed in this section are concerned, neither Fara’s nor a rigidified version of the IR-theory offers substantial advantages over standard-contextualist theories and that, on the contrary, they suffer from some significant disadvantages in that they license claims that do violence to our ordinary conception of what tallness depends on.

12.7 Conclusion

Beyond issues of detail pertaining to the specifics of Fara’s version of the IR-theory, there are two main points that we hope to have impressed on the reader: the implausibility of the semantic and psychological assumptions that one needs to make in order to get a version of the IR-theory with some hope to engage with the phenomena of vagueness (such as e.g. a version equipped with SALIENT SIMILARITY) and the great gap that would actually still remain between such version and the phenomena

of vagueness. Although this is a topic for another essay, we believe that analogs of these two points hold for a wide variety of contextualist theories. If so, other, in certain respects more radical approaches to the phenomena of vagueness begin to look more appealing. One such approach disputes the status of sacrosanct data that we have accorded so far to *SORITES SUSCEPTIBILITY* and *IGNORANCE OF CUTOFFS*: it argues that it is exactly focus on the context-dependence of vague expressions that allows us to see how one can quite competently draw in context a perfectly sharp boundary between positive and negative cases (see Sweeney 2010). At the opposite end of the spectrum, another such approach takes those two phenomena much more seriously than most contextualist theories do: by revising classical logic, it declares true soritical principles such as (S_0) , from which it is then able to derive explanations of *SORITES SUSCEPTIBILITY* and *IGNORANCE OF CUTOFFS* that exhibit a shocking simplicity in comparison with the convolutions and epicycles with which contextualist theories are typically saddled (see Zardini 2008, 2010).

That languages spoken by human beings contain vague expressions at least partly because that allows them to pursue certain practical interests which could not otherwise be (easily) pursued is a nonobvious but appealing hypothesis to which we are very sympathetic. But it is not the IR-theory. That theory rather holds that what underlies certain phenomena of vagueness is the variability of those interests. That is what we have argued against in this chapter, making the two main points mentioned above. On the one hand, no sane man's practical interests are as volatile as they would have to be in order to generate the required variability; on the other hand, that great riddle surfacing in *SORITES SUSCEPTIBILITY* and *IGNORANCE OF CUTOFFS* reaches much deeper into our thoughts about reality than anything that any such variability could ground.

Notes

* Earlier versions of the material in this chapter were given in 2005 at the Arché Vagueness Seminar (University of St Andrews) and at the Arché Reading Party in Garve (University of St Andrews); in 2008, in a course on *Context-Dependence and Relativity: Foundations and Applications* (co-taught with Herman Cappelen) at the Arché Summer School (University of St Andrews); in 2009, at the 17th Meeting of the European Society for Philosophy and Psychology in Budapest (Central European University). We would like to thank all these audiences for very stimulating comments and discussions. Special thanks go to María Cerezo, Laura Delgado, Paul Égré, Delia Graff Fara, Patrick Greenough, Nathan Klinedinst, John MacFarlane, David Ripley and Crispin Wright. In writing this chapter, the first author (the order being merely alphabetical) has benefited from an AHRC Doctoral Award, the second author, at different stages, from an AHRC Doctoral Award

and an AHRC Postdoctoral Research Fellowship, as well as from partial funds from the project FFI2008-06153 of the Spanish *Ministry of Science and Innovation* on *Vagueness, Physics, Metaphysics and Metametaphysics* and from the project CONSOLIDER-INGENIO 2010 CSD2009-00056 of the Spanish *Ministry of Science and Innovation* on *Philosophy of Perspectival Thoughts and Facts* (PERSP). Both authors have also benefited from partial funds from the European Commission's 7th Framework Programme FP7/2007-2013 under grant FP7-238128 for the European Philosophy Network on *Perspectival Thoughts and Facts* (PETAF).

1. Focusing on three prominent semantic categories in philosophical logic, we assume (roughly) that the extension of a singular term is an object, the extension of a predicate a set, the extension of a sentence a truth value. Throughout, we also assume that the reader is familiar with the broad outlines of a standard Kaplan-style semantic framework (see Kaplan 1989).
2. A general caveat: we will try to be as faithful as possible to Fara's writings, but, despite our best efforts, we might still be misrepresenting her arguments and views. Even if that were the case, we still believe that the arguments and views we do discuss have enough independent interest.
3. Fara is of course aware that expressions belonging to other syntactic categories can also be vague (and indeed paradigmatically so, think of the noun "heap"). It is beyond the scope of this chapter to investigate whether and how Fara's version of the IR-theory can be extended to other syntactic categories.
4. Throughout, when appropriate we disregard world- and time-induced context-dependence.
5. Throughout, we do not assume that the ordinary phrase "the class of *F*s" (substituting an NP for "*F*") as used in theorizing about comparison classes denotes the *set* of *F*s (or anything else that is *extensionally* individuated in terms of what things happen to be *F*). The ensuing discussion in this section will make clear why this assumption would be moot.
6. This is relevant to (i) if we assume (as we will do) that the comparison class of *F*s is picked out by the PP "for an *F*."
7. The point was first emphasized by Enç (1986: 410–12). A staple example is a sentence like "Every fugitive sits in jail."
8. The reader is warned that, from now until the end of this section, some fairly far-fetched possibilities will be contemplated, just to bring the point home with the utmost clarity.
9. On this way of achieving a conception of typicality for the *F*s, where it wholly depends on a conception of typicality for each of the things that happen to be *F*, since it may be contingent and temporary what is *F*, it may be contingent and temporary what is typically *F*. We think that this consequence is independently plausible.
10. Throughout, by "standards" we will mean, roughly, the *thresholds* that contextually mark extensions and anti-extensions. For example, a standard for "tall" is a specific height threshold.
11. Notice that the examples of intensional comparison classes we have given in our discussion of (ii) and (iii) need not give any support to (i). For the denial of (i) is perfectly compatible with the acknowledgement that the comparison class of *typical F*s just is the set of typical *F*s, and that such comparison classes are those that are contextually selected in our examples.

12. We do not want to suggest, however, that things look particularly rosy for the comparison classer. See DeRose (2008) for a recent attack based on some interesting examples.
13. See Fara (2000: 57). The importance of the first three constraints for a theory of vagueness was first emphasized by Fine (1975: 269–71).
14. We assume classical logic throughout. The same for classical semantics (in particular, for the principle of bivalence). Hence, keeping fixed ABSOLUTE CASES, we will assume throughout that there are in effect cutoffs.
15. Thanks to David Ripley for probing questions that led to a much-improved version of this section.
16. Throughout, we will mostly leave implicit the relevant respect of similarity (in this as in many other cases, height).
17. Throughout, we will use square brackets to disambiguate constituent structure.
18. As we will see in section 12.4, that understanding is necessary for Fara's explanation of the phenomena of vagueness, but it might actually be problematic if the notion of salience in question is supposed to coincide with notions already in use in psychology and linguistics, where often a notion of salience is used which is such that active consideration by no means suffices to raise salience (and actually, if abstract and cold enough, can on the contrary lower it – see Sherman et al. 1985, Oppenheimer 2004). We will not fuss further over this interesting issue here. We also henceforth drop the rather obscure qualification “active,” as we do not think it could ultimately do a very useful work: as we see again in section 12.4, the understanding of salience necessary for Fara's explanation of the phenomena of vagueness should not require anything more than mere consideration.
19. Some prominent theories of vagueness, such as the *supervaluationist* theory advocated in Fine (1975), reject bivalence. Fara (2000: 62) observes that, on that kind of theory, the entailment from SALIENT SIMILARITY to TRUE INSTANCE WHEN CONSIDERED still holds if one understands SALIENT SIMILARITY as a constraint on contextually admissible, precisified, bivalent interpretations (rather than simply as a constraint on vague, nonbivalent interpretations). However, that understanding of SALIENT SIMILARITY is extremely problematic, as, in a context c where the similarity of at least two borderline cases is salient, it would count an extremely plausible principle like:

(SV) A borderline case for being tall is not definitely not the cutoff for being tall

false on every contextually admissible, precisified, bivalent interpretation, and hence false *simpliciter* (let x and y be saliently similar in c and both borderline, with x taller than y : then it is true on every c -admissible, precisified, bivalent interpretation that x is not the cutoff for being tall, and hence it is true *simpliciter* that x is definitely not the cutoff for being tall).

20. Henceforth, we use “instance” (of a quantified sentence φ) in a broad sense, so as to encompass (also) any assignment of values to the variables of φ 's matrix (see note 23 for a justification of this choice).
21. Puzzlingly enough, Fara might disagree. In Fara (2000: 58) she gives a case where she claims that, in effect, a subject in conditions similar to Hero's can

“truly proclaim” of a cutoff that it is a cutoff. Since she explicitly contrasts that case with a case where she claims that the subject is “unable to locate” the cutoff, one might infer that Fara thinks that in her first case the subject *is able to locate* the cutoff, which in turn is usually taken to imply that the subject *knows* where the cutoff is. If so, Fara can only uphold a watered-down version of IGNORANCE OF CUTOFFS (one which does not encompass subjects like Hero). We find that more robust versions of IGNORANCE OF CUTOFFS (ones which do encompass subjects like Hero) are fully warranted and find Fara’s claims about her first case very counterintuitive. Alleged examples where speakers successfully draw sharp boundaries in a soritical series actually go back to Sainsbury (1990: 259–60). This is not the place to delve further into such examples (but see section 12.7 for a parting thought on this).

22. We are assuming throughout a reading of the restricted quantification in SALIENT SIMILARITY that makes it equivalent with an unrestrictedly quantified-into material implication. One way to try to get around this problem would be by arguing that the considerations in favor of SALIENT SIMILARITY actually support a suitably stronger reading of that constraint. In section 12.5, we will discuss in detail Fara’s own justification of the constraint and flag the spot where, in this respect, much of the action is (see note 33).
23. If we had worked instead with the ordinary notion of an instance (roughly, the result of stripping off the initial quantifier and uniformly replacing the now free variable with a congruent interpreted expression), we would have also observed another kind of failure of that implication: in that sense of “instance,” one is certainly inclined to accept every instance of “Everything is the designation of a singular term in English,” but one is typically not inclined to accept the universally quantified sentence itself (see note 20).
24. Keefe (2007: 281) makes a similar point. Revising Fara (2000: 59), Fara (2008b: 15–16) replies in effect by saying that the correct psychological-explanatory principle that should be assumed in the explanation is that, roughly, if one is inclined to accept the universally quantified sentence itself, that can be explained by the fact that one is inclined to accept every instance of it (rather than the principle that, if one is inclined to accept every instance of a universally quantified sentence, one is inclined to accept the universally quantified sentence itself). Setting aside broader issues in the philosophy of explanation, the new principle is even more clearly wrong than the old one: we are inclined to accept that every set of positive real numbers has a greatest lower bound, but that cannot be explained by the alleged fact that we are inclined to accept of every set of positive real numbers that it has a greatest lower bound – on most reasonable understandings of “inclined to accept,” the alleged fact does not even hold. And even if, in some sense, that fact held, the direction of explanation would certainly go the other way round. Similarly, in many cases in which one is (immodestly) inclined to accept that every proposition about philosophy that one believes to be true is true, that is explained by facts other than the completely trivial fact that one is (not immodestly) inclined to accept of each proposition about philosophy that one believes to be true that that proposition is true. Fara (2008b: 7) also offers the thought that the old principle might still hold defeasibly and mentions among possible defeaters knowledge that the extension of a certain

predicate might have changed in the course of the reasoning and, more relevantly for our counterexamples from aggregation failure, knowledge that the conclusion has a very low probability on one's evidence. We wholeheartedly endorse the (rather uncontroversial) defeasibility of the explanatory link (and add that insofar as the new principle is supposed to be a gloss of that, it is a bad one). All this would however only be helpful in conjunction with an explanation – which is still lacking – of why the relevant defeaters are present in the counterexamples but not when one is assessing whether to move from (S_0) 's instances to (S_0) itself.

25. Heck (2003: 118–20) gives voice to a widespread concern when he considers constructions that give rise to analogs of SORITES SUSCEPTIBILITY and IGNORANCE OF CUTOFFS but which are not plausibly taken to be context-dependent. While we certainly share Heck's broader worries here, none of those constructions is however a gradable AP, and hence Heck's examples do not immediately tell against Fara's version of the IR-theory, which – for better or worse – is explicitly restricted to GAs (and, we are assuming, to gradable APs headed by GAs and adjoined by PPs). This explains our particular choice of example in the text.
26. Considering a finite soritical series, SALIENT SIMILARITY would be straightforwardly inconsistent with ABSOLUTE CASES if all the finitely many cases of similarity could be considered at the same time. With reference to this problem, Fara (2000: 70–1) mentions with approval the idea that all the *finitely* many cases of similarity are too many to be all salient at the same time (and hence too many to be all considered at the same time). But then certainly *uncountably* many cases of similarity should also count as being too many!
27. We owe the last suggestion to David Ripley.
28. And even if it were claimed that our inclination to accept (S_2) is wholly parasitic on our inclination to accept (S_0) , this strategy would first need some refinement, since clearly inclination to accept is not generally closed under logical consequence, not even a single-premise one.
29. Having raised some trouble at the beginning of this section for an explicit suitable-similarity clause, henceforth we will oftentimes leave it out, implicitly assuming that γ is de facto 1 inch shorter than x .
30. More precisely, that claim is stronger if one assumes (as we will do throughout) that the cutoff is a borderline case and that there is more than one borderline case.
31. In response to this problem, one might be tempted to propose for “negative” GAs like “short” a truth condition different in structure from (TALL), so that, roughly, “ x is short” is true (given interest bearers *ii*, norm *n* and comparison class *c*) iff x does not have significantly (for the *ii*'s interests) more height than *n* as applied to *c*. Setting aside the (pressing) questions as to whether the notion of a “negative” GA is well-defined (which we doubt) and whether the alternative truth condition is independently plausible (for “short” at least it does not seem to be so), such a maneuver would hardly get off the ground with “red” and “orange,” as there is no plausibility to the idea that one of them is a “positive” GA and the other one a “negative” GA and, more generally, it is very hard to see how to justify a break in their apparent symmetry by assigning to them structurally different truth conditions.
32. From now on we mostly omit the “with respect to height” qualification.

33. Notice that the reading of the bi-conditional occurring in (c) is crucial if one wants to argue for a stronger reading of SALIENT SIMILARITY than we have been assuming, which might help with a problem we mentioned at the beginning of section 12.4 (see in particular note 22). Defending (c) under a suitably strong reading of its bi-conditional would however only exacerbate the difficulties discussed in the text.
34. The nondiscrimination in question is *practical* rather than *epistemic*, but it is still extremely plausibly taken to be nontransitive. The point is easily missed by taking nondiscrimination between x and y to consist in x 's being treated in the *same* way as y and then fanatically assuming that the latter relation involves a transitive-identity relation. It clearly does not, as x 's being treated in the same way as y tolerates minute and insignificant differences in their treatment, just as x 's walking in the same way as y tolerates minute and insignificant differences in their gait. One could of course stipulate to be using a notion of nondiscrimination that, contrary to the ordinary one, requires perfect match, but then (b) would be blatantly false.
35. With respect to this specific set of problems, Fara's version of the IR-theory would seem to us to work smoothly only if one assumes a relatively *solip-sistic* view of the contextual interpretation of GAs, such that the value of the interest-bearers parameter fully coincides with the speaker (or at best, as mentioned in the text, with a group of people containing the speaker and in which the objects salient to the speaker are also salient to everyone else). This seems to us to be in tension with the insights in the literature about the *communal* contextual interpretation of GAs (as found e.g. in Lewis 1979: 351–4). We are indebted here to conversations with Laura Delgado.
36. An eternalist analog of Fara's version of the IR-theory would instead require that the content of the elided verb phrase changes with respect to the time it involves. There is independent reason to think that that is possible ("Andy is in the room"; after observing Bill coming in, "And Bill is too"; after observing Charlie coming in, "And Charlie is too"... ; after observing Zac coming in, "And Zac is too"). For those who think that in such cases the time is really contributed by the verb, there is also independent reason to think that the same kind of change is possible even in the case of *bare-argument ellipsis* ("Andy is in the room"; after observing Bill coming in, "And Bill too"; after observing Charlie coming in, "And Charlie too"... ; after observing Zac coming in, "And Zac too").
37. In a particular, *precise* and *predetermined* fashion (which in the text we do not bother to specify). To fix ideas in a natural way, let us say that they change at a constant ratio which in effect coincides with the least one needed in order for the last interests to determine the most relaxed threshold.
38. As one might expect, analogous *temporal* sorites can be constructed. We leave the details of this to the interested reader.
39. As an aside, we would extend this claim to another competitor of standard-contextualist theories, namely nonindexical-contextualist theories positing in the circumstances of evaluation a *distinctive standard parameter* governing the application of a GA (as we have seen in section 12.5 and at the beginning of this section, Fara's version of the IR-theory need not posit any other parameter beyond worlds and times). These theories agree with Fara's version of the IR-theory that, even though the content of the elided verb phrase

cannot change throughout (VPES), its extension can and so they too would seem to have a substantial advantage over standard-contextualist theories. While our arguments in the text are tailored to specifically showing that this is not the case for Fara's version of the IR-theory, they can easily be adapted to this other ilk of nonindexical-contextualist theories. To elaborate a bit, on such theories there should be no objection to introducing an indexical intensional operator "by current standards," such that "by current standards, φ " is true in a context c at a circumstance of evaluation e iff φ is true in c at the circumstance of evaluation which is like e save (possibly) for its standard parameter being that of c . It is then easy to see that substituting "tall by current standards" for "tall" in (VPES) raises for these theories the same problems as those that the original (VPES) raises for standard-contextualist theories. Discussions with John MacFarlane helped to bring out this point.

References

- DeRose, K. (2008). Gradable adjectives: a defense of pluralism. *Australasian Journal of Philosophy*, 86: 141–60.
- Enc, M. (1986). Towards a referential analysis of temporal expressions. *Linguistics and Philosophy*, 9: 405–26.
- Fara, D. (2000). Shifting sands: an interest-relative theory of vagueness. *Philosophical Topics*, 28: 45–81.
- Fara, D. (2008a). Profiling interest relativity. *Analysis*, 21: 1–8.
- Fara, D. (2008b). Generalizing from the instances within "contextualist" accounts of vagueness. MS, Princeton University. <http://www.princeton.edu/~dfara/papers/generalize.pdf>
- Fine, K. (1975). Vagueness, truth and logic. *Synthese*, 30: 265–300.
- Gaifman, H. (2002). Vagueness, tolerance and contextual logic. MS, Columbia University. <http://www.columbia.edu/~hg17/VTC-latest.pdf>
- Heck, R. (2003). Semantic accounts of vagueness. In Beall, J. C. (ed.), *Liars and Heaps*, pp. 106–27. Oxford University Press, Oxford.
- Kamp, H. (1981). The paradox of the heap. In Mönnich, U. (ed.), *Aspects of Philosophical Logic*, pp. 225–77. Reidel, Dordrecht.
- Kaplan, D. (1989). Demonstratives. In Almgong, J., Perry, J. and Wettstein, H. (eds), *Themes from Kaplan*, pp. 481–563. Oxford University Press, Oxford.
- Keefe, R. (2007). Vagueness without context change. *Mind*, 116 (462): 275–92.
- Keefe, R. and Smith, P. (eds) (1997). *Vagueness: a Reader*. MIT Press, Cambridge, Mass.
- Kennedy, C. (1999). *Projecting the Adjective: the Syntax and Semantics of Gradability and Comparison*. Garland Press, New York.
- Kyburg, H. (1961). *Probability and the Logic of Rational Belief*. Wesleyan University Press, Middletown, Conn.
- Lewis, D. (1979). Scorekeeping in a language game. *The Journal of Philosophical Logic*, 8: 339–59.
- Makinson, D. (1965). The paradox of the preface. *Analysis*, 25: 205–7.
- Oppenheimer, D. (2004). Spontaneous discounting of availability in frequency judgment tasks. *Psychological Science*, 15: 100–5.

- Raffman, D. (1994). Vagueness without paradox. *The Philosophical Review*, 103: 41–74.
- Sainsbury, M. (1990). Concepts without boundaries. In Keefe, R. and Smith, B. (eds), *Vagueness: a Reader*. MIT Press, Cambridge, Mass.
- Shapiro, S. (2006). *Vagueness in Context*. Oxford University Press, Oxford.
- Sherman, S., Cialdini, R., Schwartzmann, D., and Reynolds, K. (1985). Imagining can heighten or lower the perceived likelihood of contracting a disease. *Personality and Social Psychology Bulletin*, 11: 118–27.
- Soames, S. (1999). *Understanding Truth*. Oxford University Press, Oxford.
- Soames, S. (2002). Replies. *Philosophy and Phenomenological Research*, 62: 429–52.
- Stanley, J. (2003). Context, interest relativity and the sorites. *Analysis*, 63: 269–80.
- Sweeney, P. (2010). Contextual intolerance. MS, University of Aberdeen.
- Weatherson, B. (2009). Conditionals and indexical relativism. *Synthese*, 116: 333–57.
- Weatherson, B. (2010). Vagueness as indeterminacy. In Dietz, R. and Moruzzi, S. (eds), *Cuts and Clouds*, pp. 77–90. Oxford University Press, Oxford.
- Williamson, T. (1994). *Vagueness*. Routledge, London.
- Zardini, E. (2008). A model of tolerance. *Studia Logica*, 90: 337–68.
- Zardini, E. (2010). First-order tolerant logics. MS, University of Aberdeen.

13

Vagueness and Domain Restriction*

Peter Pagin

13.1 Tolerance principles

In the introduction to their vagueness reader, Rosanna Keefe and Peter Smith classified accounts of vagueness with respect to how they handle the sorites paradox. The sorites paradox is set out in the standard way with reference to a sorites *sequence* s of objects s_1, \dots, s_n and an associated vague predicate F . In S , there is a very small and seemingly negligible difference between any two adjacent elements s_i and s_{i+1} with respect to the dimension that is relevant to satisfying F (for instance, if F is "... is tall," then the dimension is height). This suggests that if s_i satisfies F , then so does s_{i+1} . Since S is a sorites sequence for F it is also stipulated that s_1 satisfies F and that s_n does not. Let t_i denote s_i , $1 \leq i \leq n$. Then the sorites argument is set up as

- (1) 1 $F(t_1)$
2 $\forall i(F(t_i) \rightarrow F(t_{i+1}))$

3 $F(t_n)$

where premise 1 is apparently true, premise 2, the inductive premise, is apparently true, the argument is apparently valid, and the conclusion 3 is apparently false. Keefe and Smith (1996: 10) then say:

The standard alternatives for coping with the sorites paradox are

- (a) reject the validity of the argument
- (b) question or reject the (strict) truth of the inductive premise
- (c) reject the truth of the minor premise or the falsity of the conclusion

- (d) accept the whole reasoning and conclude that the vague predicate is incoherent.

I think it is fair to say alternative (b) has been the most popular one in the literature of recent decades, but this is not the place to review the alternatives.

In Pagin (2010) I presented an alternative strategy for handling the sorites paradox.¹ This strategy is not stated with respect to induction over a sorites sequence but with respect to *general* inductive principles. Roughly, the strategy consists in accepting alternative (a), *reject validity*, for ordinary uses of vague vocabulary, and alternative (d), *incoherence*, for “extreme” uses. The use of a vague predicate in a sorites sequence, on this view, is an extreme use. Because of this, the proposed strategy is not an attempt to *solve* the sorites paradox. Rather, the contradiction is accepted. The main idea is to provide a coherent semantics for ordinary uses of vague vocabulary, *with* (and adequate formulation of) the inductive principle in place.

The inductive principle, stated for the predicate “tall” and with respect to a difference of 1 mm, can be stated as:

- (2) If a man of $n + 1$ mm is tall, then a man of n mm is tall.

This principle is an inductive counterpart to the statement that

- (3) 1 mm cannot make the difference between a man’s being tall and a man’s not being tall.

(3) expresses the intuition that “tall” is insensitive to small differences, together with the intuition that 1 mm is small enough. The property of being insensitive to small differences is what Crispin Wright has called *tolerance*. Wright (1976: 156) says:

What is involved in treating these examples as genuinely paradoxical is a certain *tolerance* in the concepts which they respectively involve, a notion of a degree of change too small to make any difference, as it were. ... Then F is *tolerant* with respect to ϕ if there is also some positive degree of change in respect of ϕ insufficient ever to affect the justice with which F applies to a particular case.

For reasons to be given later, I prefer to speak of linguistic *expressions* as being vague or tolerant, rather than concepts. I shall refer to inductive principles such as (2) as *tolerance principles*. I shall refer to the particular difference given in the principle as the *tolerance level*.

Tolerance intuitions regarding vague expressions are widespread. Most accounts of vagueness involve a rejection of tolerance. Usually, they are coupled with explanations of why we have those intuitions despite the falsity of the tolerance principles.² In some cases, tolerance is nominally “accepted,” but then in the form of a principle for *judging*. For instance, Stewart Shapiro (2006: 8) says, “Suppose that two objects, a , a' differ only marginally in the relevant respect (on which P is tolerant). Then if one competently judges a to have P , then one cannot competently judge a' in any other manner,” and refers to this as a principle of tolerance. Shapiro does not, however, accept tolerance in the present sense, i.e. *semantic* tolerance, as stated for example in (2).³

The leading idea of the present approach is to accept the tolerance intuitions, and to devise a semantics that nonetheless gives a coherent representation of ordinary use of vague vocabulary. Two simple observations lie at the bottom of the approach. First, an informally stated tolerance principle such as (2), is ambiguous between two different regimentations into first-order logic:

- (T1) $\forall y(\text{Man}(y) \ \& \ H(y) \geq n + 1 \rightarrow \text{Tall}(y)) \rightarrow$
 $\forall y(\text{Man}(y) \ \& \ H(y) \geq n \rightarrow \text{Tall}(y))$
 (T2') $\forall x \forall y((\text{Man}(x) \ \& \ \text{Man}(y) \ \& \ H(x) \leq n + 1 \ \& \ H(y) \geq n) \rightarrow$
 $(\text{Tall}(x) \rightarrow \text{Tall}(y)))$

where $H(x)$ is the height of x (where H is the relevant measure function for the predicate “tall”; see below). These principles are nonequivalent. The second is equivalent to

- (T2) $\exists y(\text{Man}(y) \ \& \ H(y) \leq n + 1 \ \& \ \text{Tall}(y))$
 $\rightarrow \forall y(\text{Man}(y) \ \& \ H(y) \geq n \rightarrow \text{Tall}(y)).$

The sharpest difference between (T1) and (T2) concerns domains with *central gaps*.⁴

13.2 Measure functions and central gaps

To define the concept of a central gap we first need the concepts of an *associated weak order*⁵ and a *measure function*:

- (AWO) Let F be a one-place predicate over a domain D . We call \lesssim_F an F -associated weak order iff
- (i) \lesssim_F is a weak order over D

- (ii) For all $d, d' \in D$, if F is true of d and $d \lesssim_F d'$, then F is true of d' .

By condition (ii), the so-called *penumbral* connections of F must hold with respect to the F -associated weak order. Next, the definition of a *measure function*:

(AMF) Let F be a one-place predicate over a domain D . \mathcal{H}_F is a *measure function* for F iff

- (i) \mathcal{H}_F is a function from D to Re (the set of real numbers)
- (ii) there is an F -associated weak order \lesssim_F over D such that for all $d, d' \in D$,

$$\mathcal{H}_F(d) \leq \mathcal{H}_F(d') \quad \text{iff} \quad d \lesssim_F d'.$$

Here $\langle Re, \leq \rangle$ is the relational structure of the Reals with the *less than or equal to* relation. Thus, the measure function is a homomorphism from $\langle D, \lesssim_F \rangle$ to $\langle Re, \leq \rangle$. The existence of an associated weak order is usually assumed for vague predicates, and for any finite or countable weak order, a homomorphism into $\langle Re, \leq \rangle$ exists (cf. Krantz et al. 1971: 15–17, 39).

However, if we only have the condition that the measure function be a homomorphism, then any order-preserving function between the two relational structures qualifies as a measure function. That is, the measure function is unique only up to order-preserving transformations, which is to say that we have a so-called *ordinal scale*. An ordinal scale may be enough for the sorites susceptibility of tolerant predicates in case we only consider measure functions into N , the set of natural numbers. This may be the case, for example, when we compare entities with respect to baldness (number of hairs) or heap-hood (number of grains), where only whole numbers matter for the comparison. For then the tolerance of “bald” amounts to the fact that if anyone is bald, then there is or can be someone who has more hair and still is bald. Having more hair consists in having at least one more individual hair. Then, by repeating the tolerance step any finite boundary of baldness can be crossed, and we will have a sorites contradiction.

This does not work where the order is dense or continuous, as with height. The *taller than* order (over possible objects) is dense, for between any two objects x, y such that x is strictly taller than y , there is a third object z such that z is strictly taller than y and strictly shorter than x .

If in this case tolerance does not amount to more than what can be captured by an ordinal scale, it only consists in the circumstance that for any tall object there is or can be an object that is strictly shorter but still tall. But it is then possible that for a given tall person A of 180 cm there are infinitely many shorter persons of different heights that are tall, although no one of a height less than 179 cm is tall. This is so simply because there are infinitely many rational and hence also real numbers between 180 and 179. So in this case, tolerance in the weak, ordinal, sense does not induce a sorites contradiction.

What we need for tolerance in the normal sense expressed by principles such as (2) is the idea of a margin of a certain fixed size, such as 1 mm, such that no two objects whose difference is within that margin can differ with respect to satisfying some particular predicate.⁶ An ordinal scale does not offer that. From assignments on an ordinal scale we cannot get the information that the difference between two objects is at least as great as the difference between two other objects. We need an *interval scale*, i.e. a measure function that is unique *up to linear transformations*. On an interval scale the ratios between differences are fixed. If

$$\frac{\mathcal{H}_F(d) - \mathcal{H}_F(d')}{\mathcal{H}_F(a) - \mathcal{H}_F(a')} = r$$

and \mathcal{H}' is an admissible (hence linear) transformation of \mathcal{H} on an interval scale, then also

$$\frac{\mathcal{H}'_F(d) - \mathcal{H}'_F(d')}{\mathcal{H}'_F(a) - \mathcal{H}'_F(a')} = r.$$

Specifically, if $r = 1$, the differences are identical. So interval size identity does not vary between admissible functions. In order to construct an interval scale measure function, one needs a procedure for setting intervals as equal (cf. Krantz et al. 1971: Ch. 1). The most direct way of doing so is given by the existence of a *concatenation operation* \circ , i.e. an operation of combining two objects such that the combination can be compared with other objects or combinations. In the case of weight we can weigh two material objects together, and in the case of length we can put them end-to-end. An operation \circ on a domain D qualifies as a concatenation operation iff for all $d, d' \in D$

$$d \circ d' > d, d'$$

(where $>$ is *strictly greater*). The further condition on a measure function \mathcal{H} is then that

$$\mathcal{H}(d \circ d') = \mathcal{H}(d) + \mathcal{H}(d').$$

With a concatenation operation, we have *extensive* measurement.⁷

Given the relevant kind of measure functions, i.e. measure functions that satisfy conditions at least as strict as interval scale uniqueness, we can define the concept of a *central gap*. Let $E(F)$ be the extension of F :

(CG) A domain D has a central gap, with respect to a tolerant predicate F , a measure function \mathcal{H}_F and a tolerance level $\mathcal{F}_F > 0$ for F just in case there is a bounded real number interval (i, j) such that

- (1) $\forall d \in D(d \in E(F) \leftrightarrow \mathcal{H}_F(d) > i)$
- (2) $\forall d \in D(d \notin E(F) \leftrightarrow \mathcal{H}_F(d) < j)$
- (3) $i - j \geq \mathcal{F}_F.$

For the example of height and “tall” and some domain D , this amounts to saying that a central gap in D is an interval (i, j) in, say, millimeters, such that everyone in D is tall iff he is more than i mm, and everyone in D fails to be tall iff he is less than j mm in height, and the difference between i and j is at least as great as the tolerance level. For instance, with a tolerance level of 1 mm we might have a domain with a central gap of 20 mm; everyone is for example: either taller than 180 cm or shorter than 178 cm.⁸ Note that, despite what the term suggests, a central gap need not be “central”: the definition allows that the set of F s in D is either D or \emptyset , in which case there is no sorites sequence anyway. Let us say that a central gap for F in D is *significant* just in case there are both F s and non- F s in D .

The second simple observation is that in a domain with a significant central gap the (T1) principle is false and the (T2) principle is true but does *not* induce the sorites paradox. To verify the first claim, look at the lower endpoint of the gap, assume that there is a man in the domain that is 1779 mm and (regarded as) not tall, and consider an instance of (T1):

- (4) $\forall y((\text{Man}(y) \ \& \ H(y) \geq 1780) \rightarrow \text{Tall}(y)) \rightarrow \forall y((\text{Man}(y) \ \& \ H(y) \geq 1779) \rightarrow \text{Tall}(y)).$

Clearly, because of the gap, the antecedent is true: everyone in the domain that has a height of 1780 mm or higher also has a height that is more than 1800 mm, and hence he is tall. The consequent is false, however, since there is a man of 1779 mm who is not tall. Hence the conditional is false, and therefore also the principle (T1).

For the second claim, first note that at the lower endpoint of the gap and above, (T2) instances will have true consequents:

$$(5) \exists y(\text{Man}(y) \ \& \ H(y) \leq 1781 \ \& \ \text{Tall}(y)) \rightarrow \\ \forall y(\text{Man}(y) \ \& \ H(y) \geq 1780 \rightarrow \text{Tall}(y)).$$

This is because everyone who has a height of at least 1780 mm also has a height of more than 1800 mm, and hence is tall. Second, note that at the upper endpoint of the gap and below, (T2) instances will have false antecedents, since there is no man in the domain of height at most 1800 mm who is tall. Hence, every instance has a false antecedent or a true consequent. Hence, (T2) is true.

For the third claim, observe that in order to detach a false consequent

$$\forall y(\text{Man}(y) \ \& \ H(y) \geq n \rightarrow \text{Tall}(y))$$

we need a true antecedent

$$\exists y(\text{Man}(y) \ \& \ H(y) \leq n + 1 \ \& \ \text{Tall}(y)).$$

But if the consequent is false, then $n < 1780$, and because of the gap, there is no true antecedent with $n \leq 1800$. Hence, even if (T2) is taken as an axiom, in a domain with a significant central gap no instance will occur in a modus ponens inference with a true minor premise and false conclusion. With this form of tolerance, gappy domains are safe from the sorites.

Let us say that a domain is *relaxed* with respect to a vague predicate F , or F -relaxed, iff it has a central gap, and that it is *tight* with respect to F , or F -tight, if it does not. If the domain is F -tight, both (T1) and (T2) induce the sorites paradox for F . These are domains that contain sorites sequences. There are many predicates such that the domain of the real world is tight. For instance, given that no man has much more than 200,000 hairs on his head, it is highly probable, given the population of men on Earth, that for each $i < 200,000$, there is a man with i number of hairs on his head. If one hair cannot make the difference between being bald and not being bald, then it is probable that the domain of men on Earth is tight with respect to the predicate *bald*. How can we then make the use of “bald” in normal circumstances safe from the sorites if tolerance in the form of (T2) is to be accepted?

The proposal in Pagin (2010) was to apply the two familiar ideas of *standard of comparison* and *quantifier domain restriction* in the interpretation of natural language utterances and implement the strategy of central

gaps by means of them. This proposal is in part further motivated by a plausible view of the psychology of applying vague predicates. When a speaker says something like

(6) Sam is tall

she typically tacitly has in mind some clearly positive range of tallness and some clearly negative range of tallness, and places Sam in the upper category. She *disregards* the problematic intermediate range of heights, since there is no need to take a stand on heights in that range in order to properly classify Sam.

To be in “the intermediate range” is not the same as being a *borderline* case in any strict sense. If x is a borderline case of tallness, in some context, then x is in the intermediate range of tallness for that context, but the converse does not hold. What I call “the intermediate range” is a device used to model the psychology of a language user. In what follows I shall *not* assume that a context of utterance uniquely *fixes* an intermediate range, only that there are contextually *admissible* ways of using a central gap semantics for representing that feature of speaker psychology.

Such a semantics can be said to provide an account of vagueness provided it gives reasonable interpretations of the vast majority of ordinary uses of vague expressions, even though it does not in the traditional sense offer a solution to the sorites paradox. Rather, it shows how vague terms can be used in ordinary discourse without generating the paradox. But if too many ordinary uses turn out to be incoherent on the proposed semantics, or implausibly interpreted, then (unless there are better implementations of the basic idea), the approach itself will have failed.

I now turn to the task of setting out and developing this idea.

13.3 Comparisons: classes, standards and gaps

The present account will rely heavily on an appeal to *context-dependence*. Many different ideas about context-dependence have been used in accounts of adjectives in general, and vague adjectives in particular. One well-established idea in this area is that of a *comparison class*. It has some similarity to the present proposal, since one could think of implementing the central gap strategy by means of comparison classes.

The idea of a comparison class comes out naturally when considering utterances of (6) in different contexts. Uttered in the context where the conversation concerns students in a high school class, the statement may be true, while uttered in a different context where the conversation

concerns the members of a basketball team, the statement made *there* may yet be false. This comes out in qualifications with *for*-type prepositional phrases:

- (7) Sam is tall for a high school student.
- (8) Sam is tall for a high school basketball player.

In (7) one (typically) uses the set of high school students as comparison class for heights, and this is made explicit in the sentence by means of the PP. Analogously with (8).

The same comparison classes can also be used implicitly with just (6). This idea was developed by Ewan Klein (1980). Klein's basic idea was to have a function \mathcal{U} that for each context of utterance c picked out a subset $\mathcal{U}(c) \subseteq U$ of the general domain U of individuals (Klein 1980: 14). (6) is then true in a context c just in case John counts as tall with respect to the comparison class $\mathcal{U}(c)$.

There are several problems with this simple model. For one thing, it does not allow for variation of comparison class within quantified sentences. Syntax may require contextual updates without any corresponding change in the *external* context.⁹

A second problem is that the use of complex predicates may involve *different* standards or comparison classes, each related to simple constituent:

- (9) Sam is neither tall nor strong

under a reading where a basketball team is intended to provide the comparison class for height and a wrestling team the comparison class for strength.

A third problem, as stressed by Delia Graff Fara (2000) and by Kennedy (2007), is that a comparison class does not automatically give you a *standard* for determining truth value. For instance, to be tall with respect to some comparison class does not necessarily mean to be above the *average* height of members of the class. It seems then, as Kennedy concludes (2007: 9–10), that we need a contextually determined standard of comparison anyway. It may seem as if the comparison class that is salient in the context suggests some standard, since both the comparison class and the standard are made salient in the context. However, if we do have a standard, then we have what is needed for assessing applications of the context-dependent predicate and do not need a comparison class for this purpose over and above the standard.

Contextually determined standards are usually taken to impose sharp boundaries. The existence of sharp boundaries *prima facie* runs counter to intuitions about vague predicates and about applications of them in ordinary practice. This is compensated for in standard contextual theories by letting the context shift and its associated boundary *move*, as a function of judgments made, of attention shifts, or of assertions made in the conversation. One problem with this strategy is that boundaries are taken to move unbeknownst to or even against the intentions of the speakers.¹⁰ Postulating context shifts under such conditions in order to save consistency may seem *ad hoc*. For example, it is not independently clear why mere change of attention should change the extension of a vague predicate.

In Kennedy's account, the standard of comparison is contextually determined "in such a way as to ensure that the objects that the positive form is true of 'stand out' in the context of utterance, relative to the kind of measurement that the adjective encodes" (Kennedy 2007: 17). This is to ensure by a semantic-pragmatic mechanism that assertions made are safe, i.e. do not apply the vague predicate very close to the boundary. As far as I understand, predications close to the boundary will then have the effect of shifting the standard, therefore moving the boundary and hence (perhaps) changing the extension.

One of the goals of the present strategy is that the shift of contexts be controlled by syntax and pragmatics in ordinary ways. Syntax plus semantics may require that the context c at a particular time t of a conversation is a function of the immediately preceding context c_0 , for example, by letting the domain of c be a restriction of the domain of c_0 . I shall not postulate context change in order to save consistency. On the contrary, controlled and conservative context shifts may, as we shall see, be a way in which inconsistency is produced (in this framework, context shift is not a method of salvation, but a possible road to damnation). This is not quite the whole story, however, since some context shifts may liberate the speaker from constraints induced by earlier contexts.

In the present proposal, central gaps will serve two interrelated functions. One of these functions will be a counterpart to the function of a standard of comparison in other accounts, such as Kennedy's. With respect to a predicate F and a context c , the central gap will partition the contextual domain so that the extension of F at c consists of those objects in the contextual domain whose relevant measures are above the gap, and so that the anti-extension consists of those objects whose measures are below the gap. In this role the central gap will serve to fix the truth values of atomic open or closed sentences Ft .

13.4 Precision: standards, tolerance levels and gaps

David Lewis (1979) introduced the idea of a *standard of precision* as a truth-determining factor of the context of utterance. The standard of precision settles how lax the application of predicate is allowed to be, i.e. how much an object can deviate from strictly having the corresponding property and still count as satisfying the predicate in the context. For instance, J.L. Austin's sentence

(10) France is hexagonal

is true with respect to a low standard of precision and false with respect to any high standard of precision (Austin 1975: 143, Lewis 1979: 352).

It would seem natural to simply equate Lewis's idea of a standard of precision with the graded conception of tolerance, i.e. with the idea of a *level* of tolerance. In particular, this would seem natural, since a standard of *maximal* precision corresponds to a *zero* tolerance level: there is a sharp boundary for the predicate and no deviations from that boundary are allowed.

However, on closer inspection, the two ideas are rather different, at least in their immediate application. That the standard of precision can vary, and that it can be low, does not entail tolerance, and is even inconsistent with tolerance (at least of the (T1) variety). On Lewis's account, a nonmaximal standard of precision still determines a sharp boundary of application of the predicate, only one that allows a wider extension than does the maximal precision standard. With respect to some standard, France is *just* hexagonal-like *enough* to be hexagonal by that standard; anything less hexagonal-like, however little, would not be.

The idea of a standard of precision is analogous to the idea of a *margin of error* in measurement: it tells you the maximal distance between the measured value of some quantity and the true value of that quantity. It does not tell you that it would be correct to assign to the quantity any other value within the margin of error. So in one respect, the idea of tolerance is the opposite: it is analogous rather to an idea of a *margin of accuracy*. With tolerance, it is a necessary condition of an object *o* to satisfy some predicate that any object *o'* that differs from it along the relevant dimension, up to a certain maximal distance, satisfies it too. Since this necessary condition applies again to *o'*, the sorites can be generated.¹¹

Because of this, there is a sense in which tolerance is *reverse precision*. Precise predicates have zero tolerance, maximal reverse precision, while

vague predicates have nonmaximal reverse precision, nonzero tolerance, with a value that depends on the context. So while tolerance level, as a contextually determined value, cannot be fully assimilated to the independently motivated idea of a standard of precision, it is still clearly related to it.

13.5 Domain restriction

The second function of central gaps in the present account is to effect a domain restriction for quantifiers of the object language. The need for quantifier domain restriction is usually introduced by means of examples such as

(11) Everyone left at midnight

where what is meant is that everyone who was present, for example at some particular party, left (the location of the party) at midnight. The domain of “everyone” in (11) is therefore restricted to the set of those taking part in the event. The need to take account of these phenomena in formal semantics has been acknowledged for a long time. A seminal early contribution is Westerståhl (1985).

As suggested by Westerståhl, a domain restriction is effected by means of a context set that is intersected with the antecedently given domain. In case we have binary quantification, of the form QAB , where Q is a determiner, the context set is intersected with the set denoted by the first argument, A . To take an example from Jason Stanley and Zoltan Szabó (2000), with a context set M , the sentence

(12) Every man runs

is given the truth conditions

(12') $T(\text{“Every man runs,” } c) \text{ iff } (\mathbf{Every}(\mathbf{man} \cap M))(\mathbf{run})$

Stanley and Szabó (2000) argue that the phenomenon of domain restriction should be given a semantic rather than a syntactic or pragmatic account. As they define the terms, on a *semantic* account, the context set is given as a contextually determined value to variables in logical form. They further argue that this variable should belong in the nominal that constitutes the first argument to the determiner (“every” in (12)).

The present account will be semantic in the more general sense that context-dependence involves assigning values to context variables provided in the semantic theory. It is often immaterial, however, whether context-dependence is represented in the syntax, and therefore also immaterial whether the account is semantic in the more narrow sense of Stanley and Zsabó.¹²

Here I shall follow the tradition in representing a contextually determined domain restriction as an intersection with the first argument of the determiner denotation. In order to implement the domain restriction of the central gap account, there will be two further restrictions, assuming that both the first and the second argument of the determiner are vague. In a particular context c , a central gap will be associated with each predicate in the lexicon that has been used in the context. Let us assume that each vague lexical item is associated with a measure function. Where F is a predicate, let $\nabla(c, F)$ be the F domain cut, the set of objects whose measures are within the central gap associated with F in c . We take it to return the empty set in case the predicate in question is not vague. Let $\nabla(c, \bar{F})$ be

$$\bigcup_{i=1}^n \nabla(c, F_i),$$

the union of domain cuts for predicates F_1, \dots, F_n in context c . For two atomic predicates used together, the union is represented in Figure 13.1 as comprising two intersecting bands in the plane.¹³ In the restricted domain, no object has a pair of measures in either band. With respect

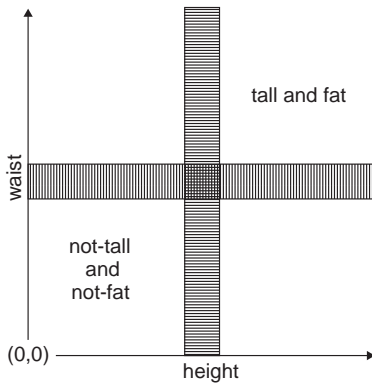


Figure 13.1 Union of gaps

to example (12) we will then have a domain that contains (at most) $(\mathbf{man} \cap M) - (\nabla(c, \mathbf{man}) \cup \nabla(c, \mathbf{runs}))$.

With such extra domain restrictions, we will have what is needed to avoid the sorites paradox in relation to the (T2) principle. This is the first function of the central gap, which is achieved by means of the second: let the extension contain objects with values above the gap, the anti-extension objects with values below the gap, and let there be no objects in the domain with values in between.

13.6 Outline of the semantics

The semantics that I shall propose has the following properties: a *gap function* \mathcal{G} has the role of an interpretation function. It is defined for a predicate F in a domain D iff D is F -relaxed or allows a *restriction* that makes it F -relaxed. When this condition is met for a context c , the interpretation of speakers in c gives a classical (bivalent) semantics where the (T2) principle is true but there are no sharp boundaries for the predicates.

It will be assumed that in each context c , for each vague predicate F that is used in c , a central gap is determined for F . This is an idealization. In any actual occasion of speech, it will be left undetermined exactly what the tolerance level is and where the gap is. It would be more realistic to speak of what is *admissible* on a certain occasion. This is all the more realistic since the contextual parameters will have to depend to some extent on what is uttered on the occasion.

It will be required that a full semantics is given in each context for the full *fragment* of a language that is used in the context, but not for linguistic material outside that fragment. This means in particular that any referring singular term that is used in c must have referent in the quantifier domain of c . For maintaining bivalence, the relevant measure of that term for the predicate F cannot fall in the central gap for F at c . Hence, the central gap is required not to include the measure of that referent. So the location of the gap depends on the terms used in the context. This will be a requirement on any function \mathcal{G} that maps contexts and predicates on central gaps. Nonetheless, it is formally an unnecessary complication to consider a *family* of such gap functions and assume that one of these functions is selected in each context. I shall therefore operate with just one such function.

This function must to some extent be arbitrary (a choice function), but must also be subject to a number of restrictions, ranked according to importance.

- (GAP) (i) The central gap must be selected so that the full fragment of language used in the context, including pragmatically determined contextual updates, is taken account of.
- (ii) The size of the central gap must be at least equal to the tolerance level.
- (iii) The location of the central gap must be selected so as to ensure consistency.
- (iv) The location of the central gap must be selected so that what the speaker says comes out as true, to the extent this is possible and reasonable.

The order of these selection principles is meant to reflect the order of importance. Indeed, (GAPi) cannot be compromised, but this has the consequence that the other three may not be jointly satisfiable, given collateral facts. When they are not, it is not always clear for example that consistency of the speaker should take precedence over the truth of individual judgments. But this is a matter of further investigation.

It is psychologically plausible that the gap is greater, and even considerably greater, than the tolerance level, but for the purpose of setting out the semantics, it is more elegant to reduce the two contextual parameters, tolerance level and gap size, to a single parameter. We can let the gap size *equal* the tolerance level. Then, the requirement (GAPii) will be automatically met.

We can further make the natural stipulation that the *standard of comparison* lies in the middle of the gap. Because of this, the gap is fully determined by two contextual parameters, the standard of comparison, which settles that position of the gap, and the tolerance level, which settles its size: where i is the value of the standard of comparison and k is the tolerance level, the gap is simply $(i + k/2, i - k/2)$. Conversely, by these stipulations, given the upper and lower endpoints of the gap, both tolerance level and standard of comparison can be recovered.

We assume a basic domain D of individuals, and a domain C of contexts. We shall let each atomic n -place predicate, for $n \geq 1$, be associated with an $n + 1$ -place function from $D^n \times C$ to the set of truth values $\{0, 1\}$. However, the vagueness of many-place predicates has not been much discussed in the literature, and I too shall restrict the discussion to one-place predicates, among the atomic ones.

We assume a language L with a finite stock L_A of atomic one-place predicates. We also assume a classical language with negation, conjunction, and universal quantification as primitive, and the others classically defined as usual. We also have a set L^* of *fragments* of L , where each

fragment $l \in L^*$ has the same basic vocabulary as L except for containing only a subset of the set of closed singular terms of L and a subset of the set of vague one-place predicates of L . Each fragment is closed under the same syntactic operations as L itself. We further have a *fragment function* $\mathcal{F} : C \rightarrow L^*$ from contexts to fragments of L . Let $\mathcal{F}_0(c)$ be the set of one-place predicates in $\mathcal{F}(c)$ and s^0 be the set of one-place predicates in the sentence s . Then we shall require that

(FP) If s is used in context c , then $s^0 \subseteq \mathcal{F}_0(c)$.

With each atomic predicate $F_i \in L_A$ is associated a weak order \lesssim_{F_i} and an admissible (at least interval scale) measure function \mathcal{H}_{F_i} . We further have a function $\mathcal{T}_c(F)$ that for a context c and an atomic predicate F returns a *tolerance level* for F at c . We also have a *central gap function* $\mathcal{G}_c(F)$ from a predicate F and a context c to pairs (i, j) of real numbers, $i \geq j$, with respect to the measure function \mathcal{H}_F . If F is vague, then $i > j$. We shall refer to the left member as $\mathcal{G}_c^+(F)$ and to the right member as $\mathcal{G}_c^-(F)$.

To maintain consistency we need the condition that:

(GT) For every predicate $F \in L$ and context $c \in C$ it holds that $\mathcal{G}_c^+(F) - \mathcal{G}_c^-(F) \geq \mathcal{T}_c(F)$.

That is, the central gap of the context must be at least as great as the tolerance level of the context. As already suggested, however, it will be convenient to let the size of the central gap *equal* the tolerance level, since then an increase of precision, i.e. a reduction of the tolerance level, automatically reduces the central gap. Hence, we stipulate that:

(GT⁼) For every predicate $F \in L$ and context $c \in C$ it holds that $\mathcal{G}_c^+(F) - \mathcal{G}_c^-(F) = \mathcal{T}_c(F)$.

When several atomic predicates are used in the same context $c \in C$, it will be required that individual terms do not have referents with measures in *any* of the gaps associated with the predicates used in c , hence that they have referents outside $\nabla(c, \bar{F})$.

Let f, f' etc. be assignment functions from variables x_1, x_2, \dots in L to objects in D . The semantic function $\llbracket \cdot \rrbracket$ maps a sentence s , a context c and an assignment function f onto a truth value in $\{0, 1\}$. Let $\llbracket F \rrbracket_c^+ = \{a : \mathcal{H}_F(a) > \mathcal{G}_c^+(F)\}$, and let $\llbracket F \rrbracket_c^- = \{a : \mathcal{H}_F(a) < \mathcal{G}_c^-(F)\}$. For a singular term t , let $\llbracket t \rrbracket_f$ be $\llbracket t \rrbracket$ if t is a constant, and $f(t)$ if t is a variable. The relation $f'[x]f$ holds iff the assignment f' differs from f at most in what

it assigns to x . Then we have the truth definition:

- (GS) (i) Where t is a constant, $\llbracket t \rrbracket \notin \nabla(c, \bar{F})$
 (ii) Where F is an atomic one-place predicate,
 $\llbracket Ft \rrbracket_{f,c} = 1$ iff $\llbracket t \rrbracket_f \in \llbracket F \rrbracket_c^+$
 (iii) $\llbracket \neg A \rrbracket_{f,c} = 1$ iff $\llbracket A \rrbracket_{f,c} = 0$
 (iv) $\llbracket A \& B \rrbracket_{f,c} = 1$ iff $\llbracket A \rrbracket_{f,c} = 1$ and $\llbracket B \rrbracket_{f,c} = 1$
 (v) $\llbracket Qx(Ax, Bx) \rrbracket_{f,c} = 1$ iff
 $\llbracket Q \rrbracket_{f'}(f'[x]f$ and $\llbracket Ax \rrbracket_{f',c} = 1$ and $f'(x) \notin \nabla(c, \bar{F})$, $\llbracket Bx \rrbracket_{f',c} = 1$).

According to (GSv), in a context c a sentence $Q(Ax, Bx)$ is true just in case $\llbracket Q \rrbracket$ many individuals of which Ax is true in c and that do not belong in the contextually determined cut $\nabla(c, \bar{F})$ are such that Bx is true of them as well.

We can verify that a tolerance principle such as

- (13) Some $x(\text{man}(x) \& \text{height}(x) \leq n + k \text{ mm}, \text{tall}(x)) \rightarrow$
 All $x(\text{man}(x) \& \text{height}(x) \geq n \text{ mm}), \text{tall}(x)$

is true in any context c where the tolerance level for “tall” is k mm or greater. We assume that the relevant measure function maps men on their heights in mm. Let us assume here that “man” is nonvague. Then the antecedent of (13) is true just in case some individual a in the restricted domain is a man and has a height above the upper edge of the “tall” gap in c . Then any individual b in the c -restricted domain that is a man and has a height at least that of the height of a minus k mm, itself has a height above the upper edge of the gap, for there is no individual in the restricted domain that has a height in the gap, and any individual c in the domain with a height below the gap is more than k mm shorter than a . So the consequent of (13) is true.¹⁴

13.7 Some consequences

In this section I shall consider five consequences of the gap semantics that *prima facie* are negative for the account.

13.7.1 Filling the gap

The function \mathcal{G} is designed to determine a central gap in a context c so that no referent of a term used in c belongs to the domain cut that results. Moreover, \mathcal{G} is to satisfy *charity* as far as is reasonable (GAPiv). And everything explicitly talked about, for example by direct enumeration

of individuals, must count. As a result, so many objects might be mentioned that no gap that agrees with the verdicts of the speaker still exists. Then the domain may contain a sorites sequence of individuals d_1, \dots, d_n such that d_1 has a height above the gap initially given as the discourse starts, d_n has a height below the initial gap, and between each d_i, d_{i+1} the difference in height is at most that of the tolerance level.

This is then a sorites sequence with respect to the current tolerance level τ and one particular standard of comparison κ that best agrees with the height verdicts of the speaker. The result will be a sorites contradiction. As \mathcal{G} is specified, a standard of comparison must be chosen that makes one or more of the speaker's verdicts come out false, but preserves consistency, such that the gap is placed above the upper end of the measures of the sorites sequence, or below the lower end. The speaker will then be represented as mistaken in her verdicts about many of the individuals in the domain.

An alternative scenario makes use of the requirement of respecting conservative context updates. This may happen for example by various anaphoric constructions that enforce a conservative context update. One possibility is verb-phrase ellipsis, as considered by Jason Stanley (2003) in his objection to contextualism about vagueness. We can have a discourse by a speaker S that evolves in the following manner:

(14) John is tall. And so is Alice, and Noah, and ..., but not ..., although ... etc.

which intuitively demands that all the individuals named belong in the domain talked about. Clearly it can happen that the resulting domain contains a sorites sequence. The result is again a contradiction. In this case, however, the value of the \mathcal{G} function is pragmatically restrained by the discourse, and must be the same at the end as in the beginning. But that value is not admissible. As a result, no admissible interpretation can be given. The speaker must be represented as not fully interpretable.

In some cases there will be a trade-off between representing the speaker as guilty of errors of estimates, or of inconsistency. This seems to me to be perfectly in order, given that the contextual tolerance level really reflects the disposition of the speaker to accept tolerance principles, and thereby reflects the speaker's lack of precision.

13.7.2 Emptying the domain

An opposite negative effect of the semantics, together with lack of precision of the speakers, is that if many vague predicates are used in some

particular context c , and the lack of precision of the relevant speakers enforces a considerable domain cut for each or many of those predicates, then the result may be a virtually empty contextual domain, contrary to what the speakers in the context may think about it.

Maybe this result is intuitively acceptable to a somewhat lower degree. On the other hand, this consequence, as well as the consequence above, would be expected to be rarely exemplified, since it relies on a combination of very low precision, the use of many vague predicates, and pragmatic discourse features that forces these predicates to be treated as used in one and the same context.

13.7.3 Overly exact statements

Suppose we have an initially given domain, such as the domain of Swedish male citizens between 20 and 60 years old. Suppose this domain contains exactly 2,500,000 individuals. Suppose speaker Nils in context c says

- (15) Exactly 1,600,000 Swedish male citizens between 20 and 60 years old are tall.

Suppose further that Nils speaks with a tolerance level of 20 mm with respect to “tall” in c . Now, we can be sure that such a large domain as this will be tight with respect to that tolerance level, and hence contain a sorites sequence. Because of this, the central gap determined by the gap function \mathcal{G} and the contextual parameters will induce a nonempty cut in the domain. Say that the cut $\nabla(c, \text{tall})$ contains 50,000 individuals.

Now, the salient reading of (15) is that which entails

- (16) For any number k , if k is the number of Swedish male citizens between 20 and 60 years old, then exactly $k - 1,600,000$ Swedish male citizens between 20 and 60 years old are not tall.

The result is that the sum of those that are said to be tall and those that are entailed by this not to be tall, is larger than the size of the *restricted* domain, which has a size of 2,450,000 individuals.¹⁵ Hence, under the domain restriction interpretation, (15) is false. Either there are fewer tall individuals in the restricted domain than what is explicitly claimed, or there are fewer nontall individuals in the restricted domain than what is entailed by (15) together with background truths, or both.

I think this an intuitively correct consequence. Exact claims about such large populations should not be made with such a high level of tolerance.

With high tolerance levels you can make statements that are capable of being approximately true, but claims of such exact character will turn out false.

Sentences like (15) will in general pose problems for contextualist accounts of vagueness, such as Raffman (1994) or (1996), or Fara (2000). They will be true only if there is an overly sharp boundary between two individuals of almost the same height, and those two individuals are identified by the sentence, e.g. number 1,900,000 and number 1,900,001 in the height ranking of Swedish male citizens between 20 and 60 years old. Such judgments are not deemed acceptable by contextualist accounts, and so there must be mechanisms that move the boundary elsewhere, and hence the claim must be considered false. It cannot be the ordinary mechanism (by which considering two highly similar individuals induces a boundary shift so that the two individuals are placed in the same category), simply because no two highly similar individuals are considered. Moreover, the shifting usually appealed to is supposed to be *true-making*, not *false-making*. In some contextualist accounts, such as Shapiro's, the speaker who asserts (15) must be considered *incompetent*.

But it has a peculiar status also on other accounts. On an epistemicist account, (15) is unknowable (even if true), unless one employs a very precise measuring method. On a supervaluationist account, (15) is necessarily not super-true (since that is inconsistent with the existence of truth-value gaps), but is still not super-false, since it is true under one sharpening (assuming the example is in the borderline area). On a degree-theoretic account, it must have a truth value close to that of a contradiction, since it entails a conjunction $p \& \neg q$ concerning the tallness status of individuals number 1,900,000 and number 1,900,001 in the height ranking, where p and q must have almost the same truth value.

13.7.4 Complication of inference

A feature of the gap account is the need for tracking the use of vague predicates across contexts. This complicates the account of overt reasoning. Consider a discourse with the following three sentences uttered in sequence:

- (17) a. Most men are rich.
 b. Most men are tall.
 c. Hence, some man is rich and tall.

Suppose we are considering an initially given domain where even with a reasonable richness cut, (17a) is true. Similarly, even with a reasonable

tallness cut (17b) is true. The third sentence, (17c), is a logical consequence of the first two. However, it contains both ‘rich’ and ‘tall’, and so the richness cut and the tallness cut must be combined for evaluating (17c). But this may have the effect of reducing the domain so that some rich men go out because of the tallness restriction and some tall men go out because of the richness restriction. It may then happen that this leaves the domain without any man that is both rich and tall. Hence, in the context, (17c) is false.

This shows that to get the logic right, we must somehow ensure that the same domain is considered across the discourse. This is troublesome in cases where the required restriction is effected *retroactively*: a sentence occurring later in the discourse restricts the domain for a sentence that has occurred earlier, for the sake of the inference.

Note, however, that it happens only to the extent that the earlier sentence is used as a premise in reasoning that is performed at the later stage. It is standard in reasoning that one must check that premises used at an earlier stage of a discourse employ words without equivocation compared with uses in later premises or the conclusion, that context-dependent expressions are used with the same values across the discourse, etc. Therefore, when using an earlier stated premise, one needs to check that it still holds in the later context, and perhaps reaffirm it in the new context. From this perspective, the complication is not radically new.

13.7.5 Content preservation

The last consequence I shall consider here is the one I see as potentially most worrying. It concerns whether the intuitive *utterance content* is preserved by the gap account.

Normally, when domain restriction is introduced in natural language semantics, it is introduced to get the intuitively perceived content right. For instance, in

(11) Everyone left at midnight

we introduce the domain restriction to arrive at the intuitive content that everyone *at the party* left at midnight.

In the gap semantics, it is not so clear that the domain restriction tracks the intuitive content of the utterance. When the average speaker Nils says

(18) Most Swedish women are tall

it is not likely that he *intuitively* wants to convey that most Swedish women in a *suitably restricted* subset of the set of Swedish women are tall. And if not, it appears that the semantics simply gets the intended content wrong.

The intuition behind this objection is pretty strong. However, it also seems that, on reflection, there is more to be said on this matter in defense of the proposal. First, it is rather clear, I think, that the speaker Nils in this case does speak with nonmaximal precision. He does not have a sharp demarcation of tall women, or tall Swedish women, in mind. It would therefore *not* be fully accurate to represent Nils as saying something equivalent to

- (19) For any number k , at the current time t , if k is the number of Swedish women at t , then the number of *tall* Swedish women at t is greater than $k/2$.

Rather, it seems intuitively right to say that Nils has some kind of *typicality* judgment in mind: it is *typical* of Swedish women that most of them are tall. Since it is typical but not precise, it should be *approximately* true. And taking it as approximately true is taking it as literally true *with respect to* an approximation. Again, it seems reasonable to equate the idea of an approximation with the idea of a standard of precision, and therefore in the case of vagueness, with the tolerance level. Statement (18) is meant to be true with respect to the accepted imprecision of “tall.”

The next step is to ask how *truth* is approximated by means of a tolerance level. *One* way of approximating truth by way of the tolerance level is to make a statement that is *literally* true with respect to a domain that is approximately the same as the original domain, and where the difference is induced by the standard of precision.

On this way of seeing the matter, the gap semantics offers one way of spelling out the approximation idea such that the intuitive truth of tolerance principles is preserved.

Notes

*The ideas in this chapter were presented at the conference *Vagueness and Language Use*, in Paris, April 2008. I am grateful to several of the participants for comments, including Paul Égré, Robert van Rooij, Agustín Rayo, and Stephen Schiffer. I am as grateful for helpful comments from an anonymous referee, as well as from the editors. And, as always I am indebted to Kathrin Glüer-Pagin for discussion of the problems over the years.

Accounts of vagueness similar to the present one, giving an important role to *gaps* in the domain, have been proposed by Ruth Manor (2006), Haim Gaifman (2002), Robert van Rooij (this volume), and Mario Gómez-Torrente (2010). For brief comparisons between those accounts and the present proposal, see note 14.

I first came up with these ideas when in 1997 reading Williamson (1994). I presented them in two international conferences in 1998 (Putnam conference in Karlovy Vary, *Meaning and Interpretation* conference in Stockholm). I took up the theme again in a grant application in 2004, and presented the basics of the account in a vagueness workshop in St Andrews in 2006 (published as Pagin 2010). I was told about Manor's work by Tim Williamson after he learned about mine, while Gaifman's work was unknown to me until after completing the first draft of this chapter, brought to my attention by an anonymous referee.

The work was sponsored by the Swedish Research Council (VR) grant *Vagueness and Context Factors*.

1. Concerning similar accounts and the question of originality, see note 14.
2. For instance, Tim Williamson (1994) offers his elegant margin-of-error principle as justification for explaining why we cannot have knowledge of the application of vague predicates in borderline cases. But this does not by itself account for the fact that speakers *believe* in the truth of those principles.
3. Tolerance is accepted by Elia Zardini (2008). Zardini avoids the sorites paradox by a restriction on logical consequence.
4. (T1) and (T2) are characterized model-theoretically in Pagin (2010).
5. A weak order \lesssim over a domain D is a binary relation over D satisfying the conditions

- (i) For all $d, d' \in D$, $d \lesssim d'$ or $d' \lesssim d$ (connectedness)
- (ii) For $d, d', d'' \in D$, if $d \lesssim d'$ and $d' \lesssim d''$, then $d \lesssim d''$ (transitivity)

6. Strictly speaking, the sorites contradiction arises even if the tolerance margin d_i between s_i and s_{i+1} depends on i and is monotone decreasing ($d_{i+1} < d_i$), as long as the series $d_1 + d_2 + \dots$ diverges.
7. One can then create a so-called *standard sequence* by choosing an object e as having a unit measure, $\mathcal{H}(e) = 1$, and comparing arbitrary objects with n -ary concatenations of exact replicas of e (cf. Krantz et al. 1971: Ch. 1).
8. More realistically, the right-hand side of condition (CG2) should be weakened: either the measure of d is less than j , or the measure function is *not defined* for d , since not everything has a length. I shall ignore this complication in the rest of the chapter.
9. There are quantified sentences like

- (i) Everyone in my family is tall

which Christopher Kennedy (1999; 2007, 8) glosses as

- (i') For every x in my family, x has a height greater than the norm for someone like x , where the relevant kind of similarity (same age, same sex, etc.) is contextually determined.

If there is only one comparison class determined by context, then this reading where comparison class or standard may depend on the individual to which tallness is ascribed, is not available.

It is not so clear how to handle this phenomenon. I am inclined to think that in the case of (i) there is tacit pragmatic enrichment to “tall,” which in fact makes the property ascribed *different* in each universal instantiation (like e.g. *tall for her age*), rather than the same, and this calls for a treatment that is anyway orthogonal to issues discussed here.

10. Shapiro (2006: 35) calls this “the Heraclitus problem” and objects (rightly) that rapid changes in extension of predicates reduce the usefulness of classical logic (I guess it would reduce the usefulness of several other logics as well). Shapiro’s own conversational score account is meant to minimize this problem with contextual accounts.
11. In Williamson (1994), there is in this sense a margin for accuracy of *knowledge*: a subject can know that an object *o* has a property *F* only if any object *o'* differing from *o* in a certain maximal way, has the property *F* as well. This margin of accuracy is in turn induced by a margin of error of the method of measuring employed (determining the number of leaves on a tree by looking at it from a distance).
12. This is exemplified in Pagin (2005).
13. This can also be regarded as the two-dimensional gap for complex predicates such as “is tall and fat” and “is tall or fat.”
14. As mentioned above, similar accounts have been proposed by Ruth Manor, Haim Gaifman, Robert van Rooij, and Mario Gómez-Torrente. On Manor’s account (2006), sentences with vague predicates are used with the *presupposition* that there is a gap that is large enough. In contexts where the presupposition is true, the standard of comparison is determined to be in the gap, so that objects both in the extension and in the anti-extension have measures outside. In contexts where the presupposition is false, sentences involving the predicate are all considered false (because the existence of a gap is stated with large scope in the logical form). Gómez-Torrente’s account (forthcoming, 2010) is similar to Manor’s but more radical: on occasions of use where speaker preconceptions of tolerance, clear cases and domain of individuals do not induce a suitable gap in the domain, the relevant vague predicate lacks extension and intension, and so the sentence uttered lacks truth conditions.

Gaifman’s account (2002) is similar to these in that it relies on the existence of actual gaps. It differs in that in contexts where there *is not* any gap for a particular tolerant predicate, the predicate simply loses its tolerance and thus becomes sharp. van Rooij (this volume) differs from Manor and Gaifman, in that the gap that is required belongs in a contextually relevant *comparison class*.

These alternatives have serious problems. In Manor’s case, the drawback is that all clear case applications of tolerant predicates counterintuitively come out false in case there is no actual gap. In Gómez-Torrente’s case, they lack truth conditions. In Gaifman’s case, although he claims that tolerance is a consequence of meaning-determining rules (p. 18), it simply disappears in case nonlinguistic circumstances are not accommodating, a combination of views that I find somewhat bizarre. van Rooij’s account does not have such obvious problems, and works well for singular predications. It does not,

however, in general work for quantification, since if the quantifier domain itself contains a sorites sequence, the comparison class does not help.

15. If you dispute that (16) is a consequence of (15), then complicate (15) with (16) as a second conjunct.

References

- Austin, J. L. (1975). *How to do Things with Words*, 2nd edn. Oxford University Press, Oxford.
- Fara, D. G. (2000). Shifting sands: an interest-relative theory of vagueness. *Philosophical Topics*, 28: 45–81.
- Gaifman, H. (2002). Vagueness, tolerance, and contextual logic. Unpublished MS.
- Gómez-Torrente, M. (2010). The sorites, linguistic preconceptions, and the dual picture of vagueness. In Dietz, R. and Moruzzi, S. (eds), *Cuts and Clouds* pp. 228–53. Oxford University Press, Oxford.
- Keefe, R. and Smith, P. (1996). Introduction: theories of vagueness. In Keefe, R. and Smith, P. (eds), *Vagueness: a Reader*, pp. 1–57. MIT Press.
- Kennedy, Christopher (1999). Gradable adjectives denote measure functions, not partial functions. *Studies in the Linguistic Sciences*, 29(1).
- Kennedy, C. (2007). Vagueness and grammar: the semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy*, 30: 1–45.
- Klein, E. (1980). A semantics for positive and comparative adjectives. *Linguistics and Philosophy*, 4: 1–45.
- Krantz, D. H. et al. (1971). *Foundations of Measurement*, Vol. I. Academic Press. Republished by Dover Publications, New York, 2007.
- Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, 8: 339–59. Reprinted in Lewis (1983), pp. 233–49.
- Lewis, D. (1983). *Philosophical Papers*, Vol. I. Oxford University Press, Oxford.
- Manor, R. (2006). Solving the heap. *Synthese*, 153: 171–86.
- Pagin, P. (2005). Compositionality and context. In Preyer, G. and Peter, G. (eds), *Contextualism in Philosophy*, pp. 303–48. Oxford University Press, Oxford.
- Pagin, P. (2010). Vagueness and central gaps. In Dietz, R. and Moruzzi, S. (eds), *Cuts and Clouds* pp. 254–72. Oxford University Press, Oxford.
- Raffman, D. (1994). Vagueness without paradox. *Philosophical Review*, 103(1): 41–74.
- Raffman, D. (1996). Vagueness and context-relativity. *Philosophical Studies*, 81: 175–92.
- Shapiro, S. (2006). *Vagueness in Context*. Oxford University Press, Oxford.
- Stanley, J. (2003). Context, interest relativity and the sorites. *Analysis*, 63: 269–80.
- Stanley, J. and Szabó, Z. (2000). On quantifier domain restriction. *Mind and Language*, 15: 219–61.
- Westerståhl, D. (1985). Determiners and context sets. In van Benthem, J. and Alice ter Meulen, A. (eds), *Generalized Quantifiers in Natural Language*, pp. 45–71. Foris, Amsterdam.
- Williamson, T. (1994). *Vagueness*. Routledge, London.
- Wright, C. (1976). Language-mastery and the sorites paradox. In Evans, G. and McDowell, J. (eds), *Truth and Meaning*, pp. 223–47. Clarendon Press, Oxford.
- Zardini, E. (2008). A model of tolerance. *Studia Logica*, 90: 337–68.

Author Index

- Austin, J. L., 293
- Balci, F., Freestone, D., and Gallistel, C. R., 38
- Barker, C., 9, 14, 15, 109, 129, 130, 142 n. 3, 143 n. 10, 161, 165–74, 178–89, 181–3, 188 n. 16, 197
- Barker, C. and Taranto, G., 14, 165, 169, 191
- Barth, H., Kanwisher, N., and Spelke, E., 37–8
- Bartsch, R. and Vennemann, T., 5, 10, 80
- Beall, J. and van Fraassen, B., 8
- Beck, S., 155
- Beck, S., Oda, T., and Sugisaki, K., 25–6, 46 n. 2
- Bierwisch, M., 46 n. 3, 144 n. 20
- Black, M., 31, 47 n. 8
- Blackburn, P., Rijke, M., and Venema, Y., 220 n. 4
- Bogal-Allbritten, E., 95 n. 7, 96 n. 12
- Bogusławski, A., 75, 81
- Bonini, N., Osherson, D., Viale, R., and Williamson, T., 2
- Bosch, P., 4
- Boysen, S. T. and Berntson, G. G., 38
- Brannon, E. M. and Terrace, H. S., 38
- Burgess, J. A., 220 n. 1
- Chierchia, G., 95 n. 8
- Cobrerros, P., 7, 9, 14–16, 220 n. 8, 222–3, 233–8, 245–7
- Cohen, A., 8–9, 14–15, 161, 191
- Cordes, S., Gelman, R., Gallistel, C. R., and Whalen, J., 38
- Cresswell, M. J., 95 n. 7, 155
- Dehaene, S., 38
- Delgado, L., 280 n. 35
- DeRose, K., 276 n. 12
- DeVault, D. and Stone, M., 2
- Douven, I., 172
- Dummett, M., 59
- Edgington, D., 46 n. 3, 220 n. 7
- Égré, P., 47 n. 20, 94 n. 6, 182, 188 n. 18
- Égré, P. and Bonnay, D., 47 n. 20
- Enç, M., 188 n. 22, 276 n. 7
- Fara, D., 4–5, 7, 9, 14–18, 26, 44, 73–5, 81–2, 84, 87, 95 n. 10, 105–6, 207–8, 210, 212, 214–16, 218–19, 220 nn. 7, 12, 246 n. 2, 249–56, 259–60, 262, 264–74, 276 n. 2, 277 nn. 18, 21, 278 nn. 22, 24, 279 nn. 25, 26, 31, 280 nn. 35, 36, 39
- Feigenson, L., Dehaene, S., and Spelke, E. S., 38
- Field, H., 4
- Fine, K., 4, 7, 30, 59, 101, 220 n. 6, 246 n. 3, 266, 276 n. 13, 277 n. 19
- Fox, D. and Hackl, M., 163 n. 8
- Frege, G., 40, 53, 70 n. 2
- Fulst, S., 8, 9, 11, 18, 46 nn. 1, 2
- Gaifman, H., 61, 70 n. 10, 249, 305, 306 n. 14
- Gajewski, J., 163 n. 8
- Gallistel, C. R. and Gelman, R., 38, 39, 47 nn. 11, 14
- Gallistel, C. R. and Gibbon, J., 28, 45
- Gelman, R. and Gallistel, C. R., 47 n. 11
- Gescheider, G. A., 47 nn. 9, 11, 12, 18
- Geurts, B., 161, 163 nn. 10, 11
- Gibbon, J. and Fairhurst, S., 38
- Ginzburg, J., 176
- Gomez-Torrente, M., 71 n. 10
- Goodman, N., 7
- Grano, T., 95 n. 8, 96 n. 12

- Groenendijk, J., Stokhof, M., and Veltman, F., 114–15
- Halpern, J., 7, 175
- Hampton, J., 4
- Hauser, M. D., MacNeilage, P., and Ware, M., 38
- Hawthorne, J., 15
- Heck, R., 220 n. 7, 279 n. 25
- Heim, I., 46 n. 3, 95 n. 7, 155
- Hobbs, J. and Kreinovich, V., 127
- Hyde, D., 8
- Jaegher, K. de and van Rooij, R., 70
- Kadmon, N. and Landman, F., 158–9
- Kamp, H., 4–5, 10, 26, 30, 52, 59–61, 70 n. 10, 73, 87–8, 101, 106–7, 110, 122–3, 249, 270
- Kamp, H. and Partee, B., 4, 26, 75, 106
- Kanazawa, M., 142 n. 9
- Kaplan, D., 276 n. 1
- Katz, G., 151, 153, 155
- Keefe, R., 1–3, 7, 9, 60, 220 n. 6, 246 nn. 1, 3, 247 n. 10, 278 n. 24
- Keefe, R. and Smith, P., 47 n. 7, 283
- Keller, J. A., 187 n. 4
- Kennedy, C., 2, 4–6, 8–13, 26, 29, 46–7, 51–4, 69, 70 n. 3, 76–82, 94 nn. 1, 5, 6, 95 nn. 8, 9, 101, 116 n. 1, 122, 142 n. 1, 252, 265, 291–2, 305 n. 1
- Kennedy, C. and McNally, L., 78, 124, 132, 135, 136, 139
- Klein, E., 4, 5, 10, 30, 33, 47 nn. 10, 21, 51–3, 56–7, 59–60, 62, 67–8, 73, 75, 76, 88–90, 102, 116 n. 1, 162 n. 6, 291
- Kölbl, M., 169
- Krantz, D. H., Luce, R. D., Suppes, P., and Tversky, A., 47 nn. 9, 10, 14, 55, 69, 286–7, 305 n. 7
- Krantz, S. G., 188
- Kratzer, A., 166–7, 197
- Krifka, M., 5, 125–6, 142 n. 6, 171, 174, 188 nn. 19, 21
- Kripke, S., 12, 100, 103, 108, 112
- Kyburg, H., 15, 168, 172, 196, 261
- Lakoff, G., 4, 12, 122
- Landman, F., 108
- Lappin, S., 7
- Laserson, P., 6, 122–3, 131, 141, 143 n. 13, 169
- Lassiter, D., 130, 142 n. 8
- Lewis, D., 3, 4, 8, 26, 52, 100, 102–5, 108, 113, 257, 280 n. 35, 293
- Liu, C-S. L., 96 n. 12
- Luce, R. D., 7, 11, 39, 47 nn. 14, 15, 58, 66
- Machina, K. F., 46 n. 3
- Manor, R., 305–6
- Marsack, C. C., 94 n. 3
- McCready, E., 144 n. 17, 149–50
- McGee, V. and McLaughlin, B., 220 nn. 9, 11
- McLachlan, G. and Peel, D., 220 nn. 9, 11
- Moltmann, F., 116 n. 1
- Morzycki, M., 151–3, 155–6, 161, 163 n. 7
- Moyer, R. S. and Landauer, T. K., 38
- Murphy, G., 98
- Nouwen, R., 8, 13–15, 162
- Olthof, A., Iden, C. M., and Roberts, W. A., 38
- Oppenheimer, D., 277 n. 18
- Pagin, P., 9, 17–18, 47 n. 16, 70 n. 10, 284, 289, 305 n. 4, 306 n. 12
- Parikh, R., 8
- Parkman, J. M., 38
- Partee, B., 5, 7
- Penka, D., 139
- Pietroski, P., 47
- Pinkal, M., 6, 44, 47 n. 5, 59–61, 122, 142 n. 1
- Piñon, C., 170–1, 176, 183, 188 n. 19, 195
- Pollock, J., 172
- Putnam, H., 83
- Quine, W. V., 9

- Ramchand, G., 143 n. 11
 Raffman, D., 4, 7, 60, 74, 249, 270, 302
 Rapp, I. and von Stechow, A., 139
 Richard, M., 75, 81
 Rotstein, C. and Winter, Y., 132
 Russell, B., 3, 45, 116 n. 1
 Ryan, S., 172
- Sainsbury, M., 220 n. 7, 278 n. 21
 Sanford, A., Dawydiak, E., and Moxey, L., 160
 Sapir, E., 3, 4, 10, 70 n. 3, 76
 Sassoon, W. G., 8, 10, 11, 12, 14, 100, 117 n. 8
 Sauerland, U. and Stateva, P., 6–8, 12–14, 44, 109
 Sawada, O., 77, 94 n. 5
 Schiffer, S., 4
 Schwarz, B., 136
 Schwarzschild, R., 95 n. 7
 Schwarzschild, R. and Wilkinson, K., 46 nn. 3, 4
 Scott, D. and Suppes, P., 39, 58–9
 Searle, J., 112
 Searle, J. R. and Vanderveken, D., 184
 Seuren, P., 52
 Shapiro, S., 4, 7, 60, 70 n. 7, 249, 270, 285, 302, 306 n. 10
 Sherman, S., Cialdini, R., Schwartzmann, D., and Reynolds, K., 277 n. 18
 Siegel, M., 143 n. 11
 Smith, N., 1, 2
 Soames, S., 4, 12, 15, 26, 73, 249, 270
 Sorensen, R., 4, 106
 Stalnaker, R., 101, 176
 Stanley, J., 81, 270, 271, 274, 294–5, 300
 Stassen, L., 77, 93
 Stechow, A. V., 46 nn. 3, 4, 51, 52, 70 n. 2, 95 n. 7, 139, 155
 Sweeney, P., 275
- Sweeney, P. and Zardini, E., 5, 9, 17
 Sybesma, R., 96 n. 12
 Syrett, K., Kennedy, C., and Lidz, J., 76, 85
- Taranto, G., 191, 202
 Tye, M., 46 n. 3
- Unger, P., 12–13, 78
- van Benthem, J., 51–7, 62–3, 67, 88
 van Deemter, K., 4, 8, 10
 van Fraassen, B. C., 8, 101
 van Rooij, R., 3, 7–9, 11, 18, 39, 46 n. 1, 62, 65, 70, 71 nn. 13, 14, 73, 87–94, 304–5, 306 n. 14
 Vanderveken, D., 171, 175
 Veltman, F., 70 n. 9
- Waismann, F., 70 n. 9, 71 n. 13, 100, 106, 117 n. 7
 Weatherson, B., 220 n. 1
 Westerståhl, D., 294
 Whalen, J., Gelman, R., and Gallistel, C. R., 38
 Wheeler, S., 10, 73, 75, 82, 88–90
 Williamson, T., 4–5, 7, 14–15, 26, 70 n. 5, 71 n. 12, 74, 83–4, 86, 106, 122–3, 129, 213, 220 nn. 1, 8, 9, 233, 247 nn. 10, 11, 257, 305, 305 n. 1, 306 n. 11
 Wittgenstein, L., 61, 70 n. 9
 Wolf, L. and Cohen, A., 8–9, 14–15, 161, 175, 191
 Wright, C., 2, 4, 16, 29, 187, 220 n. 7, 284
 Wynn, K., 38
- Zanuttini, R. and Portner, P., 152, 161
 Zardini, E., 275, 305 n. 3
 Zwicky, A., 146–50, 153–4, 157, 162 nn. 2, 4, 7

Subject Index

- A**
- absolute
 - vs. relative, *see* adjective
 - Absolute Case (principle), 256–7, 261, 277, 279
 - definiteness, *see* definiteness
 - absolutely*, 121, 124, 128, 132–5, 144 nn. 16, 19
 - accessibility relation, 231, 233–42, 246, 247 n. 10
 - adjective, 3, 5, 6, 9, 10–14, 18, 19, 25, 27–8, 33, 41–6, 48, 49, 52–4, 57, 67, 70 n. 8, 71, 73–5, 78–84, 89, 93 n. 1, 94–5 nn. 6, 9, 97, 102, 106, 116 n. 1, 117, 132–6, 143 n. 10, 144 n. 19, 145, 147–8, 151–2, 154–5, 161, 163 nn. 7, 11, 164, 170–1, 203, 247 n. 4, 251–2, 281, 290, 292, 307
 - absolute vs. relative 12–13, 19, 49, 78–9, 94 nn. 5, 6, 95 nn. 6, 7, 97, 144, 307
 - gradable, *see* gradable
 - admissible
 - comparison class, 90
 - measure function, 287, 290, 296, 298, 300
 - precisification, *see* precisification
 - adverbs
 - degree, 9, 12–15, 146–64, 195–6
 - evaluative, *see* evaluative
 - negative, *see* negative
 - sentence, 13, 15, 146–7, 149, 161–2
 - ambiguity, 1, 9, 114, 125, 186, 277 n. 17, 285
 - analog magnitude scale, 11, 28, 33, 35, 37–46
 - approximate number system (ANS), 28, 33, 37–8, 40–6
 - approximately*, 121–45, 302, 304
 - approximation, 6, 8, 35, 124–5, 129, 142, 304
 - approximators, 12–13, 118, 119, 121–45
 - arithmetic, 28, 39, 45, 47 n. 13, 201
 - assertion, 9, 14–15, 76, 81–2, 91, 96, 118, 121, 158, 171, 173, 175–8, 181, 183–4, 186, 190, 192–4, 196, 198–203, 222, 257, 292
- B**
- B-intro (rule), 242–6
 - baldness, 1, 3, 35, 52, 121, 123, 134, 140, 142 n. 1, 208, 217, 252, 286, 289
 - belief, 14–15, 142, 150, 165–89, 191–203
 - objectivized, 165–203
 - bi-truth, *see* truth
 - bivalence, 5, 16–17, 72, 207–10, 214, 220, 223–4, 227–8, 231–5, 241, 259, 277, 296
 - bi-valuation, 227, 242
 - blurriness, 2, 29–30
 - boundary, 1–5, 7, 16, 18, 25–6, 29–30, 33, 41–2, 44–5, 84, 94 n. 6, 95 n. 10, 101, 108, 116 n. 6, 135, 218–46, 272, 275, 278 n. 21, 281 ref., 286, 292–3, 296, 302
 - sharp, *see* sharp
 - borderline
 - cases, 1–3, 5, 7, 9, 14, 16–17, 25, 28–30, 33, 41–5, 72, 100, 207–20, 222–33, 239, 242–7, 258, 264, 277, 279, 290, 305
 - operator, 228, 243
 - buffer zone, 28–33, 41–3, 45
- C**
- categories (lexical), 3, 8, 52, 100, 276 n. 1, 3
 - categorization, 4, 105
 - choice function, 54, 62, 296

- clarity, 9, 14–15, 18–19, 48, 161,
165–203
adverbial vs. adjectival, 14, 167,
170–1, 179, 186, 192, 195–6,
203 n. 1
personal, 167, 169–71, 179, 184–7,
192–6
- clearly*, 1, 3, 12–16, 60–2, 106, 129,
143 n. 10, 161–2, 170–1, 174,
179, 183, 186, 188, 189, 195
- coercion, 132–3
- common ground, 91, 116, 142–3,
176–7, 191–2, 196–7, 202, 203 n.
1, 222, 271, 290
- communication, 8, 144
- comparatives
implicit vs. explicit, 10–11, 27, 51,
56–7, 61, 65–70, 74–8, 81–3, 85,
87, 89–94
morphology, 10, 75–7, 80–1, 83,
87–8, 95 n. 7, 96 n. 11
- comparison
class, 5, 6, 11, 18, 26, 51–3, 56–7,
60–3, 67, 68–71, 75, 88–92, 162,
252–6, 265–6, 270, 274, 276,
290–1, 306–7
standard of, *see* standard
- completely*, 3, 12–13, 124, 134–7,
139–40, 144 n. 16
- completion, 101–12, 115, 116 n. 5
see also precisification
- compositionality, 5, 10–11, 13, 19, 20,
27, 46, 49 ref., 53, 78, 80, 81–3,
87–9, 91, 93, 97 ref., 117 ref., 121,
141, 307 ref.
- concept
point-denoting, 130
precise, 4, 6, 46
vague, 2, 7, 19–20 ref., 79, 122–3,
281 ref.
- conditionals, 46 n. 2, 170, 188 n. 19,
189 ref., 195, 239, 244–5, 247 n.
13, 258–61, 272–3, 279–80 n. 33,
288
counterfactual, 272
- consequence, *see* logical
- consistency, 35, 61, 71 ref., 190 ref.,
292, 297–8, 300
of belief, 181
- Consistency Constraints, 75, 88–90
of gap principles, 217–18, 246
paraconsistency, *see* logic
postulate, 53
- context
change, 53, 60, 71 ref., 103, 207,
274, 281 ref., 291–2
external vs. internal, 291
set, 63, 176, 294
structure, 53–4, 56, 61, 63, 65, 67, 69
- context-dependence, 4–5, 7, 10, 18,
19 ref., 20 ref., 26, 52, 54, 75, 86,
162 n. 5, 249–52, 255, 262, 270,
275, 276 n. 4, 290–1, 295–303
- contextualism, 4, 7, 9, 17, 18, 26,
60–1, 97 ref., 249, 250–1, 255,
260, 263, 265, 270–5, 280–1, 300,
302, 307 ref.
nonindexical, 250, 270, 280–1 n. 39
- continuous vs. discrete, 28–9, 37, 84,
286
- counterpart relation, 103, 114
- crisp judgments, 29, 77–91, 94 n. 4,
95 n. 10
- cross-world identity, *see* identity
- cutoff point, 4–5, 12–13, 17, 30, 39,
52, 59, 85, 101, 104–7, 109, 111,
116 n. 6, 169, 207, 249, 251,
258–60, 262, 264–5, 269–70, 271,
273, 275, 277 nn. 14, 19, 21,
279 nn. 25, 30
- D**
- D-introduction (rule), 16, 210, 212,
215, 222–37, 242–6
- de dicto, vs. de re, 112–13
- definite description, 77–8, 82, 85, 86,
90, 95 n. 11, 103–4, 112–13
- definitely*, 12–16, 28–9, 57, 105, 108–9,
121, 123–4, 128–30, 140, 142 n.
4, 143 nn. 9, 10, 187, 209, 211,
217–19, 222–8, 230–2, 234–6,
239, 243, 245–6, 277 n. 19
- definiteness, 16, 243–4
absolute, 216–17, 235–6
- degree
adverb, *see* adverb
of belief, 14, 172, 177, 196

- modifier, *see* modifier
 theories, 25–7, 46 n. 3, 69, 80, 92
 of truth, 20 ref., 46 n. 3
 delineation, 52–4, 143 n. 10,
 162 n. 6
 demonstrative, 103, 111, 114–15,
 281 ref.
determinately, 15–16
 see also definitely
 determiner, 3, 6–7, 294–5, 307 ref.
 difference
 just noticeable, 47 n. 18
 significant, 5, 11, 18, 32, 35–7,
 39–40, 42, 58–60, 66, 74, 81,
 84–5, 105, 116 n. 2, 267–8, 280 n.
 35, 284–5, 288–9, 300, 304
 disagreement, 168–9, 185–6, 189 ref.,
 190 ref., 193, 220, 224
 discourse particle, 142 n. 5, 145 ref.
 discrimination, 5, 20 ref., 35–6, 39,
 40–1, 49 ref., 71 ref., 268–9,
 281 n. 34
 domain
 cut, 295
 of individuals, 25, 33–5, 39, 41, 53,
 64, 75, 82, 84, 91, 93, 98–9, 102,
 109, 111, 114, 149, 161, 262, 269,
 285–92, 295–304
 restriction, 18, 269, 283, 289, 292,
 294–6, 301, 303, 307 ref.
 widening, 152, 161
 Downward Difference (principle),
 54, 62
 dynamic semantics, *see* semantics
- E**
- Ekman's law, 47 n. 18
 ellipsis, 207, 270–4, 280 n. 36
 emotives, 161
 epistemic
 modals, 14, 109–10, 115, 130,
 166–7, 172, 197
 vs. scalar (vagueness), 13, 20 ref.,
 118 ref., 121–5, 128–30, 135, 140,
 142 n. 4, 143 n. 11, 145 ref.
 epistemicism, 4–5, 7, 11, 13, 20 ref.,
 21 ref., 50 ref., 70 n. 5, 71 n. 12,
 74, 83–7, 93, 95 n. 10, 106, 115,
 122, 221 ref., 258, 264, 280 n. 34,
 302
 error, 5, 34–5, 39, 45, 58, 66, 84, 86,
 95, 293, 300, 305–6
 Euclidian principle, 230–2, 247 n. 9
 evaluative adverb/predicate, 146–8,
 152–4, 158–64
 evidence, 14, 130, 166–70, 177–80,
 182–3, 191–203, 278
 exact
 boundary, 26, 29, 108, 114, 207, 227
 quantity, 35, 43, 45, 46 n. 4, 152,
 156
 science, 4
exactly, 13, 20 ref., 27, 31, 123–5,
 127–8, 131–7, 140–1, 142 nn. 4,
 5, 143 n. 12, 301–2
 exclamative, 150, 152, 161, 164 ref.
 explicit vs. implicit comparatives, *see*
 comparatives
 External Structure (principle), 256,
 266
- F**
- factive predicate, 158, 161, 181, 182,
 185
 false vs. true, 3–4, 9, 28, 31–2, 44–5,
 51, 53, 56, 59, 61, 65, 67, 68, 73,
 75, 76, 82–5, 88–90, 94 n. 4, 95 n.
 11, 101, 106–7, 110, 116 n. 2,
 133–4, 151, 153, 156, 166–7, 172,
 174, 176, 180–1, 185, 188 n. 9,
 191, 193–8, 207–8, 210–11,
 213–14, 218–19, 253–4, 257–8,
 269, 273, 277 n. 19, 280 n. 34,
 283, 288, 289, 291, 293, 300–3,
 306 n. 14
 bi-false, 227, 235
 falsity, 31, 42, 94 n. 4, 131, 133,
 220 n. 5, 224, 227, 235, 283, 285
 see also false
 first-order
 logic, *see* logic
 vagueness, *see* vagueness
 framing, 160–1
 fuzzy
 boundary, 29, 41, 47 n. 7, 94 n. 6
 logic, 4, 12, 20 ref., 43–4, 59, 122,
 144 ref.

- G**
- gap
 central, 18, 285, 288–306
 denotation/extension, 20 ref., 30,
 60–7, 89, 101, 103, 108, 110, 112,
 116 n. 6, 237, 306 n. 14
 principle, 16, 19 ref., 208, 210–12,
 216–20, 221 ref., 222–46
 truth-value, 17, 100, 106–7, 110,
 115, 207–8, 210–15, 224, 235, 302
 vs. glut, *see* glut
- Gaussian distribution, 35, 36, 43
- genericity, 174, 188 n. 8, 189 ref.,
 190 ref.
- glut, ix, 8, 19 ref.
- gradable adjective/expression, 2, 5,
 9–14, 19 ref., 20 ref., 25–8, 33–4,
 41, 43, 46, 48–9 ref., 52–3, 71–2
 ref., 75–80, 83, 87, 93, 94 nn. 5, 6,
 95 n. 7, 96 n. 12, 96–7 ref.,
 117–18 ref., 143 n. 10, 144 ref.,
 147–8, 151, 153–5, 162 nn. 3, 6,
 163 n. 11, 164 ref., 247 n. 4, 252,
 262, 279 n. 25, 281 ref., 307 ref.
- graded property, 98, 101–3, 107, 109
 belief/evidence, 172, 177
- gradual change, 29–30
- granularity, 44, 116 n. 2, 126–40,
 144 ref.
 function, 126
- H**
- halo, 20 ref., 122, 131, 143 n. 13,
 144 ref.
- heap, 2–3, 19–20 ref., 49 ref., 71 ref.,
 121–2, 129–30, 135–41, 142 n. 9,
 221 ref., 247 n. 4, 248 ref., 276 n.
 3, 281 ref., 286, 307 ref.
see also sorites
- hedges, 12, 20 ref., 144 ref.
see also modifier; approximator;
 intensifier
- higher-order vagueness, *see* vagueness
- I**
- identity, 99, 103–5, 111–12, 115, 117
 ref., 280 n. 34
- ignorance, 4–5, 11–12, 20 ref., 50 ref.,
 97 ref., 99–103, 106–11, 207,
 220 n. 2, 249, 251, 258–9, 262,
 264–5, 270–1, 275, 278 n. 21,
 279 n. 25
- ignorance of cutoffs (principle),
 249, 251, 258–9, 262, 264–5,
 270–1, 275, 278–9 nn. 21, 25
see also unknowability
- illocutionary
 entailment, 184
 strength, 171, 177–8, 190 ref.
- imprecision, 2, 6–7, 42, 47 n. 5, 142 n.
 1, 304
- indeterminacy, 12, 101–2, 109, 114
- individual (entity), 12, 80, 89, 91, 95 n.
 7, 98–115, 116 nn. 1, 3, 143 n. 10,
 155, 162 n. 6, 291, 297–302
- partial, 108
- indifference, 55, 58–9, 60
- indistinguishability, 39, 57, 59–61, 69,
 91–2, 129
see also indifference
- induction principle, 2, 45, 49 ref.,
 117 n. 7, 284
- inductive premise, 57–61, 73, 78, 84,
 90, 121–2, 283–4
see also tolerance principle
- inexactness, 4, 7–8, 19 ref.
- intensification, 8, 150, 152, 161
- intensifier, 12, 13, 119, 152
- interest-relativity, 5, 19 ref., 26, 44, 48
 ref., 75, 81–4, 87, 95 nn. 10–11,
 96 ref., 105, 117 ref., 250–75,
 279 n. 31, 280 n. 37
- Internal Structure (principle), 256,
 273–4
- interval, 6, 46 n. 4, 50 ref., 58, 64,
 69–70, 125–8, 140, 142 n. 6,
 287–8, 298
 scale, *see* scale
- irreflexivity, *see* reflexivity
- iteration (principle), 230, 232, 238–9,
 247 n. 8
- J**
- judgment, 5, 10–11, 29, 31, 38, 57,
 73, 77, 175, 187, 189, 191, 193,
 196–7, 222, 285
- justification, 14–15, 166, 169, 171,
 181–7, 191–7, 200–3

- K**
- knowledge, 14, 19 ref., 26–8, 33, 97
 - ref., 106, 145 ref., 166, 168,
 - 172–3, 181, 184, 188 n. 11, 264,
 - 269, 278 n. 24, 305 n. 2, 306 n. 11
 - common, 101
 - partial, 108, 111
 - private, 166, 192–8
- L**
- language
 - metalinguage, 25, 262
 - natural, 3, 19 ref., 43, 45, 47 n. 10,
 - 49 ref., 60, 71 ref., 83, 93, 97 ref.,
 - 99, 121, 145 ref., 164 ref., 289,
 - 303, 307 ref.
 - pro-drop, 188 n. 22
 - of science, 3–4
 - use, *see* use
 - Leibniz's law, 100
 - lexical semantics, *see* semantics
 - local
 - truth, *see* truth
 - validity, *see* validity
 - logic
 - classical, 4–5
 - first-order, 175, 189 ref., 282 ref.
 - many-valued, 4, 8, 19, 43
 - modal, 16, 19 ref., 209–10,
 - 221 ref., 229, 238–9, 244–5,
 - 247 n. 9
 - paraconsistent, 8
 - partial, 4, 20, 60, 117
 - three-valued, 57, 60, 62, 71
 - logical
 - consequence, 208–10, 221 ref.,
 - 279 n. 28, 303, 305 n. 3
 - words, 3
 - lottery paradox, *see* paradox
- M**
- markedness, 53, 73, 75, 77, 80, 87,
 - 96 n. 12, 133, 147, 149–53
 - generalization, 149–52, 160–1,
 - 169–70
 - margin for error, 5, 39, 58, 66, 84, 86,
 - 95, 293, 305 n. 2, 306 n. 11
 - mathematics, 3, 44, 48 ref., 180
 - meaning
 - underdetermination, 4, 7–8, 208
 - vs. use, *see* use
 - measure function, 33–4, 40, 44, 69,
 - 83, 88, 285–8, 298–9, 305 n. 8,
 - 307 ref.
 - measurement, 8–13, 18, 23, 27, 31,
 - 33–5, 39–41, 43–5, 47 n. 9,
 - 49–50 ref., 55, 58–9, 61, 66, 69,
 - 70 n. 6, 71–2 ref., 81–4, 86, 88,
 - 94 n. 6, 95 n. 9, 96–7 ref., 101–2,
 - 107, 109–11, 130, 133, 155–6,
 - 164 ref., 285–301, 305 nn. 7–8,
 - 307 ref.
 - mereology, 104, 144
 - mixture model, 178, 180, 182, 185–8,
 - 190 ref., 191
 - modal
 - adjective, 170–1
 - logic, *see* logic
 - operator, 14, 16, 136, 157, 209, 231,
 - 247 n. 10
 - modifier, 3, 8, 12, 21 ref., 74–5, 97
 - ref., 109, 121, 129–30, 136, 139,
 - 145 ref., 146–63
 - degree, 146–50, 153–62, 185–6, 188
 - sentence vs. predicate, 13–14, 151,
 - 162 n. 3
 - monotone
 - downward, 157–9, 163 nn. 9, 11
 - nonmonotone, 163 n. 11
 - upward, 156–8, 160, 163 n. 11
 - see also* monotonicity
 - monotonicity, 14, 146–7, 153–60,
 - 163 n. 6, 163 nn. 9, 11, 164 ref.,
 - 180, 238–9, 305 n. 6
 - multidimensional predicate, 10, 53,
 - 102, 142 n. 7
- N**
- negative
 - adverb, 14, 146, 149, 154
 - context, 124
 - polarity item (NPI), 125, 135, 158,
 - 163 n. 9
 - vs. positive extension/denotation,
 - see* positive
 - negotation, 176–8, 190, 202

- No Reversal (principle), 54, 62
 norm, 4, 14, 265–6, 270, 274, 279 n.
 31, 305 n. 9
 normal distribution, *see* Gaussian
 distribution
 normality, 154, 166–7, 173–5, 183,
 186, 187 n. 2, 188 n. 8, 191–3,
 196–7, 254, 263, 265, 289
 numerosity, 11, 28, 33, 37–46, 48 ref.
- O**
 objectivity, *see* belief
 ontology, 44, 47 n. 22, 88, 98–9,
 117–18 ref.
 open
 meaning, 7, 112
 scale, 13, 127, 132
 texture, 70 n. 7
 ordering
 interval-order, 58, 70 n. 4
 partial-order, 64, 142 n. 7
 semi-order, 11, 20 ref., 39–40, 47 n.
 15, 49 ref., 57–60, 62–70, 71 ref.
 total-order, 53
 weak-order, 11, 55, 59, 62, 66–70,
 84, 86, 285–6, 298, 305 n. 5
- P**
 paradox
 gap-principles, 223–36, 246 n. 2
 lottery, 15, 172–4, 196–7, 202, 261
 Moore's, 179
 sorites, *see* sorites
see also gap principles
 partiality, 4, 15, 20 ref., 26, 28, 60–1,
 88–9, 99, 101, 107–8, 110–11,
 113–15, 116 n. 6, 117 n. 8, 145
 ref., 307 ref.
partially, 128, 134, 139, 144
 penumbral (connection), 19 ref., 59,
 221 ref., 248 ref., 286
 perception, 4, 177
perfectly, 128–9, 136–7, 140, 144 n. 17
 POS, *see* positive morpheme
 positive
 vs. comparative form of the
 adjective, 9–11, 19 ref., 27–9,
 41–5, 46 n. 5, 47 n. 20, 49 ref., 53,
 71 ref., 74–5, 78–93, 95 nn. 8–11,
 96 n. 12, 97 ref., 162, 164 ref.,
 292, 307 ref.
 contexts, 188 n. 18
 morpheme (POS), 10–11, 81, 83,
 86–7, 95 n. 8, 97 ref., 155
 vs. negative extension/denotation,
 18, 76, 78, 82, 106, 111–13,
 116 n. 6, 117 n. 9, 211, 217,
 220 n. 12, 257, 272, 275, 290
 vs. negative form, 14, 146, 279 n. 31
 practical interest, *see* interest
 pragmatic
 halo, *see* halo
 treatment of the sorites, 9, 17–18,
 61, 63, 70 n. 9
 pragmatics
 vs. semantics, 6, 8, 20 ref., 71 n. 12,
 74–5, 77, 86, 91, 94 n. 5, 97 ref.,
 118 ref., 122, 126, 144 ref., 190
 ref., 292, 297, 300–1, 306 n. 9
precisely, 44, 122, 125, 128, 131, 134,
 136
 precision, standard of, *see* standard
 precisification, 12, 44, 87, 89, 101,
 209, 214–17, 220 n. 3, 227–42,
 246, 247 nn. 5–7, 14
 admissible, 209, 215, 217, 223,
 227–8, 231, 233, 240–1, 277 n. 19
 presupposition, 18, 20 ref., 53, 56, 75,
 78, 82, 85–7, 90–1, 95 n. 11, 107,
 116 n. 2, 118 ref., 143 n. 9, 158,
 167, 173, 181–3, 188 n. 9, 197–8,
 202, 217, 306 n. 14
 proper name, 3, 100, 103–4, 107–15,
 118 ref.
 property
 graded, *see* graded
 intrinsic vs. extrinsic, 99, 104–5
 relational, *see* intrinsic vs. extrinsic
 values, 12, 98–116
 probability, 20 ref., 172–3, 175–81,
 189–90 ref., 196, 201, 203 ref.,
 278 n. 24, 281 ref.
 subjective, 201
 psychological predicate, 162 n. 2
see also evaluative predicate
 psychology, 3–4, 17, 19–20 ref., 48–50
 ref., 98–9, 162 n. 2, 261–2, 274,

275 ref., 277 n. 18, 278 n. 24,
281–2 ref., 290, 297
psychophysics, 38, 47 n. 20, 49–50 ref.

Q

qualitative
 expression, 2
 notion, 11
quantification, 90–1, 129, 130, 153,
 167, 189–90 ref., 228, 263, 278 n.
 22, 294, 297, 307 n. 14
quantifier, 3, 6–7, 18, 20 ref.,
 50 ref., 130, 144 n. 16, 164 ref.,
 278 n. 23, 289, 294, 296,
 307 n. 14
quantity, 2, 11, 293

R

relative vs. absolute gradable
 adjective, *see* adjective
reflexivity, 55, 58–9, 62, 70 n. 4, 209,
 215, 218, 220 n. 10, 234, 235–6,
 245
regional
 truth, *see* truth
 validity, *see* validity
 valuationism, *see* regional truth
restriction, *see* domain
reversal, *see* No Reversal
rigid designation, 103, 108–11

S

salient
 Similarity (principle), *see* similarity
 standard of comparison, 5, 25, 75,
 94 n. 6, 96 n. 12, 291
scalable, 28–9, 32–3, 42–5, 47 n. 20
scalar
 adjective/concept, 75, 98, 143 n. 11
 approximator, 13, 124–40
 structure, 3, 49 ref., 80, 94 n. 6,
 144 n. 14
 vagueness, 13, 20 ref., 44, 46 n. 5,
 50 ref., 118 ref., 121–6, 142 n.
 1, n. 4, 145 ref.
 variability, 36–40
scale
 analog magnitude, *see* analog

interval scale, 69, 287–8, 298
open vs. closed, *see* open
ordinal, 34, 286–7
ratio, 28, 34–5, 39–40, 49 ref., 69,
 287
 scale interval, 125–8, 140
 of truth, 133
seamlessness, *see* transition
semantics
 compositional, *see* compositionality
 degree, *see* degree theories
 dynamic, 18 ref., 117 ref., 130, 144
 ref., 175–7, 189 ref., 202–3
 gap, 296–99, 303–4
 lexical, 121, 152
 possible world, 101–2, 129, 167,
 175–6, 209, 231–40
 probabilistic, 14, 130, 175
 region-valuation, 222–3, 233–5
 supervaluationist, *see*
 supervaluations
 see also central gaps; regional truth
sharp boundary/cutoff/transition,
 1–5, 25, 30, 32, 39, 59, 83–4,
 94 n. 6, 95 n. 10, 101, 108,
 208, 210, 211–13, 218–20,
 224–6, 228, 230, 245–6, 263,
 275, 278 n. 21, 292–3, 296, 302,
 304, 306 n. 14
shiftiness, 1, 19 ref., 26, 44, 48 ref., 71
 ref., 83–4, 87, 96 ref., 117 ref.,
 162 n. 5, 189 ref., 249, 256, 257,
 281 ref., 292, 302, 307 ref.
signal detection theory, 47 n. 20
significant difference, *see* difference
similarity, 4, 35, 39, 59, 69, 102, 105,
 267–70, 277 n. 16, 302
 Salient Similarity (principle of),
 17–18, 251, 255–76, 278 n. 22,
 279 nn. 26, 33
 Similarity Constraint, 73–4, 78,
 80–90
singular term, 12, 100, 108, 276 n. 1,
 278 n. 23, 296, 298
skepticism, 18 ref., 167–8, 189 ref.,
 192–3, 201, 203 ref.
slack, 6
 regulator, 131
social choice, 10

sorites

paradox, 2, 4, 5, 9, 15–18, 20–1 ref.,
26, 29, 45, 50 ref., 57, 59–63,
70–1 ref., 78, 84, 86, 90, 121–2,
165, 187 n. 2, 190 ref., 224–36,
267–72, 280 n. 38, 282 ref.,
283–4, 287–96, 305 nn. 3, 6,
307 ref.

series, 17, 207–20, 221 ref., 224–36,
284, 300–1, 307 n. 14

susceptibility, 2, 26, 165, 249, 251,
258–65, 270–5, 279 n. 25, 286

standard

of comparison, 5–6, 9–10, 14,
25–33, 41–2, 52, 75, 79–86, 88–9,
94 n. 5, 96 n. 11, 101–2, 105, 107,
117 n. 7, 130, 143 n. 10, 151, 155,
165–69, 187 n. 1, 188 n. 16,
192–3, 201, 255–7, 262, 269,
280–1 n. 39, 289, 291–2, 297,
300, 306 n. 9

of precision, 9, 18, 293–4, 304

Strawson entailment, 158, 164 ref.

subjective probability, *see* probability

subjectivity, 44, 167, 171, 179, 184,
189 n. 23, 191

subvaluationism, 8

super-truth, *see* truth

supervaluations, 4, 7–9, 11–12, 16,
19 ref., 30, 52, 59, 62, 87–9, 91,
100–1, 106, 108, 116 n. 6, 118
ref., 207–21, 222–47, 277 n. 19,
302

surprisingly, 3, 13, 147–61, 162 n. 5

see also evaluative adverbs

syntax, 5, 152, 291–2, 295

morphosyntax, 95 n. 8

T

threshold, 6–7, 10, 37, 40–2, 69–70,
168, 174, 176, 178, 187, 193, 201,
203, 276 n. 10, 280 n. 37

tolerance

level, 293–304

principle, 2, 4–5, 17–18, 71 ref.,
281–2 ref., 283–304

transition

gradual, 1

point, 94–5 n. 6

seamless, 2, 16

sharp, 208–13, 219, 224

transitivity, 62, 64, 70 n. 4, 215–16,
220 n. 10, 236–9, 244–5,
280 n. 34

co-transitivity, 55

nontransitivity, 7, 10–11, 19 ref.,
39, 59–60, 235, 242, 246, 268,
280 n. 34, 305 n. 5

semi-transitivity, 58

True Instance When Considered

(principle), 259–65, 270, 277 n.
19

truth-value gap, *see* gap

truth

bi-truth, 227–41, 247 nn. 5, 7, 14

global truth, 210, 213–14

local truth, 209–10, 214

regional truth, 16, 214–20, 223,
234–46

super-truth, 16, 227–34, 240–2, 246,
247 n. 7, 302

type

semantic type, 80, 88, 135, 155

type-shifting, 88, 95 n. 7, n. 8

typicality, 19 ref., 118 ref., 254, 265,
276 n. 9

typology, 13, 49 ref., 97 ref., 129, 144
ref., 160–1

U

uncertainty, 1, 5–6, 9–10, 13, 28–43,
48 ref., 74, 83, 95 n. 10, 129

underdetermination, *see* meaning

underspecificity, 3

unknowability, 17, 106, 116 n. 4,
220 n. 2, 264, 302

see also epistemicism

unsettledness, 104, 207–8

update (discourse/context), 175–7,
188 n. 11, 202–3, 291, 297,
300

Upward Difference (principle),
54, 62

use

of language, 1–7, 14–15, 18, 26, 33
vs. meaning, 1, 6, 108, 207,
306 n. 14

V

vagueness

- definition, 1–3
- epistemic vs. scalar, *see* epistemic, scalar
- first-order vagueness, 15–16, 230, 245
- second-order vagueness, *see* higher-order vagueness
- higher-order vagueness, 9, 14–16, 20 ref., 30, 32, 39, 42, 47 nn. 16, 20, 108, 116 n. 6, 129–30, 207–20, 221 ref., 222, 225, 230, 245
- semantic vs. epistemic, 7, 11, 101, 109, 207, 281 ref., 285

validity

- definition of, 233

- global, 208, 210, 212–19, 220 nn. 6, 7, 233

- local, 208–10, 213, 215, 233–4, 240

- regional, 208, 213–20, 220 n. 8, 223, 237

valuation

- bi-valuation, *see* bi-valuation

- function, 53–4

- region-valuation, *see* regional truth

- subvaluation, *see* subvaluationism

- supervaluation, *see* supervaluations

- veridical predicate, 182, 188 n. 18

- very*, 3, 12–13, 130, 148, 160–1, 165

W

- Weber's law, 36–7, 40