FernUniversität in Hagen

Objectivity, Relativism and Context Dependence

Autor:
Max Kölbel

Fachbereich Kultur- und Sozialwissenschaften
Institut für Philosophie

000 000 000 (00/00)

Table of Contents

Introduction to the Author

Max Kölbel

ICREA Research Professor at
Departament de Lògica, Història
i Filosofia de la Ciència
Universitat de Barcelona

www.icrea.cat/Web/ScientificStaff/Max-Kölbel-482
www4.ub.edu/grc_logos/max-kolbel

Born 1968 in Berlin.

1987–1989: Zivildienst (German national duty).

1989–91: Undergraduate studies (up to Zwischenprüfung) in Philosophy and Arabic at Freie Universität Berlin.

1991–1996: Postgraduate studies in Philosophy at King's College London (MA, MPhil and PhD)

1996–1998: Postdoctoral Fellow at Instituto de Investigaciones Filosóficas, Universidad Nacional Autónoma de México, Mexico City.

1999: Postdoctoral Research Fellow, Doctoral Programme in Cognitive Science, Universität Hamburg.

1999–2000: Lecturer in Philosophy (tenured), University of Wales Swansea.

2000–2001: Lecturer at the Faculty of Philosophy, Cambridge University and Fellow and Director of Studies in Philosophy, New Hall Cambridge (now: "Murray Edwards College")

2001–2009: Lecturer, Senior Lecturer (2004) and Professor of Philosophy (2007), University of Birmingham, UK.

Currently ICREA Research Professor at the University of Barcelona since 2008. Member of the LOGOS research group.

Publications (selection):

Books: *Truth without Objectivity*, London: Routledge 2002.

Relative Truth, co-edited with Manuel García-Carpintero, Oxford: Oxford University Press 2008.

Arguing about Language, co-edited with Darragh Byrne, London: Routledge 2010.

Articles: "Two Dogmas of Davidsonian Semantics", *Journal of Philosophy* 98 (December 2001), pp. 613–35.

"Faultless Disagreement". *Proceedings of the Aristotelian Society* 104 (October 2003), pp. 53–73.

"Indexical Relativism vs Genuine Relativism". *International Journal of Philosophical Studies* 12 (Oct. 2004), pp. 297–313.

"Moral Relativism", in Dag Westerstahl and Torbjörn Tännsjö (eds), Lectures on Relativism, Göteborg University 2005.

"How to Spell Out Genuine Relativism and How to Defend Indexical Relativism". *International Journal of Philosophical Studies* 15 (July 2007), p. 281–288.

"'True' as Ambiguous", *Philosophy and Phenomenological Research* 77 (September 2008), pp.359–84.

"Truth in Semantics", in *Midwest Studies in Philosophy* 32 (2008), pp. 242–57.

"The Evidence for Relativism", *Synthese* 166 (January 2009), pp. 375–95.

"Sittenvielfalt und Moralischer Relativismus", in Gerhard Ernst (ed.). *Moralischer Relativismus*, Paderborn: Mentis, 2009.

"Literal Force: a Defence of Conventional Assertion", in Sarah Sawyer (ed), *New Waves in Philosophy of Language*, Basingstoke: Palgrave Macmillan 2010.

"Vagueness as Semantic", in R. Dietz & S. Moruzzi (eds), *Cuts and Clouds: Issues in the Philosophy of Vagueness*, Oxford: Oxford University Press 2010.

"Conversational Score, Assertion and Testimony", in Herman Cappelen and Jessica Brown (eds): *New Essays on Assertion*, Oxford: Oxford University Press 2011.

Learning objectives of this text

The primary aim of this text is to provide an accessible introduction to recent debates concerning two opposing positions which are typically called “relativism” and “contextualism”. Debates in this area are clearly related to perennial philosophical questions concerning objectivity and relativism. However, these recent debates are explicitly about the correct account of the “semantic content” of certain sentences. They are, therefore, debates in a specialized field, that of natural language semantics, and the central notion, *semantic content* is a technical notion in this field.

The present text is designed to achieve the primary aim (i.e. to introduce the uninitiated to this recent debate) by pursuing three objectives: first to explain how questions of natural language semantics engage with wider philosophical questions concerning the relationship between language, thought, societies and the world, secondly to explain the technical background needed to understand the recent debate, and thirdly to explain and contribute to the current debate. Chapter 2 is mostly dedicated to the first objective, while chapters 3 and 4 are mostly dedicated to the second. Chapters 5 and 6 serve the third objective.

The aim of this text is ambitious. One important reason for this is that it is not easy to provide an accessible introduction to natural language semantics that provides all the background needed to understand the current debate. The introductory material in chapters 3 and 4 therefore differs from standard introductions in the philosophy of language. It stresses foundational issues as well as phenomena of context dependence, while leaving aside traditional controversies on which introductions usually focus, e.g. debates regarding reference, the proper treatment of names or definite descriptions, etc. An attempt has been made to provide fully articulated formal semantic descriptions of various toy languages so that the reader is enabled to check for him or herself all the claims made about formal semantics. This means that these chapters are not easy and will require concentrated study, especially by those completely unfamiliar with semantics.

Another reason why this text is ambitious is that it tries to ground its introduction to the relativism debates on some fundamental considerations about the nature, purpose and empirical status of natural language semantics. This is itself a controversial area of debate.

A third reason why this text is ambitious is that it not only tries to introduce the reader to recent debates concerning relativism, it also attempts to make progress in these debates.

As a result readers will have to work hard to master this text. On the plus side, their efforts will be repaid not only by putting them into a position to adjudicate and take part in a cutting edge philosophical debate, but also by giving them a thorough introduction to natural language semantics, which will be useful in other areas of the philosophy of language.

References

Bach, Kent (2006). Perspective on Possibilities: Contextualism, Relativism or What?". Paper presented at the Semantics & Pragmatics Workshop Sheffield, 2 December 2006.

Blais, Michel J. (1987). Epistemic Tit for Tat. *Journal of Philosophy* 84, pp. 363-75. [This item is needed for an exercise.]

Boghossian, P. (2006). *Fear of Knowledge: Against Relativism and Constructivism*. Oxford University Press, Oxford. [Accessible book defending the view that there are objective facts.]

Borg, Emma (2004). *Minimal Semantics*, Oxford: Oxford University Press. [Seminal defense of semantic minimalism.]

Burgess, John (1984). Basic Tense Logic. D. Gabbay and F. Guenthner (eds.), *Handbook of Philosophical Logic*, Vol. II, 89-133. [Standard exposition of tense logic, cited here because of foundational remarks.]

Burns, Linda (1991). *Vagueness: An Investigation into Natural Languages and the Sorites Paradox*, Dordrecht: Kluwer. [Cited here as a classic defense of the view that vague predicates are systematically ambiguous.]

Byrne, Darragh and Max Kölbel (2010). *Arguing about Language*. London: Routledge.

Cappelen, Herman and John Hawthorne (2009). *Relativism and Monadic Truth*. Oxford: Oxford University Press. [Cited because of an objection to relativism.]

Cappelen, Herman and Ernie Lepore (2004). *Insensitive Semantics*. Oxford: Blackwell. [Classic exposition of semantic minimalism.]

Carnap, Rudolf (1956). *Meaning and Necessity*. Chicago: Chicago University Press. [Pioneering work in intensional semantics.]

Carston, Robyn (2002). *Thoughts and Utterances*, Oxford: Blackwell.

Castañeda, Hector-Neri (1967). Indicators and Quasi-indicators. *American Philosophical Quarterly* 4, 85-100.

Chalmers, David (2002). On Sense and Intension. *Philosophical Perspectives* 16, 135-82.

Coady, C. A. J. (1992). *Testimony: A Philosophical Study*. Oxford: Clarendon Press. [Classic book on testimony.]

Cohen, Stewart (1986). Knowledge and Context. *The Journal of Philosophy* 83, 574-583. [One of the classic papers articulating contextualism about knowledge.]

DeRose, Keith (1991). Epistemic possibilities. *Philosophical Review* 100, 581-605. [One of the classic papers articulating contextualism about knowledge.]

DeRose, Keith (1992). Contextualism and Knowledge Attributions. *Philosophy and Phenomenological Research* 52, 913-929. [One of the classic papers articulating contextualism about knowledge.]

Dreier, James (1990). Internalism and Speaker Relativism. *Ethics* 101, 6-26. [Defence of Harman's moral relativism.]

Dreier, James (1999). Transforming Expressivism. *Nous* 33, 558-72.

Dummett, Michael (1991). *The Logical Basis of Metaphysics*. Cambridge, Mass., Harvard University Press. [Cited as the place where Dummett explains his notion of *ingredient sense* as opposed to *assertoric content*.]

Egan, Andy (2011). Relativism about Epistemic Modals. In Steven D. Hales (ed), *A Companion to Relativism*, Oxford: Blackwell.

Egan, A., J. Hawthorne and B. Weatherson 2005: "Epistemic Modals in Context", in G. Preyer and G. Peter (eds), *Contextualism in Philosophy*, Oxford: OUP. [Paper defending relativism about epistemic modals.]

Einheuser, Iris (2008). Three Forms of Truth-Relativism, in Manuel Garcia-Carpintero and Max Kölbel (eds.), *Relative Truth*, Oxford: Oxford University Press 2008. [Mentioned as a paper that defends the view that different people may inhabit different worlds.]

Evans, G. (1985). Does tense logic rest upon a mistake? In *Collected Papers*, Oxford: Oxford University Press, 343-363.

Fara, Delia Graff (2000). Shifting Sands: An Interest-Relative Theory of Vagueness. *Philosophical Topics* 28, pp. 45-81.

Field, Hartry (1973). Theory Change and the Indeterminacy of Reference. *Journal of Philosophy* 70, 462-81.

Fodor, J. (1975). *The Language of Thought*. Cambridge, MA: Harvard University Press. [Mentioned as a source for the view that Frege's view of sense is unscientific.]

Frege, Gottlob (1892). Über Sinn und Bedeutung". *Zeitschrift für Philosophie und philosophische Kritik* 100, pp. 25-50. Reprinted in Günther Patzig (ed.), *Funktion, Begriff, Bedeutung*, Göttingen: Vandenhoeck

& Ruprecht 1986. English translation by Kölbel in Byrne & Kölbel 2010.

Frege, Gottlob (1918). Der Gedanke. Eine Logische Untersuchung. *Beiträge zur Philosophie des deutschen Idealismus* 1, 2, 58-77. Reprinted in Günther Patzig (ed.), *Logische Untersuchungen*, Göttingen: Vandenhoeck & Ruprecht 1986. English translation in McGuinness 1984.

Gillies, Anthony and Kai von Fintel (2006). CIA leaks. Paper presented at Pacific APA meeting in Portland, 2006.

Glüer, Kathrin and Åsa Wikforss (2009). Against Normativity. *Mind* 118, 32-70. [Recent paper attacking the view that meaning is irreducibly normative.]

Grice, Paul (1989). *Studies in the Way of Words*. Cambridge, Mass.: Harvard University Press. [Contains most of Grice's papers on meaning, articulating, among other things, his view of the interaction between semantics and pragmatics.]

Hardwig, J. (1985). Epistemic Dependence. *Journal of Philosophy* 82, 335-349. [Needed for an exercise.]

Harman, Gilbert (1975). Moral relativism defended. *Philosophical Review* 84, 3-22. [Classic statement of a form of indexical relativism about moral language.]

Hawthorne, John (2006). *Knowledge and Lotteries*. Oxford: Oxford University Press.

Hume, D. (1966). *An enquiry concerning human understanding* (2nd ed.) (L. A. Selby-Bigge, Ed.) Oxford: Clarendon Press.

Kamp, Hans (1971). Formal Properties of "Now". *Theoria* 37, 227-274.

Kamp, Hans (1981). The Paradox of the Heap, in U. Mönnich (ed.), *Aspects of Philosophical Logic*, Dordrecht: Reidel, 225-77.

Kaplan, D. (1977). On Demonstratives, in J. Almog, J. Perry and H. Wettstein (eds), *Themes from Kaplan* (1989), Oxford: Oxford University Press, 481-563. [Kaplan's classic articulation of a semantics for a language containing indexicals.]

Keefe, Rosanna (2000). *Theories of Vagueness*, Cambridge: Cambridge University Press.

King, Jeffrey (2003). Tense, Modality, and Semantic Values. *Philosophical Perspectives* 17, 195-245. [Influential paper comparing an intensional and a quantifier view of tense.]

King, Jeffrey (2007). *The Nature and Structure of Content*, Oxford: Oxford University Press. [Recent defense of structured propositions as opposed to sets of world-like entities.]

Kölbel, M. (2002). *Truth Without Objectivity*. London: Routledge.

Kölbel, Max (2003). Faultless Disagreement. *Proceedings of the Aristotelian Society* 104, 53-73.

Kölbel, Max (2004). Indexical Relativism vs Genuine Relativism. *International Journal of Philosophical Studies* 12, 297-313.

Kölbel, Max (2007). How to Spell Out Genuine Relativism and How to Defend Indexical Relativism. *International Journal of Philosophical Studies* 15, 281-288.

Kölbel, Max (2008a). 'True' as Ambiguous. *Philosophy and Phenomenological Research* 77, 359-84.

Kölbel, Max (2008b). Motivations for Relativism, in Manuel Garcia-Carpintero and Max Kölbel (eds.), *Relative Truth*, Oxford: Oxford University Press 2008, 1-38.

Kölbel, Max (2009). The Evidence for Relativism. *Synthese* 166, 375–-95.

Kölbel, Max (2010a). Vagueness as Semantic, in R. Dietz & S. Moruzzi (eds), *Cuts and Clouds: Issues in the Philosophy of Vagueness*, Oxford: Oxford University Press, 304–26.

Kölbel, Max (2010b). Literal Force: a Defence of Conventional Assertion. In Sarah Sawyer (ed), *New Waves in Philosophy of Language*, Basingstoke: Palgrave Macmillan, 108–37.

Kompa, Nicola (2002). The Context Sensitivity of Knowledge Ascriptions. *Grazer Philosophische Studien* 64, 11-18.

Kripke, Saul (1982). *Wittgenstein on Rules and Private Language*. Cambridge, MA: Harvard University Press. [Classic book endorsing an argument purportedly from Wittgenstein, that meaning is irreducibly normative.]

Kripke, Saul (2008). Frege's Theory of Sense and Reference. *Theoria* 74, 181-218.

Kuhn, Thomas (1962). *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.

Künne, Wolfgang (1992). Hybrid Proper Names. *Mind* 101, 721-731.

Lackey, Jennifer (2006). Knowing from Testimony. *Philosophy Compass* 1/5, 432-448.

Lasersohn, Peter (2005). Context Dependence, Disagreement, and Predicates of Personal Taste. *Linguistics and Philosophy* 28, 643-86. [Influential paper outlining a non-indexical relativist semantics de taste predicates.]

Lasersohn, Peter (2008). Quantification and Perspective in Relativist Semantics. *Philosophical Perspectives* 22, 305–37. [Paper addressing some important details of the non-indexical relativist view.]

Lewis, David (1970). General Semantics. *Synthese* 22, 18–67. [Influential and classic paper in which Lewis provides an overview of the form of a semantic theory for a language.]

Lewis, David (1979a). Attitudes *De Dicto* and *De Se*". *Philosophical Review* 88, pp. 513-43. Reprinted in Lewis 1983.

Lewis, David (1980). Index, Context, and Content, in Stig Kanger and Sven Öhman (eds.), *Philosophy and Grammar*, Dordrecht: Reidel. Reprinted in Lewis 1998.

Lewis, David (1983). *Philosophical Papers*, Volume 1, Oxford: Oxford University Press.

Lewis, David (1986). *On the Plurality of Worlds*. Oxford: Blackwell. [Exposition of Lewis' view that possible worlds are concrete entities.]

Lewis, David (1996). Elusive Knowledge", *Australasian Journal of Philosophy* 74, 549-567. [Paper defending contextualism about knowledge.]

Lewis, David (1998). *Papers in Philosophical Logic*, Cambridge: Cambridge University Press.

Lipton, P. and Kusch, M. (2002) Testimony: a Primer. *Studies in History and Philosophy of Science* 33, 209-17.

Locke, J. (1975). *An essay concerning human understanding* (P. H. Nidditch, Ed.). Oxford: Oxford University Press.

MacFarlane, John (2003). Future Contingents and Relative Truth. *Philosophical Quarterly* 53, 321-36. [Classic paper in which MacFarlane motivates a form of relativism by means of the view that the future is genuinely open.]

MacFarlane, John (2005a). Making sense of relative truth. *Proceedings of the Aristotelian Society, 105*, 321-339.

MacFarlane, John (2005b) The Assessment Sensitivity of Knowledge Attributions. In Tamar Szabo Gendler and John Hawthorne (eds.), *Oxford Studies in Epistemology* 1, Oxford: Oxford University Press, pp.197-233.

MacFarlane, John (2007). Relativism and Disagreement. *Philosophical Studies* 132, 17–31.

MacFarlane, John (2008). Truth in the Garden of Forking Paths, in Max Kölbel and Manuel García-Carpintero (eds.), *Relative Truth*, Oxford: Oxford University Press.

MacFarlane, John (2009). Nonindexical contextualism. *Synthese*, 166, 231-50.

MacFarlane, John (2011). What is Assertion? in Herman Cappelen and Jessica Brown (eds.), *Assertion*, Oxford: Oxford University Press.

MacFarlane, John (forthcoming a). Epistemic Modals are Assessment-Sensitive. In Brian Weatherson and Andy Egan (eds.), *Epistemic Modals*, Oxford: Oxford University Press.

MacFarlane, John (forthcoming b). *Assessment sensitivity: relative truth and its applications*.

Mackie, J.L. (1977). *Ethics: Inventing Right and Wrong*. Harmondsworth: Penguin.

Margolis, Eric, Laurence, Stephen (2008). Concepts. *The Stanford Encyclopedia of Philosophy (Fall 2008 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2008/entries/concepts/>.

Markman, Ellen M. (1994). Constraints on word meaning in early language acquisition. *Lingua* 92, 199-227.

Markman, Ellen M., and Jeanne E. Hutchinson (1984). Children's sensitivity to constraints on word meaning: Taxonomic versus thematic relations. *Cognitive Psychology* 16, 1-27.

Matthews, Robert (1994). The Measure of Mind". *Mind* 103, 131-46. [Influential paper comparing the role of propositions, i.e. abstract entities, in the description of minds to that of numbers on scales in the measurement of quantities.]

Matthews, Robert (2007). *The Measure of Mind*, Oxford: Oxford University Press. [Booklength treatment of the same topic.]

McGuinness, Brian (1984). Frege: Collected Papers on Mathematics, Logic and Philosophy, Oxford: Basil Blackwell.

Menzies, Peter (2007). Causation in Context. In Huw Price & Richard Corry (eds.), *Causation, Physics, and the Constitution of Reality: Russell's Republic Revisited*. Oxford University Press.

Neale, Stephen (1990). *Descriptions*. Cambridge, Mass.: MIT-Press. [Excellent exposition of many issues relating to Russell's theory of definite descriptions.]

Ninan, Dilip (2010). *De Se* Attitudes: Ascription and Communication. *Philosophy Compass* 5, 551-567.

Pagin, Peter (2010). Central Gap Domain Restriction, in R. Dietz & S. Moruzzi (eds), *Cuts and Clouds: Issues in the Philosophy of Vagueness*, Oxford: Oxford University Press.

Pagin, Peter and Dåg Westerståhl (2010a). Compositionality: Definitions and Variants. *Philosophy Compass* 5, 250-64.

Pagin, Peter and Dåg Westerståhl (2010b). Compositionality: Arguments and Problems. *Philosophy Compass* 5, 265-82.

Perry, John (1977). Frege on Demonstratives. *Philosophical Review* 86, 474-497.

Perry, John (1979). The Problem of the Essential Indexical. *Nous* 13, 3-21.

Perry, John (1997). Indexicals and Demonstratives. In Bob Hale and Crispin Wright (eds), *Companion to the Philosophy of Language*, Oxford: Blackwell Publishers.

Predelli, Stefano (1998). I Am Not Here Now. *Analysis* 58, 107-115.

Predelli, Stefano (2005). *Context: Meaning, Truth and the Use of Language*, Oxford: OUP.

Price, Huw (1983) 'Does "probably" modify Sense?'. *Australasian Journal of Philosophy* 61: 396-408.

Prior, Arthur (1967). *Past, Present and Future*, Oxford: Clarendon Press.

Raffman, Diana (1994). Vagueness without Paradox. *Philosophical Review* 103, 41-74.

Raffman, Diana (1996). Vagueness and Context Relativity. *Philosophical Studies* 81, 175-92.

Raffman, Diana (2005). How to Understand Contextualism about Vagueness: Reply to Stanley. *Analysis* 65, 244-48.

Recanati, François (2004). *Literal Meaning*. Cambridge: Cambridge University Press.

Recanati, François (2007). *Perspectival Thought*. Oxford: Oxford University Press.

Reid, T. (1764). *Inquiry into the human mind and the principles of common sense*. There are many editions. The best one is by Derek Brookes,

Edinburgh: Edinburgh University Press. Standard editions are those by Hamilton (1895) and Duggan (1970). [Needed for an exercise.]

Russell, Bertrand (1905) On Denoting. *Mind* 14, 479-93.

Russell, Bertrand (1918/9). *Introduction to Mathematical Philosophy*. London: Allen and Unwin.

Rysiew, Patrick (2007). Epistemic Contextualism. *The Stanford Encyclopedia of Philosophy* (Fall 2007 Edition), Edward N. Zalta (ed.), forthcoming URL = <http://plato.stanford.edu/archives/fall2007/entries/contextualism-epistemology/>.

Schlenker, Philippe (2003). A Plea for Monsters. *Linguistics and Philosophy* 26, 29-120.

Searle, John (1969). *Speech Acts*. Cambridge: Cambridge University Press.

Shapiro, Stewart (2003). Vagueness and Conversation, in J.C. Beall (ed.), *Liars and Heaps*, Oxford: Oxford University Press.

Smith, Michael (1995). *The Moral Problem*. Oxford: Blackwell.

Soames, Scott (1999). *Understanding Truth*, Oxford: Oxford University Press.

Soames, Scott (2002). Replies. *Philosophy and Phenomenological Research* 62, 429-52.

Sorensen, Roy (1988). *Blindspots*, Oxford: Clarendon.

Stalnaker, Robert (1976). Possible Worlds. *Noûs* 10, 65-75.

Stalnaker, Robert (1978). Assertion. In P. Cole (ed), *Syntax and Semantics*, ix. *Pragmatics*, 315–22. Reprinted in Stalnaker 1999.

Stalnaker, Robert (1984). *Inquiry*. Cambridge, MA: MIT-Press.

Stalnaker, Robert (1987). Semantics for Belief. *Philosophical Topics* 15. Reprinted in Stalnaker 1999.

Stalnaker, Robert (1999). *Context and Content*. Oxford Oxford University Press.

Stanley, Jason (1997). “Names and Rigid Designation”. In Bob Hale and Crispin Wright (eds.), *A Companion to the Philosophy of Language*, Oxford: Blackwell, 555–85. [Cited because of its discussion of the notion of *ingredient sense*.]

Stanley, Jason (2000). Context and Logical Form". *Linguistics & Philosophy* 23: 391-434. [Important paper denying the existence of "unarticulated constituents.]

Stanley, Jason (2003). Context, Interest-Relativity and the Sorites. *Analysis* 63, 269-80.

Stanley, Jason (2005). Semantics in Context. In G. Preyer & G. Peter (eds.), *Contextualism in Philosophy: Knowledge, Meaning, and Truth*, Oxford: Oxford University Press, 221-53.

Stanley, Jason (2006). *Knowledge and Practical Interests*. Oxford: Oxford University Press.

Stanley, Jason and Zoltan Szabó (2000). On Quantifier Domain Restriction. *Mind and Language* 15, 219-261.

Travis, Charles (1985). On What Is Strictly Speaking True. *Canadian Journal of Philosophy* 15, 187-229.

Travis, Charles (1997). Pragmatics. In Bob Hale and Crispin Wright (eds), *A Companion to the Philosophy of Language*, Oxford: Blackwell.

Tugendhat, Ernst (1970). The meaning of 'Bedeutung' in Frege. *Analysis* 30, 177-189.

Welbourne, M. (2001). *Knowledge*. Chesham: Acumen.

Williamson, Timothy (1994). *Vagueness*, London: Routledge. [Classic book on vagueness, defending the view according to which vagueness is a form of ignorance.]

Williamson, Timothy (1996). Knowing and Asserting. *Philosophical Review* 105, 489-523. This article reappears in revised form as chapter 11 of Williamson 2000.

Williamson, Timothy (2000). *Knowledge and its Limits*, Oxford: Oxford University Press.

Wright, Crispin (2001). On being in a quandary: Relativism, vagueness, logical revisionism. *Mind*, 110, 45-98.

Wright, Crispin (2003). Vagueness: A Fifth Column Approach, in JC Beall (ed), *Liars and Heaps*, Oxford: Clarendon Press.

Wright, Crispin (2008). Relativism about Truth Itself, in Manuel García-Carpintero and Max Kölbel (eds.), *Relative Truth*, Oxford: Oxford University Press.

Yourgrau, Palle (1983). Knowledge and Relevant Alternatives. *Synthese* 55, 175-190.

1. Introduction

On November 22nd 1963, John F. Kennedy was shot and killed. Was it Lee Harvey Oswald who shot him? Many believe that it was, and many believe it wasn't. If those who believe that Oswald did it are right, then those who think he didn't can't be right. And conversely, if those who think Oswald didn't do it are right, then those who think he did can't be. In other words: it's an *objective* question whether Oswald shot Kennedy. We may not know who is right, but we immediately recognize that, if one is right, the other can't be. In this sense it is an objective matter whether Oswald shot Kennedy. We share the same objective world and our beliefs are answerable for their correctness to that shared world. The world can be such that Oswald did shoot Kennedy or such that he didn't, but not both. We recognize this immediately and without further investigation because it is part of our competence as thinkers to recognize the objective status of a matter like this.

The issue of whether a question is objective should not be confused with the question of whether there is any good or conclusive way of establishing the answer, or whether anyone actually knows the answer. There may well be objective questions to which no-one knows the answer, or the answer of which cannot be established conclusively. What I mean, when I say that a question is objective, is just that it is a priori (i.e. can be known merely on the basis of conceptual competence) that if one person answers "yes", and this answer is correct, then anyone who answers "no" is wrong, and conversely: if anyone correctly answers "no", then anyone who answers "yes" is wrong). Whether anyone has good evidence for a given answer, is justified in giving a certain answer, or whether anyone knows the answer, is a separate question. Perhaps no-one ever knew or will ever know whether Oswald shot Kennedy (perhaps not even Oswald himself, if he suffered from amnesia or some other cognitive impairment at the time). Nevertheless, we are convinced that if anyone correctly believes that he did it, then anyone who believes that he didn't must be wrong.[1]

There are countless objective questions. For example the question whether Kennedy was shot, the question whether I left the lights on before I departed on my vacation, whether you are at this moment travelling on a train, whether Miró was born in Barcelona, or whether the first human in-

[1] How can we be so sure of this? – We have been trained to use the concepts in question in this way, i.e. to put down any divergence of view on matters such as this (who shoots a gun when at whom etc) to some kind of mistake.

habitants of America came from Asia[2]. Examples abound. Objectivity seems to be an everyday phenomenon.

Nevertheless, quite a few people are sceptical of the idea of objectivity. Why? Perhaps the sceptics are impressed by widespread disagreements on certain subject matters, disagreements that are persistent and cannot easily be resolved by recourse to commonly accepted evidence. If there is widespread disagreement on some question, and the question is objective in the way described above, then there must be a large number of people whose beliefs are wrong. This means that those who believe the answer is "yes" will have to believe also that those who answer "no" are wrong, and vice versa. This may seem presumptious or feel uncomfortable. For example, consider disagreements on whether Olafur Eliasson is a good artist, or whether Barbie dolls are suitable toys for young children. In each case it might seem rash and presumptious to say that at least one side to the dispute must be wrong. Who is to say which side is wrong? In what does their mistake consist? And if there is such a mistake, why does the disagreement persist? Consider the question whether it can be morally permissible not to intervene when someone is attempting suicide, or whether German colonial ambitions were an underlying cause of WW1. Again, it seems unclear how the correct answer to these questions should be determined. So why say that one party is making a mistake? If we give up the idea that these questions are objective then we no longer need to say that at least one party to the dispute is in the wrong.

Different viewpoints, especially in politics, aesthetics, ethics or history seem not to leave room for objective standards. There seems to be no reason to assume that there is only one correct, objective standpoint. Our views seem to be the product of causal historical factors that influence and bias us, and these factors often do not seem to track an independently existing reality. Moreover, we can't make ourselves free from these irrational influences. Our views seem coloured and biased from the start. So assuming that those disagreeing with us in these matters must be wrong would seem to be unwarranted.

In addition to questions of aesthetic or moral value, or of history, even science itself gives rise to doubts about objectivity. Experience shows that scientific experts can disagree (synchronically and diachronically), and Kuhn has claimed that the transition from one scientific paradigm to the next is not governed by rational considerations (Kuhn 1962). Thus it seems equally presumptious to say that scientific questions are objective.

2 This example is taken from Boghossian 2006, which is an excellent discussion of contemporary scepticism about objectivity and knowledge.

Who is to say that in a scientific dispute at most one party can be right? Perhaps our theoretical beliefs depend so profoundly on accidental presuppositions that it would be rash to say that one framework must be wrong.

Robert Nozick (2001, p. 22) suggests a further explanation for scepticism about objectivity. Objective facts can be seen as constraining our freedom. For example, if it is an objective fact that lack of sleep causes tiredness in humans, then this limits our possibilities. It means that we cannot stay up indefinitely without getting tired. So Nozick conjectures that reluctance to accept objective facts may be explained by a desire that the facts were different. In this case, those who would like to stay up indefinitely without getting tired, might have a desire that it wasn't an objective fact that sleep deprivation caused tiredness. The flip side of this explanation, of course, is that objective facts can be seen to enable us to achieve our goals just as much as they can be seen to constrain us. Thus, the very same fact that lack of sleep causes tiredness may be felt as a relief by those who would like to get tired, or those who would like someone else to get tired at a certain point.

But I believe that Nozick's diagnosis points in the direction of a better motivation for scepticism about objectivity, one also mentioned by Nozick. This motivation, however, depends on a confusion between objectivity and unalterability. Many facts are the product of human activity and human social interactions. Thus, to take a banal example, the fact that most Spaniards have dinner later than 8.30 pm is an objective fact, but one that depends in many ways on the choices made by many people. If they chose differently then it wouldn't be a fact. To take a more momentous example, it is an objective fact that in most societies, people can acquire property rights over certain goods, a fact that has certain consequences for the options available to individuals, for example that they cannot just take an apple from a fruit stand without paying or having to deal with the threat of sanction. Such facts are "socially constructed" in the simple sense that they depend on a certain customs and social institutions, which ultimately depend on the choices made by individuals. We may very well lament some of these socially constructed facts, i.e. facts about how the individuals in some society interact. If we want to change the situation, we may very well come up against the prejudice that these facts are simply given and unalterable. Thus, someone who wants to abandon the institution of property, before they will even be able to make their case that abandonment is desirable, will come up against the view that property rights and duties are simply "an objective fact", that "that's just the way the world is". What this conservative view comes to is probably just that the institution of property is inevitable, and that it could not (or not feasibly) be abandoned by individuals making different choices. This is what the oppo-

nent of property will need to argue against. She will argue that property is socially constructed and is therefore not a fact that we just have to accept. Thus, what she needs to argue against is *not* that there is currently an institution of property and that this is an objective fact. What she needs to argue is simply that it is not necessary (and in a second step: that it is not desirable) that this institution should continue.

Nozick may well be right that it is opposition to established views or practices that often motivates scepticism about objectivity and objective facts in general. But as I pointed out, denying the existence of objective facts altogether is not needed to make the point that some practice depends on the decisions made by members of a society.

Exercise 1: *Read chapter 3 of Boghossian's Fear of Knowledge (2006). Essay question: "What does Boghossian mean by 'fact-constructivism', and do his 'three problems' show that fact constructivism is wrong?"*

Despite a certain sceptical tendencies, the vast majority of people will readily concede that at least some things are objective—who for example would deny that it is an objective matter whether Oswald shot Kennedy? Where the boundary between the objective and the non-objective lies, by contrast, is controversial. Thus, even those who admit in principle that there are some objective questions may diverge in their views as to whether, for example, moral questions are objective. Similarly, they may diverge on whether it is an objective matter whether some piece of music is aesthetically more valuable than another, or whether some dish is tasty. Whatever these controversial areas are, most people will agree that there is an unspectacular range of truths that are objective. At least this is so in our everyday experience. We assume, and our competence seems to require, that we regard the question of whether the Butler did it, or whether the lights are still on, or whether there are three chairs in the room, as objective. Even those with extreme metaphysical views, that reality is a construct etc, will still treat a range of matters as objective in that they will just assume, in an a priori manner, without empirical evidence, that either the butler did it or he didn't, and that if one person thinks that he did it, and another thinks that he didn't, then one of them is wrong. All except philosophical extremists will admit that it is an objective matter whether the lights are still on, or whether there are three chairs in the room, and even that in many cases we can easily establish the correct answer.

This does not mean that we shouldn't take these extremists seriously in philosophy. Their arguments may well be worth considering. Parmenides argued that there could only exist one thing, on the grounds that any other view was incoherent. His argument is one worth examining in detail. There is a place for these considerations. However, this text is not such a place.

In this text, we shall be starting from the assumption that there are several things. We shall be assuming that there are many things, that there are many objective truths concerning them, and that we can have knowledge of many of those. We will also assume that there may well be things we say and think that are not capable of objective truth, which may come as a relief for those who want to say that there is no such thing as objective history, objective facts about matters of taste, or objective moral requirements etc. We will be concerned with an account of our language and our thought that allows for both objective and non-objective areas of enquiry.

In forming views about the world, we are not alone. We rely heavily on what others tell us. This cognitive interdependence is deep. We acquire most of our views about the world from others. But it goes beyond merely receiving information from others via linguistic communication. For we also acquire our conceptual tools from others and hone them through mutual exchange of views. These conceptual tools are part of a human legacy that our predecessors pass on to us, that develops in our hands, and that passes on to new generations.

In pooling our cognitive resources, objectivity is a helpful assumption. If we know that a question is objective then we know that if another thinker correctly answers the question then it will also be correct for us to answer the question in this way. We can make use of others' answers, thus saving ourselves the efforts they made in arriving at their answer– assuming, of course, that their efforts were well-directed, that they used good methods and made no mistakes in doing so. Conversely, the assumption of objectivity puts constraints on the methods regarded as appropriate for forming beliefs: the correct methods must be such that anyone employing them correctly will arrive at the same answer, i.e. that arriving at divergent answers is a symptom of some mistake which will prompt a search for the location of the mistake so that it can be avoided. If, despite a divergence of answers, no mistake in the application of the method can be found, then this motivates a revision of the method, or perhaps motivates us to re-think the status of the subject matter as objective. Such corrective interplay helps hone our conceptual apparatus.

If this picture of our cognitive interaction with others and the world is right, then it makes sense for our conceptual and linguistic tools to allow objective as well as non-objective subject matters. Language is the primary medium with which we conduct our cognitive interactions. We use language to convey information, language learning is the principal conduit by which we mutually calibrate our conceptual repertoires. So an account of linguistic communication that makes room for communication about objective as well as non-objective matters is needed. This text attempts to outline an account of linguistic meaning that meets this requirement. That is, it tries

to show how a semantic theory of a natural language (i.e. a theory that describes the meanings of the expressions of that language) can make room for objectivity as well as non-objectivity.

There is a certain general theoretical framework or paradigm within which much semantic theorising has been and is being carried out. The notion of a “proposition”, “propositional content” or “semantic content” plays a central role in this framework: the main meaning property of a sentence type[3] consists in it expressing such a semantic content. These propositions or semantic contents are bearers of truth-values, and one central idea of the framework is that the proposition expressed by (an utterance of) a sentence allows us to make certain predictions about the correctness or incorrectness of the utterance. Some versions of the framework hope to do without propositions or contents, so they try to arrive at these predictions by merely specifying the conditions under which (an utterance of) the sentence is true. However, I shall mostly be addressing the standard version which does allow propositional contents. Usually, propositions are thought to function not just as the contents of sentences or utterances, but also as the contents of thoughts and linguistic acts. Thus, the content of an utterance of the sentence is also the potential content of thought, for example the content of a belief or of a desire; as well as the content of an assertion or of a command. For example, if in uttering the sentence “Sam smokes.” I express a proposition, the proposition that Sam smokes, I might also be asserting that proposition and – if my assertion is sincere – express a belief with that content.

As it turns out, the way a semantic theory deals with communication concerning non-objective questions depends on the kind of propositions it postulates as the contents of utterances. Much of this text is devoted to teasing out various different ways in which non-objective discourse can or should be treated, and what role various different notions of propositional content would play in such an account.

The next chapter, therefore, deals in a general way with the idea of propositional contents as abstract entities that can be employed to characterize language and thought. It tries to justify this way of theorizing. Chapter 3 then explains the rationale behind the typical form most semantic theorising takes, namely the form of a definition of a semantic truth-predicate, and rehearses some of the considerations that have led theorists from an extensional to an intensional version of this approach. Many sentences of natural languages depend for their correctness on the context in which they are used. Chapter 4 explains a standard framework for incorporating

3 I.e. a string of repeatable signs of which many instances or tokens can be produced.

context dependence into the semantics of a language, namely a framework roughly along the lines of Kaplan's "Demonstratives" (1977).

Up to that point, the material expounded is fairly well-known, even though I am perhaps putting a different stress on certain matters. In Chapter 5, we finally move into an area of recent controversy, namely the proper treatment of certain types of context-sensitivity. I shall be advocating the coherence of what has recently been called a "relativist semantics", i.e. a semantics that postulates propositional contents whose truth-value varies with non-standard parameters. However, I shall also be considering competing approaches, of which some will stand up well to my critical examination.

Chapter 6 enters into a discussion of the coherence of "relativist" semantics that takes its starting point from Gareth Evans' critique of tense logic (Evans 1985). This will also be the place where I consider whether forms of relativism that have been called "radical" can be motivated and are coherent.

Finally, in chapter 7, I draw some conclusions from the forgoing discussions.

2. Sharing a World

2.1 Introduction

The Basic Picture

Like most people, I like to think that I am not alone in the world. There are other people as well, and they are part of the very same world of which I am part. These other people occasionally notice the very same things I notice about our shared environment. Often they tell me things about it, or I tell them. Sometimes I believe what they say, sometimes they believe what I say, but not always. My beliefs, and others' beliefs, about the world can be more or less successful. In particular, they can be true or fail to be true. Another dimension of success is justification: a belief can be true by luck or because the believer was justified. Being justified in believing something involves having good evidence or reasons for the belief.

This basic picture of our position in the world is a very common, probably held by most of the readers of this text. To be sure, certain radical thinkers will doubt aspects of the basic picture. Sceptics of different sorts doubt, respectively, the existence of the external world, the existence of other people, or the possibility of having justified beliefs or knowledge about the world. These radical thinkers may have good reasons or arguments for their views, reasons and arguments well worth our attention and discussion. However, radical scepticism is not our present topic, so I shall just assume here that the basic picture is correct: we are not alone, and we share the same reality about which we can have more or less successful beliefs, with truth and justification being two important dimensions of attainable success.

Even with this basic picture in place there is plenty of room for philosophical debate and controversy. For example, the extent and nature of the shared world is debatable: is all of the shared world ultimately material or natural? Or does the shared world contain immaterial or non-natural things or facts? Are there normative facts, e.g. facts as to what is good, bad, what one ought or ought not to do? If so: are these normative facts material or natural? How do we access these facts? This text will not address all these questions, but it will develop a framework in which thought and linguistic communication about the shared world can be described and explained.

Testimony

Let us return to the basic picture: we – that is you, the readers, I, the author, and everybody else in the world – share the same reality, and we can have more or less successful beliefs about that reality, beliefs that we sometimes share. Now, two thinkers can coincide in beliefs about the

world not only because they arrive at the same belief independently, but also because they frequently believe things others have told them. You and I might both notice independently that the corner shop is closed for inventory, or one of us might notice and tell the other. We face the world together. This is true in more than one sense: first, it's the same world about which we all have views. Secondly, our views about the world are collaborative in countless ways. People rely on what they are told by others. An informant perhaps relies on sense-perception in arriving at the view they later pass on to others who are not in a position to rely on sense-perception in this matter. In using sense-perception, the informant may have had to rely on information they in turn received from their informants. For example, if they use sense-perception to find out that the corner shop is closed for inventory, it may be that they need to rely also on information they earlier received from others – such as the information they got from the sign "closed for inventory" on the shop's door, or the information needed to interpret such a sign. This second way in which we face the world together, i.e. the fact that we derive our information about the world from one another, is one I would like to explore a little more here. What others tell us can be a source of knowledge about the world. In the philosophical literature this potential source of knowledge is often called "testimony"– a rather grandiose word for an everyday occurrence.

The extent of reliance on testimony

Our reliance on the testimony of others in forming beliefs about the world is quite extensive, and by no means restricted to beliefs about essentially institutional or social facts such as the opening hours of shops. Consider, another example. I, like most people, have views about who my (biological) parents are. This is not an institutional matter. What is my justification for these views? It's testimony: others, including my parents, have told me who my parents are. Sure, I could improve my justification by drawing on additional evidence, and in certain situations such further evidence may even be demanded of me (e.g. when I apply for a passport). I might rely on a birth certificate. But again, that evidence involves the testimony of others, though in this case the testimony is given in a certain legally prescribed form by people who, in virtue of their special status (registrar), are deemed especially reliable witnesses. But these people, again, will have relied on the testimony of others, such as a physician, midwife or the parents themselves. I could improve my evidence further: I might rely on the testimony of a forensic scientist, who in turn is in the possession of DNA evidence. Suppose this expert has herself taken DNA samples from my parents and from myself, and has herself carried out the DNA analysis of these samples. Would the forensic scientist's view as to whether these are my biological parents be justified independently of the testimony of others? Arguably not: even the forensic scientist will rely on the testimony of others in assessing the significance of the DNA evidence. In order to have evidence as to whether I am the biological son of a certain couple, evidence

that is completely independent of the word of others, the believer would not only need to be a forensic scientist who takes the samples and carries out the analysis herself. This scientist would also have to have carried out by herself experiments that provide evidence for the scientific theory on which the DNA test relies. It is hard to see how there could be such a scientist. For empirical science is in large part a collaborative activity.[4]

***Exercise 2**: Read Hardwig 1985 and Blais 1987. Essay question: What is Hardwig's conclusion, and has he offered convincing arguments for it? Is the objection by Blais successful?*

A moment's reflection will reveal the enormous extent to which the beliefs of individual thinkers depend for their justification on communication with other thinkers. To be sure sense perception, memory and reasoning play a crucial role in the acquisition and justification of belief and knowledge. But if an individual thinker had to rely only on his or her own senses, memory and reasoning, he or she would not get very far. As individual thinkers by ourselves, we can justify only beliefs that are immediately justified by our own observations, in so far as we remember them, and by our own powers of reasoning. This is precious little compared to the wealth of information we all rely on every day. I have many beliefs about places I have never been to, based on the testimony of others. I acquire further beliefs about how I can get to these places, if I want to, again by relying on the testimony of others. When I buy food in the supermarket, I do not just rely on others to produce the food that I want to consume, I also rely on others in forming my beliefs on what types of food will nourish me.[5] My belief that bread is made from wheat or rye, that wheat and rye will nourish me, etc are backed up by the experiences of others. So is my belief that certain fruits, vegetables, roots, or fungi are edible, that certain drugs cure certain diseases, that there is a railway to Siberia and that humans and other animals are mortal.

Human societies derive great benefits from a division of labour in the production of goods. Readers will be familiar with the idea of economic spe-

---

4 There are those who believe that the justificatory dependence of individual thinkers on others is so great that we should abandon the idea of individual knowing subjects and instead regard entire communities of thinkers as the appropriate subjects of knowledge, a view sometimes called "communitarianism" (e.g. Hardwig 1985, Welbourne 2001). There is a good account of more recent debates on testimony in Lipton and Kusch 2001 and Lackey 2006. Recent debates about the nature of testimony go back to Reid 1764, ch. 6, Hume 1966, section 10, Locke 1975.

5 I might, to a certain extent, experiment with various foods available to me. I might also on occasion experiment with different methods for preparing the food. To this minimal extent, my shopping choices are supported by beliefs I can justify by my own sense experience and my own reasoning.

cialisation, and may have views on the way their society divides labours and distributes goods. But the epistemic division of labour that I am talking about is independent of the type of economic specialisation that any reader of this text will be familiar with. Epistemic division of labour concerns even societies that are not specialized in the production of goods in the way that ours are. For any human society will use language to communicate, and linguistic communication will involve at least one thing: the transmission of information.[6] Humans pool their epistemic resources so that the sphere of knowledge and belief of individuals is extended far beyond the reach of their individual powers of observation, memory and reasoning. This sharing is both synchronic between individuals who live at the same time, and diachronic, namely when later generations rely on the beliefs acquired by earlier generations. This sharing of information is to a large extent independent of economic specialisation. An agricultural society, or a hunter-gatherer society in which individuals (or perhaps family units) produce by themselves virtually all the goods they use will still be a society in which the individuals rely for their beliefs on what they can learn from others. In fact, such a society requires *greater* epistemic dependence simply because individuals require a wider range of knowledge and know-how in order to be able to master this way of life. The range of knowledge that needs to be passed from old to young will increase, perhaps even because the level of economic specialization is lower.

Testimony requires co-ordinated conceptual capacities

Now, individual thinkers can pool their epistemic resources effectively only if they have at their disposal a means of communication and if their faculties of belief formation and reasoning are sufficiently co-ordinated. On a simple model, sharing of epistemic resources consists in some individuals *expressing* some of their beliefs by means of language and other individuals as a result acquiring beliefs with the same content. Thus, the recipient of the information needs to be able to form a belief with the same content as the belief expressed by the source. This may be an overly simple model. However, it makes clear that the belief forming faculties of informant and recipient need to be suitably related. Thus, for example, if someone tells me that the shop is closed, in order for the transmission of information to succeed, I need to be able to form the belief that the shop is closed, which involves having the concept of a shop and of being closed. If I lack the concepts needed to form a belief of the right sort, I cannot benefit epistemically.

6 Another essential function of human language is that of co-ordinating action in a group, such as when one individual tells another what to do. This may or may not be a special case of the information-sharing function.

On a more complex model of success in communication, the recipient may perhaps acquire a belief that is different in some respects from the belief expressed by the source. However, even on such a more complex model, the belief acquired by the recipient cannot be *arbitrarily* related to the belief expressed by the source, but rather must be non-accidentally related to it, for otherwise it would be difficult to explain how this process could serve for the transmission of information, i.e. how the recipient's new belief could non-accidentally inherit some of the merits and benefits of the belief expressed by the source. Thus, in any case, the recipient of testimony will need to have the capacity to acquire beliefs of the right kind if she is to benefit from the reception of testimony. For example, if someone tells me "*your* pants are on fire.", successful communication might require that I form the related belief that *my* pants are on fire.[7] On some views, in this case the belief formed by the recipient of the information is different from the belief expressed by the provider of the information: the provider believes something they would express by saying "your pants are on fire" while the recipient acquires a belief they would express by saying "my pants are on fire". Even if we want to say that these are *different* beliefs, or that their contents differ, we will still want to maintain that they are non-accidentally related and that they require at least in part similar capacities for thought. Intuitively, at least the concept of *pants* and of *being on fire* employed by the two sides need to be similar – whatever detailed account of concepts we adopt. Moreover, it is not a coincidence that the concept *yours* employed by the source should be mirrored by the concept *mine* employed in the recipient's newly acquired belief. The belief expressed must be co-ordinated with the recipient's newly formed belief.

2.2 Propositions and Concepts as the Common Property of Many

When we share our epistemic resources, what exactly is it that we share? When we tell others how things are or they tell us, how should we describe these interactions? The basic picture I described allows us to assume, at the very least, that individuals have beliefs, they say something to some audience, as a result the audience forms a belief, and this resultant belief can inherit some of the virtues (e.g. justification, truth) from the source's belief. It is clear that in order for any such episode to be a case of successful pooling of epistemic resources, certain benefits or advantages of the source's beliefs must non-accidentally have been passed on to the recipient and have done so by bringing about a new belief in the recipient.

7 See for example Castañeda 1967, Kaplan 1977, p. 533, Perry 1979, Lewis 1979. A recent overview is provided in Ninan 2010.

For example, suppose that there is a bear near the clearing and that Anna believes that there is a bear near the clearing. This belief will have obvious benefits and advantages for Anna, for example she'll be able to avoid or seek an encounter with the bear deliberately, depending on what she is after. In this situation Anna can, if the circumstances are right, make these advantages available to someone else, say Ben, by telling Ben that there is a bear near the clearing and thereby getting him to believe that there is a bear near the clearing.

In the philosophy of language and related disciplines it is common to theorize about beliefs and certain other mental states, as well as about linguistic acts, by describing them in terms of their *propositional content*, i.e. to treat them as *propositional attitudes* and *propositional acts*. Mental states are described as having a certain proposition as their content. Linguistic acts are described as *semantically expressing* propositions. But what is a proposition? One of the basic features of propositions is that they are publicly accessible, objective entities, rather than private or subjective. Frege, a pioneer of the propositional framework, emphasized this right from the start. He thought that linguistic expressions express "Senses", and in the special case where the expression is a sentence, the sense expressed is called a "Thought", i.e. what would nowadays standardly be called a "proposition" or a "content". Frege made quite clear that by "Senses" and "Thoughts", he did not mean anything essentially private or subjective, such as what he called "ideas" ("Vorstellungen"):

> The idea is subjective: one person's idea is not that of another. By virtue of this alone, there are many differences among the ideas that are associated with the same Sense.… In this, the idea differs essentially from the Sense of a sign, which can be common property of many and is not a mode of an individual soul. For one can hardly deny that humanity has a common treasure of thoughts that it transmits from one generation to another. (Frege 1891, p. 29)

Propositions, Frege's anti-psychologism.

Frege was an "anti-psychologist", which means that he rejected the idea that logic deals with laws of thought in the sense of describing psychological regularities, or describing any concrete psychological occurrences such as experiences or mental images. Those, he thought, were essentially private and inaccessible to anyone but the thinker him or herself, if at all. Rather, Frege held, logical laws are relations amongst certain abstract entities which do not depend on the actual psychological features of anyone. Logic deals with that aspect of a thinker's thoughts that is objective in the sense of being accessible to many people, namely the *content* of the thinker's thoughts. These contents can be thought by different thinkers (though Frege thought there were exceptions to this general rule).

Several people can believe (doubt, reject, suppose, etc) the same content or proposition. Thus, one person can believe the proposition that there is a bear near the clearing, and another person can believe (disbelieve, consider, etc) the very same proposition. Propositions, in turn, are structured in the sense of having constituents that can recur in different propositions – constituents that I want to call "concepts" here. Thus, the very same concept – say the concept of a clearing – can occur as a constituent part in different propositions.

Propositions as public

Frege was primarily interested in the logical properties of languages, and his main interest weren't natural languages but ideal scientific languages. Nevertheless, his anti-psychologistic approach has found many followers amongst theorists interested in the semantics of natural language. This following continues until the present day. Natural language semanticists still postulate public abstract entities – calling them "propositions", "intensions", "contents", "truth conditions" – and describe the central semantic properties of natural language sentences by saying that sentences (or sentences in contexts) express these contents (propositions, intensions etc). There is plenty of debate about the exact nature of these entities, and about the legitimacy of postulating them, but most theorists concur in the general anti-psychologistic approach on which publicly accessible abstract entities play a key role in semantic theorising.

Propositions "measure" mental states

Propositions are abstract entities.[8] According to the approach we are exploring, propositions are individuated in terms of the concepts that constitute them and concepts are in turn individuated in terms of certain rules of correct application.[9] This means that in principle, propositions and concepts can be examined by purely a priori methods, just as numbers or artificial languages can. However, this should not deter us. It always remains an empirical issue which proposition or concept a given thinker is employing or expressing at a given time, so that ultimately our theorizing remains answerable to experience. Some theorists have compared propositions in the description of people's mental states to numbers on a scale with which we measure some quantity such as, for example, the speed of an object. While plausibly the numbers with which we measure speed stand in relationships that can be examined in purely a priori ways – it is a priori that a speed of 5 meters per second is numerically greater than a speed of 3 me-

8 Even if, according to some, propositions can have concrete constituents.

9 The exact construal of propositions remains a controversial matter: some would prefer to say that propositions are constituted not by concepts but rather by objects and properties (e.g. Salmon 1986), or that propositions are sets of possible worlds (e.g. Stalnaker 1987). It is not crucial for current purposes whether we adopt a Russellian direct-reference approach or conceive of propositions as having concepts or senses as constituents.

ters per second – it remains an empirical matter which objects have which speeds, what causes and effects objects with given speeds have etc.[10]

Psychologistic alternatives

Frege's approach contrasts with a more psychologistically oriented outlook that treats beliefs and their conceptual building blocks as concrete psychological representations in the minds of individual thinkers. This is exemplified by the empiricists Locke and Hume, as well as contemporary theorists like Fodor (1975), Dretske (1981) or Laurence and Margolis (1999). Fodor, for example, believes that while mental states represent certain objects, it does not make sense to theorize about mental content in ways that go beyond the representational function. While the general term "bear" represents bears, or has bears in its extension, and a token of "the clearing" represents, or refers to, a particular clearing, there are no abstract "senses" or modes of presentation by which these representational relations are mediated. Thus Fodor and other defenders of psychologistic theories wish to theorize about the mind only in terms of causal representational relations. I do not wish to deny that it makes sense to approach the human mind in these ways: clearly, one can examine the very same psychological processes in several different empirical ways.

Argument for the legitimacy of shared contents

In the next chapter, I will discuss the role of abstract entities like propositions and contents in semantic theorizing, i.e. in theorizing about the meaning of natural language expressions. At this point, I am addressing the more basic issue of whether we can legitimately theorise about the content of speech and thought in the non-psychologistic way proposed by Frege and pursued by many semanticists today. I want to defend the legitimacy and naturalistic respectability of theorizing about cognitive states in this way. Describing beliefs and other cognitive states as having abstract propositions and concepts as contents, that are governed by shared norms, helps explain the possibility of epistemic resource sharing. It makes no difference for my defence whether the beliefs of individual thinkers are ultimately constituted by certain states, processes or dispositions of the brains of these thinkers – states, processes and dispositions that are governed by psychological, and ultimately physiological and physical laws.

The argument starts from our basic picture, namely from the commonsense idea that we share information about the world by talking to one another. It seems clear that at a minimum, typical cases of successful information sharing involve the following correlations: one person, who has a true and justified belief, expresses this belief by making an assertion to an

10 Robert Matthews (1994, 2007) examines the analogy between measurement and propositional descriptions of the mind.

audience, and as a result the audience comes to believe what they have been told and thereby acquires a belief that is itself true and justified. For example, a testimonial source might tell an audience that there is a bear near the clearing, the audience can, as a result come to believe that there is a bear near the clearing, and the truth and justification of this resultant belief will often be a non-accidental product of this interaction.

Transmitted advantages

We will typically describe both the original belief and the newly caused belief as beliefs *that there is a bear near the clearing*. I am not taking this to establish that there is a single object that they both believe. However, it does seem clear that acts of communication like this can benefit the recipient in ways that are not random but quite systematic. Certain epistemic advantages have been extended from the first person to the second, advantages that are the result of the truth of the original belief. What are these advantages? If Anna truly believes that there is a bear near the clearing, she can use this belief in her theoretical and practical reasoning, and if the reasoning is good and other beliefs on which she is relying are also true, she will have a greater chance of achieving her goals. For example, if she wants to find a bear, she can look for one in the vicinity of the clearing in question, and her chances of finding one will be increased. Similarly if she wants to avoid an encounter with a bear, she will keep away from the clearing and the truth of her belief will, *ceteris paribus*, increase her chances of success. These practical advantages are extended from the testimonial source to the recipient. If Anna tells Ben that there is a bear near the clearing, and Ben believes her, then Ben will enjoy the same or similar advantages. This seems to be part of the point of information sharing.

Now, for this transfer of epistemic advantages to work, Anna and Ben need to have capacities for thought that are in some sense commensurable. Anna's and Ben's respective bear-thoughts need to be similar in some sense, and so must be their thinking about the clearing. Of course there will also be individual differences between them. Anna may know or believe things about bears or about the clearing that Ben doesn't know or believe, and vice versa. Moreover, they may have different goals, so that each of them might make quite different use of the belief in question. However, it seems clear that success in communication involves that the concepts they employ in the belief are in some sense similar – in a sense that allows them to share the advantages of believing truly that there is a bear near the clearing. What is this sense? Since the advantages transmitted from a testimonial source to a recipient of testimony are ultimately practical advantages, the similarity between the belief expressed by the source and the belief acquired by the recipient must be a similarity in the practical impact of the beliefs. The two beliefs will, *ceteris paribus*, have the same role in reasoning. In other words, if they both had the same fur-

ther beliefs and the same preferences, then the belief would have a similar effect on their practical and theoretical reasoning.

Similarities of source's and recipient's belief

It is difficult to see how else testimony, in successful cases, could involve a transmission of cognitive advantages from source to recipient. For the non-accidental advantages we gain when acquiring true beliefs are ultimately advantages that result from a fit between our actions and environmental conditions. Having a certain belief is non-accidentally advantageous for Anna when the world is a certain way (i.e. as the belief represents it as being). If the advantageousness of the belief Ben acquires as a result of Anna's testimony did not depend on the world being exactly the same way, then one could hardly speak of the cognitive advantages being *transmitted* from Anna to Ben. Since these advantages are ultimately practical, Anna and Ben must be similar in the role that the two beliefs play, respectively, in their practical reasoning. Since this is a *similarity* in some respect, it is unproblematic to abstract and speak of "what the two beliefs have in common". In the simplest case, we might say that two beliefs *have the same content* just if they are similar in the way described.

Sensitivity to similar norms

The obvious way to explain how two thinkers can non-accidentally come to have beliefs that are similar in this way is to say that they have been calibrated by certain social processes to have belief faculties that are compatible for information exchange with other thinkers in the same community. These calibration processes are partly processes that rely on the natural development of young humans in general, and partly they are cultural in the sense that they have aspects that depend on the contingent history of a particular community. The two thinkers have acquired their capacities and methods of belief acquisition and reasoning through certain learning processes, and these processes ensure that its participants employ similar concepts. The similarity consists in the thinkers' being sensitive to similar rules of correct application of concepts and correct belief. In particular, they are sensitive to similar norms of correctness for their respective beliefs that there is a bear near the clearing, and to similar norms of correct application for constituent concepts such as *bear*, *nearness* and *the clearing*. From the point of view of effective information sharing, we can think of these similarities as the purpose of the processes of calibration.

Being sensitive vs being subject to norms

The norms to which each thinker is sensitive may differ, though they must be similar enough for information sharing to be feasible. Thus, if we wanted to say that the concept a thinker is employing in a given episode is the concept individuated by the norms of correct application to which that thinker is *sensitive* at that moment, then we would have to say that these two thinkers employ distinct (though similar) concepts. In the anti-psychologistic model I am proposing, however, the concept a thinker is

employing in a given episode is individuated in terms of the norms she is *subject* to, and I want to say that both are subject to the same norms. The notion of *being subject to* a norm is an idealisation that results from abstracting from the individual differences in *sensitivity* to norms that I mentioned. Thinkers are *subject* to those norms that they would be *sensitive* to if the processes of calibration I mentioned were maximally successful. On this view, then, Anna and Ben are employing the same concepts of *bear*, *nearness*, and *the clearing*. On a simple elaboration of this idea, we might say, moreover, that the belief expressed by Anna has the same *content* as the belief acquired by Ben. But as we will see, this is not the only elaboration.

No irreducible normativity required

I would like to emphasize that talk of correctness conditions for concepts and for believing contents does not amount to a claim of *irreducible* normativity. Sure, talk of correctness conditions and saying that thinkers are subject to these norms of correctness does amount to making normative claims. However, whether this amounts to irreducible normativity depends on how we understand talk of thinkers being subject to these norms. One way of understanding it is to say that individual thinkers are sensitive to certain norms and that their sensitivities are co-ordinated as a result of social processes. These are, on the one hand, the natural processes of language learning and other processes of child development. Children come to employ the words and concepts of their social environment. On the other hand, these are the socio-cultural processes of education and deference to experts. The norms thinkers are *subject* to are then an idealization of the *sensitivities* as they would arise from these processes under ideal conditions.

No categorical normativity required

Moreover, there is no need to view the normativity involved as being categorical. For thinkers may be subject to a requirement to comply with the norms they are subject to only on some condition. For example: *if* a thinker is to be maximally calibrated with respect to other thinkers in her community, then she ought to comply with the norms to which she is subject. Or: *if* a thinker is to employ concepts as they would do if they had been ideally trained, then they ought to comply with the norms to which they are subject.[11] But the condition need not itself be a requirement, so that no categorical *ought* can be derived. Or perhaps the condition is a requirement, but only a conditional one: *if* a thinker is to maximise the likelihood of benefit from communicating with other thinkers, he or she ought to maximise calibration with other thinkers, and if he or she is to maximise

11 In the wake of Kripke 1982, there has been extensive debate about the alleged normativity of language and thought. For a recent argument against the normativity of thought, see Glüer & Wikforss 2009.

calibration, then she ought to comply with the norms he or she is subject to. Thus, the picture I am sketching is compatible with all the norms involved being merely conditional norms.[12]

Individual thinkers are sensitive to certain norms. They will accept that calling something a “clearing” or a “bear”, or saying that one is near the other is correct only under certain conditions, and they will largely agree in what these conditions are. Individuals will also have concepts that they tend to apply only under certain conditions, and other individuals will tend to apply these concepts only under the same or similar conditions. The fact that different individuals are *subject* to the same norms explains that they are *sensitive* to *similar* norms (for sensitivities are the result of a calibration process of which the norms to which they are subject are an idealisation). And this in turn explains how one thinker can easily acquire the epistemic advantages another thinker has by believing what the other says. If both are subject to the same norms of application of *bear*, and both are consequently sensitive to similar norms, and consequently both apply the concept *bear* to the same kind of thing then there will be no unpleasant surprises.

Natural and institutional calibration

Different communicators can possess the same concepts in the sense explained of being sensitive to the same or similar norms of correct application, whatever these are.[13] How exactly is it that different thinkers in the same community have conceptual tools that overlap in this way? I offered an obvious explanation: thinkers do not develop their intellectual skills in isolation, but rather they acquire their capacity to form beliefs in a way that ensures similarity with other thinkers. Novice thinkers acquire their intellectual methods and habits from mature thinkers, be it in the course of the natural process of language acquisition, by imitation or through deliberate training. In addition there are certain congenital similarities among all hu-

12 The idea of a conditional requirement can be formally captured by stressing that the *ought* in question takes wide scope with respect to the conditional: *Ought* (if *p*, then *q*), rather than being part of the conditional’s consequent: if *p*, then *Ought q*. The important difference is that in the case of the former, the consequent cannot be detached via modus ponens and a non-normative premiss *p*.

13 There is room for considerable disagreement about the kind of norms of correct application that govern concepts, and to which thinkers and language users are subject. Thus the conditions for the correct application of the concept *bear* might involve that bears have a range of observable features, or they might involve that bears must have descended from the specimen to which the concept was originally applied. For the current discussion we can just leave this matter unresolved. The important fact to which I want to draw attention is simply that it is crucial for the success of communication that language users are mutually calibrated in their conceptual repertoires, which then facilitates (and explains) successful sharing of epistemic resources.

mans that facilitate this calibration.[14] Even mature thinkers take part in processes of developing and maintaining their methods, processes that ensure a certain minimum level of conceptual calibration within an intellectual community. There even are elaborate social institutions that serve this process of conceptual calibration. Some members of a community will have a special status in the maintenance of the conceptual tools in certain areas – by counting as experts. Non-experts in an area will defer to experts for the correct application of concepts in that area. This ensures that different thinkers in the same community are sensitive to concurrent norms of correct application.

Concurrence between experts and lay people

In fact, the existence of experts forces us to make more precise what kind of similarity in sensitivities is fostered by the community's training schemes. The expert's sensitivity might well differ from the sensitivity of the non-expert who defers to the expert in her employment of a concept (even under ideal conditions). I am not, but an expert is, able reliably to tell a leopard from a jaguar in the wild. Thus, the expert will apply the concept *leopard* on occasions when I wouldn't. Does this mean that our sensitivities are out of line or dissimilar? That something has gone wrong? Of course, my knowledge is not as great as it might be when it comes to leopards. But that's normal, and it's the point of distinguishing experts from non-experts. Nothing has gone wrong as far as my sensitivity to norms in the case of the concept *leopard* is concerned. I defer to the experts, which means that I will be prepared to call something a "leopard" only if I believe that an expert would. If we want to say that the expert's sensitivity is different, this does not mean that our sensitivities are out of line. Since mine defers to hers, we concur.

Faulty sensitivities

Now, while sensitivity to norms is a matter of the dispositions and psychological profiles of individuals, being *subject* to norms is a matter of the individual being part of a community of thinkers and taking part in the mentioned processes of conceptual calibration. This means that the norms to which an individual is subject and the norms to which she is *sensitive* can in principle diverge, even though they are not wholly independent of one another, for in general it is *because* individuals are *subject* to certain norms that they become *sensitive* to the same or similar norms. The norms an individual is *subject* to are those that she would be sensitive to if she had had maximum exposure to the community's processes of conceptual training and maintenance, and had fully absorbed all the lessons. For example, an individual, call him "Hans", might mistakenly apply the con-

14 Markman & Hutchinson 1984 and Markman 1994 present experimental results that show children in the early stages of language learning to have clear biases in extending a newly acquired word to new cases.

cept he expresses by “bear” to a boar, and this might be in accordance with his own sensitivities to conceptual norms, i.e. Hans will consistently apply the concept to boars even in other situations. This would be a case in which the norms the individual is *sensitive* to diverge from the norms he is *subject* to. This will be the result of incomplete or faulty training, or of errors in applying the lessons learned in training. As far as his use of the word “bear” and the associated concept goes, this individual would not be a fully functioning cog in the community’s processes of information sharing.

Mistakes in employing a non-faulty sensitivity

This case must be distinguished from a case where an individual makes a mistake by the lights of her own, non-faulty sensitivity. Thus, a city kid might mistake a boar for a bear simply because she lacks the information to employ her own sensitivity correctly. She might defer in her use of the concept to more expert users, so that her own sensitivity is best summarised by saying that she has the tendency to apply the concept *bear* to whatever the experts call a “bear”. She might then believe of a boar in front of her that this is an animal that would be classified as a bear by experts, thus apply the concept. However, the city kid would correct her judgement once she learns that an expert would not classify this as a bear, and this correction would not involve a change in her sensitivity. Thus, in this case the discrepancy is between the individual’s performance on a particular occasion and their general sensitivity to a norm, where this general sensitivity need not be out of line with the norms that she is subject to by virtue of being a member of an epistemic community. Thus there is a difference between having a faulty sensitivity and making a mistake in bringing to bear one’s sensitivity on a particular case.

Aspects of idealization

Talk of the conceptual norms to which thinkers are subject is thus an idealization which captures the common standards of concept application maintained by a community of thinkers via its training and calibration schemes. This idealization suffers from the same problems as other idealizations: just as there are no frictionless plains there are no individuals that are exactly sensitive to the norms they are subject to. However, individuals tend to approximate the ideal, and it is this fact that allows us to explain the evident success of epistemic resource sharing. Another aspect of the idealization is that the processes of calibration I described are messier in reality than in my story. The boundaries of communities of thinkers are fuzzy and complicated, and individuals are typically members of many overlapping groups that take part in information exchange and conceptual calibration. Individuals are members of families, social groups, countries, professional organizations, the worldwide community of humans etc. Thus, it is an enormous simplification to speak of just one community and its processes of conceptual calibration.

I said earlier that I was going to defend Frege's idea that we can profitably theorize about the beliefs and linguistic acts of thinkers by describing them as propositional attitudes and propositional acts, and that this involved describing beliefs and linguistic acts by assigning them propositions as contents, where propositions are publicly accessible entities. I also said that there is an analogy between describing thinkers by relating them to abstract propositions, and the measurement of magnitudes by relating them to numbers on a scale. We describe the speeds or temperatures of objects by assigning to them numbers on a scale (meters per second, degrees Celsius), and in this assignment some of the properties of, and relations among, numbers represent some of the properties of, and relations among, the objects thus measured. The notion of being subject to norms that I have just introduced allows us to say a little more about the types of properties of thinkers that are measured by ascriptions of propositional content to their beliefs and linguistic utterances.

Propositional contents as representing the norms to which thinkers are subject

Let us think of propositions as constituted by concepts, and of concepts as individuated by their norms of correct application. Then describing a thinker as believing a proposition amounts to describing her as someone who is employing the concepts that constitute the proposition and being in this employment subject to the constitutive norms of these concepts. Whether the thinker is subject to such norms in not directly manifest in her behaviour. But it is indirectly manifest in that she is a member of a community that enforces these norms, and is as such likely to be sensitive in her employment of these concepts to norms that concur with these norms to which she is subject.

2.3 A Puzzle about Expertise

Before I move on to make some observations about two broadly different types of calibration, I would like briefly to discuss a puzzle that arises in connection with the idea of deference to experts. Consider the concept of a blueberry. In applying the concept to a given object, I might use as a criterion that that thing must have a certain appearance, that it must look a certain way. For example, if it grows on a shrub of a certain sort, has a certain size, shape and colour, or tastes a certain way, or comes in a box labelled "blueberries", then any of these will be good grounds for me to apply the concept expressed by "blueberry" to that object. But not everything that looks like a blueberry or tastes like a blueberry, or comes in a box labelled "blueberries", is a blueberry, even though having any of these attributes is normally a good reason for thinking that it is. A responsible and competent user of the concept must use, in applying the concept, some criterion that is at least a good indication that the conditions for correct ap-

plication are met. Nevertheless, the user may ultimately defer to experts: he or she might use the criteria in question only as a result of believing that they reflect the experts' superior criteria, thus hoping to manifest a sensitivity to norms that concur with the experts.

An expert on blueberries might use certain essential criteria that overrule considerations of the shape, colour and size of a candidate blueberry. A particular blueberry might not have the shape, the size, the colour or the taste of a blueberry, but nevertheless be a blueberry because it is the fruit (perhaps grown under nonstandard circumstances) of a blueberry plant. And again there will be essential criteria, maintained by the experts, for deciding what shrubs are genuine blueberry plants. Perhaps it is part of the concept of a blueberry plant, as maintained by biological experts, that blueberry plants are a natural kind, where perhaps experts in biological systematics have again the authoritative view on what counts as a natural kind of the relevant sort.

Puzzle: mistaken experts

Now, this dependence on experts gives rise to a puzzle. Suppose that the vast majority of users in a community happily use the term "blueberry". They almost always agree on what counts as a blueberry, they pick them when ripe and bake delicious blueberry pancakes and blueberry muffins. In their wisdom, the experts of this community, i.e. those users to which the majority defer in their judgements as to what is a blueberry and what is not, go beyond the criteria successfully used by the majority. They develop a biological theory according to which blueberry plants fulfil a number of scientific criteria. With this theory, they try to capture something that those plants from which the ordinary users pick the berries have in common. Now suppose that as a matter of fact this theory is mistaken: nothing has the characteristic that the experts think purported blueberry shrubs have. Should we conclude that all the ordinary users have got it wrong?

This is indeed puzzling. If we say that the concept shared by the community and expressed by the word "blueberry" is individuated in terms of the conditions of correct application favoured by the experts, we end up saying that every one of the users misapplies the concept. But this seems wrong: the judgements everyone would express using the word "blueberry" are quite successful otherwise: everyone relies on these judgements and no-one ever experiences any unpleasant surprises. The blueberry pancakes turn out exactly as they should, and no-one has reason to be unhappy with the way their actions are guided by the judgements they would express using the word "blueberry". It would be much more plausible to say that the experts made a mistake about the essential application conditions of the concept expressed by "blueberry", and that everyone else was applying that concept correctly. However, if we don't defer to the experts in what the correct application conditions of a concept are then

how should it be decided whether in a particular case people are applying the concept correctly? Or rather: how should we decide which concept people are using, the one defined by the experts' theory, or a concept that is defined by the non-experts' criteria (and if so, how?)? Which norms are users really subject to?

Resorting to natural kinds

One way out of this puzzle is to say that there is a certain natural kind (where this kind need not coincide with anything the erroneous experts would recognize), and that the concept correctly applies only to members of that kind because in the existing practice it is members of that kind that are causally responsible for the vast majority of judgements involving the concept expressed by "blueberry". The error of the experts is an error concerning that natural kind. However, this way out presupposes a theory of natural kinds, and it will only work for those cases where there is a relevant natural kind according to that theory. Do we have a theory of natural kinds which covers all the cases? In the example given, the experts had an erroneous scientific theory of certain biological kinds, so we can't simply resort to a theory of natural kinds provided by scientific experts to solve this particular puzzle. For those experts might again be mistaken. Moreover, there seem to be concepts in use that do not lend themselves as naturally to the natural kind treatment as the concept we express by "blueberry". Consider for example the concept expressed by "hill", or that expressed by "hip-hop" or "cake".[15] So I propose not to rely only on a theory of natural kinds for the resolution of this puzzle.

Another way out: our conceptual repertoire develops.

Another way out is to concede that it is sometimes not clear which exact concept a given word expresses. Concepts are abstract entities that are individuated merely by the norms of application that govern them. Whether anything falls under a concept does not in any way depend on any user of the concept, not even on the most authoritative of experts. Thus, our puzzle does not affect the abstract realm of concepts. Concepts are pure and they come with definite norms of correct application. The mess is in the realm of human thought and human communication. We hope to maintain, via various processes of calibration, a harmonious communal repertoire of concepts, and the mechanism of deference to experts is part of this system. However, the system does not always work without problems. We might say: "concepts develop over time". But strictly speaking I should say: the conceptual norms we are sensitive to develop over time, and so do the conceptual norms we are subject to. This means that we sometimes un-

15 Though one might take the view that in areas where the concept of natural kind does not help, there are various concepts of artificial kinds which are equally suitable for the resolution of the puzzle.

See http://web.mac.com/cranetim/Tims_website/Jaffa_cakes.html for an expert opinion on the rules governing the concept of a jaffa cake.

whittingly abandon concepts and adopt new ones, for the concepts we use are the ones individuated by the norms we are subject to, and which norms we are subject to can change. In some cases, the concept of a natural kind can stabilize these developments and save us from having to say that concepts change when our criteria change. But this only works in those cases where a sufficiently robust notion of a natural kind is available.

2.4 Unisono and Polyphonic Calibration

I have proposed that one way to describe our beliefs and speech acts is by describing them as having propositional content, and that these contents are individuated in terms of communal norms of correct application to which users are subject. What norms users are subject to is in turn a result of social processes which ensure that different individuals are sensitive to concurrent norms. I have argued that this provides a way of explaining how it is possible for communities to share epistemic resources, for the mutual calibration of different individuals in the same community makes them in principle suitable as recipients of each others' testimony.

Unisono and polyphonic calibration

One example was the concept of *blueberry*. Different thinkers will concur in the norms to which they are sensitive in applying this concept, and they will do so as a result of being trained in the same way, deferring to experts, etc, i.e. as a result of being subject to the same communal norms. In the case of this concept, the norms of correct application are such that if one person correctly applies the concept to an object then it will be correct for anyone else also to apply the concept to that object. This means training up a new user of the concepts will involve making sure that the novice applies the concept to the same objects to which the experienced user is applying it. For if it is correct for the experienced user to apply the concept then it is also correct for the novice to apply it. Let me call this way of inducing new users to the correct use of a concept "unisono", for teacher and novice, as it were, sing the same notes. What unisono calibration achieves is concepts that are objective in the following sense: if one of two people applies the concept to an object and the other denies it of the same object, then it is a priori that one of them is misapplying the concept. This is an a priori matter because it is built into the calibration scheme for the concept that it can be correct for one person to apply the concept only if it would also be correct for others to apply it.

Many concepts are calibrated in the unisono way and are consequently objective – *tractor*, *green*, *water*, *oak* are examples. But quite obviously not all concepts are like that. Consider the concept *I* (i.e. the one expressed by the first person singular personal pronoun), or the concepts

mine, *yours*, *there*, *here*, *left*, *right*, etc. Clearly, successful induction into the use of these concepts does not require that teacher and novice come to apply the concepts to the same objects. In fact, in some cases it requires the opposite. This is Sesame Street territory. Consider the sketch where Grover is going to help Kermit explain the difference between *here* and *there*: [16]

Sketch from Sesame Street

[Grover is carrying an armful of heavy bricks, Kermit has been trying to tell him where to put them down.]

Kermit [pointing]: Why don't you put the bricks over *there*.

Grover [can't see where Kermit has pointed and is about to put them dangerously close to Kermit] Put bricks here.

Kermit: Aaahhh, Grover, Grover, you are gonna put them on my foot! I told you to put the bricks *there*. Don't you know the difference between *here* and *there*?

Grover [huffing and puffing under the weight of the bricks]: Ah, no, explain, froggy.

Kermit: Ok, well, any place you are, Grover, is *here*.

Grover: Here?

Kermit: Aha. And, you see, any place where you are not is *there*.

Grover: Ok, any place Grover is not is there, ok.

Kermit: So put the bricks over *there* [pointing backstage]

Grover [moving backstage]: Ok, put bricks over there, ok. [Having arrived, shouting from backstage] Ok, froggy, here I am. [Mumbling to himself] Here? Uuhh, here? [Shouting to Kermit] I am here?

Kermit: Yes, put them down now.

Grover [getting desperate under the weight of the bricks, coming back upstage]: Froggy, I tried to, but when I got *there*, it was *here*.

Kermit [shaking his head]: Ok listen, I'll tell you what I'll do. I'll make things easy for you, alright? Now you stay up *here*, and I will go *there*, and then you can put the bricks down *here*.

Grover [exhausted]: Whatever.

Kermit [moving backstage]: I'll go there. [Having arrived] OK, you can put them down now.

Grover: I put bricks down here.

Kermit: Yeah, right *there*. [Then moves off stage.]

---

16 The original sketch can be found on youtube http://www.youtube.com/watch?v=7iWy5IObbIE or http://www.youtube.com/watch?v=sUX_LnyyM8Y.

Grover [confused]: There? Froggy? Froggy? Oh froggy! He say there, but I am here!

[Thinking, exhausted under weight of bricks, he really needs to put them down.]

Let's see, any place I am is here, any place I am not is there. So: I put bricks some place I am not!

[He throws bricks away from himself offstage.]

Grover did it! I am *here*, but bricks *not* here, so Grover put bricks *there*! Right froggy?

[Close-up on Kermit as he is climbing out of the pile of bricks, which have landed on him.]

Kermit: Uff, right, Grover, I just wish you hadn't put the bricks *here*.

When Kermit says "Any place where you are, Grover, is *here*. And any place where you are not is *there*." he is not articulating the rule correctly, or at least misleadingly. For he can easily be understood to be articulating the incorrect unisono rule that "here" refers Grover's location, and "there" to all other locations. In fact the rule is that "here" refers to the place where Grover is only when Grover himself uses the term. When Kermit uses the term, it refers to the place where Kermit is. But they can be in relevantly *different* places, as in the sketch. Thus, when Kermit and Grover are fully competent with the concept, they will occasionally apply the concept to different places.

Rules of correct application for indexicals

What Kermit needs to achieve in teaching Grover is that Grover apply the concept he expresses by "here" to the place where he is, and this place will on many occasions differ from the place to which Kermit would apply the concept. Thus, the training will not involve unisono calibration, but what I call "polyphonic" calibration. When several voices play unisono, they all play the same note at the same time. When they play a polyphonic piece, they will often play different notes at the same time. However, they will not do this in a random fashion. Rather, the different voices, even if they play different notes, will be in harmony. Something analogous holds for polyphonic calibration in the calibration of concept use. For while different users apply the concept to different objects even on the same occasion, they do so in a co-ordinated, harmonious way. Kermit and Grover should apply "here" to different places when they are in different places, and under certain circumstances (e.g. when "there" is accompanied with the right demonstrative gesture), Grover should apply "here" exactly to the place to which Kermit applies "there", and vice versa. Grover has mastered the rules governing the concepts expressed by "here" and "there" only if he knows he has been instructed to put down the bricks where he himself is, the place he himself would call "here", when Kermit says "put them down there" while pointing to Grover's position. Other concepts, such as *I*, *you*, *mine*, *yours*, *now*, *then*, *today*, *yesterday*, *this*, *that* etc are simi-

larly co-ordinated: competent error-free use does not involve that everyone apply the concept to to the same things, rather, it often requires different people to apply it to different things. But these requirements of correct application are co-ordinated in such a way that the time to which one person correctly applies "now" is the time to which this or another person correctly applies "then" on an earlier or later occasion, and so on.

Even though the correct extensions of the concepts in question seem to vary with certain features of the occasion on which the concept is employed (such as the person employing the concept, the time or place at which they are using it) there is a very good sense in which the relevant norms of correct application are always the same. The concept *I* always applies to the person employing the concept, the concept *now* always applies to the time at which the concept is being employed, etc. In this sense fully calibrated users are sensitive to the same norms of correct application even in the case of concepts that are taught polyphonically.

Alternative: different people express different concepts with "here"

One might, slightly unhelpfully, insist that the norms are different for different people: the norm for Kermit is that he should apply the concept expressed by "here" only to the place Kermit is at the time of application, while the norm for Grover is the norm that he apply the concept expressed by "here" only to the place where Grover is at the time of application. In that way, one might therefore insist, the concepts that Grover and Kermit express by "here" are distinct concepts and *these* concepts are unisono: anyone using them correctly will apply them to the same object.[17] If we say this, we still have to account for the ways in which different users are co-ordinated. Thus, when interpreting others, users would have to 'translate' in such a way that whenever Grover uses "here", others should interpret him as having expressed the concept they themselves would express by "the place where Grover is". However, this leads to a problem. For presumably Grover also expresses some concept with the phrase "the place where Grover is". But *that* concept does not seem to be the same concept as the one he expresses with "here". For if Grover temporarily forgets that he is Grover, the belief he would express by saying "there is immediate danger *here*" will play a very different role in his thinking from the role of the belief he would express by "there is immediate danger where Grover is".[18] So it seems that Grover expresses different concepts with "here" and "the place where Grover is". If the second of these is the

17 Frege's approach is along these lines: the thought expressed by a sentence containing "I" when used on a particular occasion contains a sense corresponding to the token of "I" that is unique to the speaker of that occasion and is only accessible to that speaker. See Frege 1918 and Künne 1992. Note that Frege attempts to find a different way around the difficulty that I am about to mention.

18 See Perry 1977, 1979.

same as the concept others are also expressing with “the place where Grover is”, then we cannot say, as we were considering, that Grover expresses with “here” the same concept Kermit expresses with “the place where Grover is”.

There are other possible solutions to this problem, but for the moment I would like to stick to the idea that there are genuinely polyphonic concepts, i.e. that the concept one user expresses by “here” is the same as that which others express by “here”, i.e. that they are subject to the same norms of correct application. These norms, however, are such that places to which the concept may be correctly applied can vary from person to person, and from time to time. The aim of the process of calibration is to bring in line the sensitivities of different users in such a way that they tend to comply with these norms.

Exercise 3*: Read Perry 1979. Essay question: What kind of solution to the problem just mentioned here would Perry propose?*

Indexicals and practical reasoning

As we saw in the example just given, the concepts expressed by expressions like “I”, “here”, “now” and other indexicals have an important role in reasoning, particularly in practical reasoning. When I believe that the meeting starts *now*, and I want to take part in the meeting from the beginning, then this will motivate me, everything else being equal, to make my way to the meeting room. When I believe that *I* am the one who is making a mess, and I don’t want to make a mess, I will motivated, ceteris paribus, to take measures to stop making a mess. Corresponding unisono concepts cannot play this exact role. A belief that the meeting starts at 12 noon on March 1st 2010 will not have the same role in practical reasoning, for if I do not also believe that it is now 12 noon on March 1st 2010, it will not motivate me in the same way. Similarly, a belief that MK is making a mess will not motivate me in the same way as the belief that *I* am making a mess, unless I also believe that I am MK. It seems that certain polyphonic concepts are indispensable to practical thought.

“Delicious” as polyphonic

These considerations are well-known in the case of so-called “indexical” expressions like “I”, “here”, “now” etc. They are less well-known and less widely recognized in certain other cases, though in the recent literature they are being discussed extensively. For example, the concepts expressed by a range of evaluative terms such as “delicious” are arguably calibrated in a polyphonic way. What the novice has to learn, when she acquires the concept *delicious* is not to apply the concept to exactly the things to which her expert teacher does. Rather, she has to learn that it is appropriate to apply the concept when she herself would respond to the food in question in a certain, favourable way. People’s responses to food vary. They depend on their *taste*, and their taste varies from person to

person. So, to which objects it is correct for a novice to apply the concept *delicious* depends on her taste; her taste may vary from that of other people, including her expert teacher; so the concept is polyphonic. Again, there is a connection with motivation: when I believe something to be delicious then that constitutes a motive for me to eat it. In fact, the motivational connection provides something of a rationale for the concept's being polyphonic: what my judgements motivate me to do should match my preferences, and my preferences may differ from those of others.

Polyphony and motivationality go hand in hand

Now, there are those who will deny that the concepts expressed by "delicious", "tasty", "funny", etc (and their prima facie equivalents in other languages) are in fact polyphonic. They will insist that, even though sensitivities vary, thinkers are nevertheless subject to norms of correct application that would, if complied with, lead to everyone applying the concepts of taste to the same objects. They believe, as it were, that there is a certain standard of taste to which everyone is subject. Usually those people will also relax the motivational constraint and hold that believing that something is, for example, delicious, does not by itself provide a motive for action. I do not here want to argue that these people are wrong, or that there aren't some communities where the concepts expressed by the evaluative terms mentioned (or their bona fide translations) are treated in the unisono way. However, what I do claim is, first, that many communities do treat these concepts as I suggested and they do adopt polyphonic processes to calibrate users in the use of these concepts in the way I suggested. Secondly I want to claim that it makes sense to have concepts of this sort, and that in principle they have as much practical utility as indexical concepts. One question that will occupy us in coming chapters is how a semantic theory should account for expressions that express concepts of this sort. If some communities don't employ concepts of this sort, then this account will not concern them. But if some communities, actual or possible, do employ such concepts, then this is sufficient motivation to explore possible treatments in a semantic theory.

Moral concepts seem motivational

Generally, evaluative concepts are motivational: employing them gives the thinker reasons or motives to act in a certain way, at least given certain other beliefs. Now one class of motivational evaluative concepts is the class of moral concepts, the concepts expressed by expressions like "ought", as in "he ought to help his friend" or "(morally) wrong", as in "it is (morally) wrong to torture". When someone judges that she ought to help her friend, then that gives her, ipso facto, a reason or motive to help her friend. Moral judgements are intrinsically motivating. This thesis is often called "internalism about moral judgement". In a moment, I will draw attention to an important difference between moral concepts and other evaluative concepts. But before that, let me clarify two issues that might otherwise lead to confusion.

Internalism about moral judgement

The internalist thesis, i.e. the thesis that moral and other evaluative concepts are intrinsically motivating, is merely the thesis that judging in this way gives the judge *a* reason or motive to act in a certain way. That is, *if there are no other relevant reasons or motives*, it will motivate her. But the presence of a reason or motive to do *F* is compatible with the presence of conflicting reasons or motives to act in ways that are incompatible with doing *F*. In the case of conflicting reasons or motives, the agent will have to weigh them up against one another and thus come to an overall judgement as to what she ought to do. We can call the reasons or motives that enter this weighing up process "pro tanto" reasons or motives, and the overall reasons or motives that come out of this process "overall" reasons. The internalist thesis, then, is that judgements of what one ought to do, or what is morally good or bad are pro tanto reasons or pro tanto motives for action. We can express this by saying, for example, that if I judge that something would be morally bad to do, then I will, *everything else being equal*, avoid doing it.

Weakness of will

Now, this does not yet seem to be entirely correct, for there seem to be cases in which people act against their better judgement. That is, even though they judge it to be morally bad to do something, and there are no contravening motives, they still fail to be moved to avoid doing it. Thus, they act against their own judgement as to what they ought, overall, to do. This is usually called "weakness of will", and it is regarded as a form of irrationality. Thus, strictly speaking, the thesis that evaluative beliefs are intrinsically motivating is the thesis that evaluative beliefs provide reasons or motives that will, everything else being equal, *and if the believer is rational*, move the believer to act in a certain way.

Moral concepts seem unisono

Now, it is an interesting feature of moral concepts that they seem to be taught in a unisono way. That is, when a mature thinker trains up a novice thinker in the use of the moral *ought* or of *morally bad*, she will aim to achieve that at the end of the process the novice will tend to apply these concepts to exactly the same things to which the teacher applies them. While parents may try to teach their children to apply "delicious" in accordance with their children's individual preferences and tolerate that their children may diverge in their judgements from them, they will regard it as a sign of failure if their children diverge in their judgements of what is morally permitted or required. They will take the extent of convergence with their children in moral judgement as a measure of their success in teaching them. Thus we seem to have a case of unisono calibration. Thus moral concepts seem to be both motivational, like other evaluative concepts, and unisono, like typical non-evaluative concepts. This two-faced nature of mo-

ral concepts is at the centre of much debate in metaethics.[19]

Polyphony and motivationality go hand in hand

In the case of other evaluative concepts, I pointed out that it makes a lot of sense that these concepts should be both motivational and also polyphonic. For, I said, we should be motivated in accordance with our preferences, and our preferences vary from person to person. Thus, it makes sense that it should depend on a thinker's preferences or taste whether it is correct for her to judge something to be delicious, for such judgements will carry motivational force and drive the thinker to choose those foods that he or she judges to be delicious. If it were not so, i.e. if the concept *delicious* were a unisono concept, all thinkers would be driven to judge the same things to be delicious, regardless of their preferences. Given the considerable differences in people's taste, this would mean that many would be driven to choose what they don't like. So the polyphonic nature of these non-moral evaluative concepts is well-motivated.

Unisono despite motivational

Now, in the case of moral concepts, the motivating feature is combined with unisono calibration. Given what I just said about non-moral evaluative concepts, how can this make sense? Moral concepts seem to have an important function in the regulation of human interactions. The way moral concepts perform this function seems to involve instilling certain norms of correct application in all thinkers, and thereby – via the motivational link – also impose certain motivational constraints on these thinkers. However, in this case the primary aim is not to make sure everyone acts in line with their personal preferences[20]. On the contrary, given the regulatory function of moral thought, it can be required that their moral thought guide agents in ways that go against their individual preferences. Thus, given that moral concepts have this function, of influencing the actions of moral thinkers in ways that beneficially regulate social interaction, it makes sense that moral concepts should be unisono despite being motivational.

Universalizability trivialized?

So what's the difference between moral and other evaluative concepts? It seems to be a fundamental feature of moral thinking that everyone should be evaluated by the same standards, should be required to obey the same moral norms, regardless of individual differences. This is not so in our non-moral evaluative thinking, such as our thinking about matters of taste. Why? It seems that a basic principle of justice or equality requires us to view moral norms in this way. Moral norms are said to be "universalizable". But it is not easy to see what exactly is meant by saying that everyone

19 For a very accessible introduction to these issues, see the first chapter of Smith 1995. Exercise: Read Smith's first chapter and elaborate a tentative answer to "the moral problem".

20 See Mackie 1977, ch. 4, for an explanation of the need for universalizable norms.

should be subject to the same standards and be required to obey the same norms. There is a sense in which any system of moral norms can be made to look superficially universal. For take any two actions and any moral norm, such that one action complies and the other doesn't comply with the norm. We could easily articulate a new norm with which both actions comply. For example, it might seem that if we say that John acted wrongly in torturing the prisoner, then at first sight the principle of the universalizability of moral norms seems to require us to extend the judgement of wrongness to other acts of torture, such as Joan's act of torturing a prisoner. The idea would be that if a system of norms N does not also evaluate Joan's act as wrong, it is not universalizable. However, N might distinguish different kinds of acts of torturing a prisoner, classifying one kind as wrong, the other as permissible. N might prohibit acts of torturing prisoners that fulfil condition C, but permit acts of torturing prisoners that do not fulfil this condition. Thus, after all, N does comply with the principle of universalizability. It seems that any charge of non-universalizability can be rebutted following this recipe, for whenever two acts are distinct acts, it should be possible to find such a condition C.

Saved from trivialization

However, the idea that moral norms should be universalizable need not be completely toothless. Perhaps there are meta-rules that allow only certain, morally relevant criteria for making morally significant distinctions, and disallow all others (e.g. replacements of C in the above attempt to rescue N might be restricted to a limited range of conditions). Such a meta-rule might preclude arbitrary conditions such as the condition that the agent was born on a Tuesday or is called "Luigi". Moreover, these would presumably preclude conditions such as the condition that the agent felt like it, or the condition that the action brought pleasure to the agent. It seems to be typical of moral norms that certain personal features of individuals, such as their personal whim, cannot play a role in moral considerations. Since the point of moral norms is at least in part to create a stable and predictable community, this will preclude temporary whims from being a criterion suitable for making moral distinctions. Moral norms need to motivate people often against their temporary interests, tendencies or preferences. They also need to allow people to exert influence over one another: one thinker can challenge the action of another by reminding them of a moral norm. They might discuss that matter, and eventually come to an agreement as to whether the action in question was permitted. Such discussions make sense only if the justifications that are recognized as relevant are of a limited sort, namely of a sort that is acceptable and accessible to all thinkers. If they weren't, discussion of moral matters could not help people to influence one anothers' actions.

Probability as polyphonic

Another range of concepts that seem to exhibit polyphony of an interesting sort are broadly connected with uncertainty, incomplete evidence or in-

complete information. Arguably the concepts of probability and of epistemic possibility are polyphonic: whether an individual thinker should think that something is probable or that it is epistemically possible depends respectively on the evidence and information available to the thinker. Let us first consider the concept of probability. At least in one sense of "probable", it will depend on the evidence available to a thinker whether that thinker is correct in thinking that a certain event is probable. For example, it may be correct for one thinker to believe that probably the treasure is on the island, because she has evidence that supports this hypothesis, and no evidence that supports the hypothesis that the treasure in not on the island. If the thinker had had further evidence, or if she had lacked evidence that she in fact had, it might have been correct for her not to judge that the treasure is probably on the island. Another thinker, who in fact does possess further evidence, might correctly judge that it is not probable. Thus, whether it is correct for an individual to judge some possibility to be probable depends on that very individual's evidence. Thus different people, in so far as they have different evidence available to them, might correctly apply the concept *probable* to distinct possibilities. This is reflected in the way the concept of probability is taught, so the concept is polyphonic.

Epistemic possibility as polyphonic

Consider the judgement that the treasure might be on the island. Under what conditions is it correct to judge in this way? Arguably, the correctness of such a judgement depends on the thinker's state of knowledge. Thus, if the thinker's knowledge (or knowledge she might easily have acquired) rules out that the treasure is on the island, then it is not correct for her to judge that the treasure might be on the island. It will be correct only if her knowledge does not preclude this possibility. Again, different thinkers may have different states of knowledge. Accordingly, one thinker may correctly judge that something might be the case while another correctly judges that it cannot be the case. Again, the concept seems polyphonic.

Beliefs as to what is probable and what might be the case also have a certain role in reasoning. While they aren't intrinsically motivational in the way evaluative beliefs are, they do seem to have an important role that distinguishes them from corresponding unisono concepts. Consider the proposal that for Grover to judge that something is probable is subject to the same norms as for him (or anyone else) to judge that Grover's evidence supports it. As in the case of indexical concepts, this seems problematic. For suppose Grover has forgotten that he is Grover. Then he might perfectly coherently believe that Grover's evidence strongly supports the hypothesis in question, while at the same time believing the negation of the hypothesis to be probable. By contrast, he could not believe coherently that the hypothesis is highly probable while also believing its negation is probable. If any concept is a candidate for equivalence with *probable* at

all, then it would be something like the concept *is supported strongly by my evidence*, i.e. an indexical concept, for it contains the indexical concept *mine*. A similar argument could be devised against the proposal that epistemic modal concepts can be construed as unisono concepts, by saying, for example, that the concept Grover expresses by "it might be the case that …" is the same concept as the one he and others would express by "it is not ruled out by what Grover knows (or might easily come to know) that …".

Similar phenomena of polyphony have been alleged to exist in a number of further concepts, some of which will be discussed in chapter 5 below. To summarize, then, among concepts that are calibrated polyphonically, rather than unisono, there seem to be indexical concepts as well a range of non-moral evaluative concepts. Moral concepts take a special place in that they share the motivational pull of evaluative concepts, yet seem to differ from other evaluative concepts in that they don't seem to be polyphonic. Concepts connected with uncertainty, such as probability and epistemic possibility are also plausibly polyphonic.

To conclude this chapter, I want to draw attention to a number of differences between polyphonic concepts that are indexical, and those that are not. To recall, concepts like *I*, *you*, *tomorrow*, *here*, *my uncle*, etc, are typical examples of indexical concepts. The expressions ("I", "you", "tomorrow", "here", "my uncle", etc) that express indexical concepts have a well-known feature, namely that when reporting an utterance involving them, we cannot generally use the same indexical expression as was used in the reported utterance. For example, when Anna says at one time "I am thirsty now.", then if someone else is reporting the utterance at a later time, they cannot say, for example "Anna said that *I* am thirty *now*.", but rather, for the report to be correct, the reporter must replace the indexical expression to make up for the changed context of the report. "Anna said that *she* was thirsty *then*." would, for example be a correct report. In principle, one can imagine a language otherwise like English containing expression "I*" and "now*", such that "Anna said that I* am* thirsty now*." would be a correct report of an earlier utterance by Anna of "I* am* thirsty now*." However, as a matter of fact "I" and "now" do not function that way. Whatever the linguistic context within which we use indexicals, they will always pick out the speaker of the context and the time of the context (see Kaplan 1977). "Anna *thought* that I am thirsty now." does not report correctly the thought Anna expressed by saying "I am thirsty now." A correct report would for example be: "Anna thought that she was thirsty then.".

Direct speech seems to be the only exception: Anna said "I am thirsty." would be a correct report. However in quotation contexts, the indexical

expressions are not used but merely mentioned. What is reported here is which words the utterer has used.

Reporting indexical speech

This feature of indexicals reflects the polyphonic nature of the concepts expressed: in the mind and mouth of the reporter, the indexical concept is goverened by the same norms of correct application, but since these norms dictate that what the concepts apply to depends on the circumstances in which they are being applied, the reporter has to adjust the choice of words or concepts if he or she wants to characterize adequately what the reportee expressed.

Non-indexically polyphonic speech is reported differently

However, some polyphonic concepts, and the words used to express them, do not seem to share this feature: they can generally be reported in indirect speech without changing the words originally used in the utterance to be reported. Let's assume, plausibly, that the concept expressed by "delicious" is polyphonic, i.e. that the range of things to which it can be correctly applied varies with the personal aesthetic responses of the concept user. However, if Anna utters the words "The cake is delicious.", then one *can* later report this utterance with the words "Anna said that the cake was delicious.", no matter whether it would be appropriate for the reporter also to classify the cake as delicious, i.e. appropriate to the reporter's aesthetic response. Thus, it seems that the variation in the objects to which "delicious" can be correctly applied can be modified within the scope of an indirect report. Further reflection will reveal that the same applies to the other non-indexical polyphonic concepts mentioned above.

This difference in how polyphonic speech and thought is reported is part of a wider syndrome of differences. When I say "The cake is delicious." or "The treasure might be under the palm tree.", it can be appropriate for you to reply "No, it's not delicious." or "No, it can't be under the palm tree.", thereby signalling some kind of disagreement. This sort of reply is appropriate even if you are perfectly aware that it was correct for *me* to classify the cake as delicious given my aesthetic preferences, or even if you are perfectly aware that the treasure's being under the palm tree is not ruled out by the information available to *me* when making my remark. This is not so in the case of typical indexicals: if I say "I am hungry." it is not coherent for you to answer "No, I am not hungry.". If I say, standing in one spot "Here is where the Iceman was found.", it is not coherent for you to shout back from another spot, a few paces along the slope "No, here is not where the Iceman was found.". Such a reply by you would be coherent only if you thought that you, in using the word "here" were talking about the same spot.

A syndrome of further differences

This points to another aspect of the same syndrome of differences. When I say "He was found *here*.", and you reply "No, he was not found *here*.", but

each of us is talking about a different spot, then this counts as a misunderstanding, as a communicative malfunction, one which can be repaired, for example, if one of us points out “You are talking about that spot over there, while I was talking about this spot over here.” We have simply been talking past one another. By contrast, when I say “The cake is delicious.” and you say “No, the cake is not delicious.” then our employing different standards of taste in our respective evaluations does not render our interchange a misunderstanding, nor does it count as talking past one another.

These features of non-indexically polyphonic concepts and expressions makes it easier to enter into a discussion of what is delicious, what might be the case etc. It makes it easier to exchange our reasons for why we judge differently. It creates a superficial impression that when we are talking about what is delicious or what is epistemically possible, we are talking about aspects of an objective world shared by all parties. These features of our ways of speaking and thinking persist even when we are perfectly aware that the concepts in question are polyphonic, i.e. that it may well be that the concptual norms in play require opposing verdicts by thinkers that differ in certain respects (e.g. their aesthetic responses or their state of information). In later chapters, especially chapter 5, we shall be looking into different ways in which a semantic theory can capture this difference between indexical and non-indexical polyphony. One approach will be to assimilate apparently polyphonic non-indexical concepts to unisono ones, i.e. to claim that they are only apparently polyphonic. Another approach will be to assimilate apparently non-indexical polyphonic concepts to indexical ones, i.e. to claim that they are, after all, like indexicals. A third approach will recognize them as genuinely polyphonic and non-indexical.

3. What Semantics Does

3.1 Introduction: A Formal Model

Summary

In Chapter 2 I claimed that beliefs and speech acts can be fruitfully described in terms of their propositional contents. I construed these contents in a broadly Fregean anti-psychologistic way: contents are constituted by concepts, and these are in turn individuated in terms of norms of correct application, norms to which the thinkers and language users in question are subject. I argued that describing beliefs in terms of contents that are governed by public norms of correctness will help explain the success of testimony. If all the members of a community are trained up to employ the same concepts, i.e. to be more or less sensitive to the same norms of correctness then this facilitates information exchange. When one thinker believes there is a bear near the clearing, and does so truly and for good reason, then the advantages of thus believing can be made available to other thinkers by inducing in them a belief involving concepts that are subject to the same norms. This involves, for example, that the role the concept *bear* plays in the thinking of different members of a speech community is co-ordinated in such a way that the advantages of having correct beliefs can be passed from one thinker to the next.

Natural language semantics

Until this point I have mostly ignored the role *language* plays in these interactions. It is obvious that thinkers need to have at their disposal also a system of linguistic symbols that they can use in order to make available to other thinkers the advantages of their own beliefs. Natural languages are such systems. What, then, are the properties of languages, i.e. of systems of repeatable types of expressions, that facilitate this exchange of information? Natural language semantics is a discipline that attempts to answer this question, or at least to answer aspects of the question.

Formal languages as models

Natural language semantics has a certain history, and this history explains some of the ways in which natural language semanticists still approach natural language today. The discipline was born when people like Carnap, Montague, Lewis, Kaplan, Davidson tried to use the methods of formal logic, as applied to *formal* languages (i.e. stipulatively defined languages) in the analysis of natural languages. Formal logic is an a priori discipline because it deals only with formal languages, languages that have exactly the properties that the theorist stipulates them to have, and examines the deductive consequences of these stipulated properties. It is therefore fundamentally different from natural language semantics, which studies languages as they are actually used by human populations, which is an empirical phenomenon. One way of seeing the relevance of the methods

of formal semantics to natural language semantics is this: we can use formal languages, i.e. languages the properties of which we stipulate, as *models* of natural languages, models that represent certain properties of natural languages and thereby give us a better understanding of them.

How models represent

As in every model, not all the properties of the model are representative of properties of the phenomenon modelled. A wooden architectural model will represent certain spatial properties of a building or possible building. Thus, certain geometrical properties of the model, such as the angle between one part and another part of the model, will represent the angle between one wall and another wall of the projected building. In this case, the representing angle and the represented angle will be the same. Other properties of the model, such as the length of these parts, will represent in a less direct way: the length of the wall in the model need not be identical to the length of the represented wall. Nevertheless, the lengths of the parts of the model will be representative in that they preserve the proportions of the lengths of parts of the building. Yet other properties of the model are not representative at all: the grain and colour of the wood used to construct a model does not usually represent a corresponding structure or colour of the building's surfaces, or in the hardness of its building materials.

In order to use a model, one needs to understand the way in which it is supposed to represent the original. In the case of the architectural model, we need to know, for example, that it preserves angles and proportions of lengths, and that it does not preserve, say, the colour or microstructure of surfaces. The same goes for formal models of natural languages. In order to use these, we need to have an idea of which features or aspects of natural languages they are supposed to represent, and with which degree of accuracy. In the case of natural language semantics, it is not always obvious what exactly the semanticist wants to model. It will be my first task in this chapter to clarify this somewhat. In modelling natural languages, semanticists have long been operating at a level of detail that corresponds to an architect who is merely considering *general principles* for constructing a highly simplified model that captures only the barest outline of the building. They have actually constructed various simple models, each designed to represent a specific aspect of natural languages. Semantics is far away from a model that models all the rooms of the building, and even further away from modelling all the nooks and crannies, the furniture and plumbing.

3.2 What kind of Data does a Semantic Theory Predict?

Simple model: sentences encode the contents of the beliefs they express

The picture is quite complicated, so let me begin with a simplification that I will later abandon. Since we are assuming that different thinkers can employ the same concepts and believe propositions constituted by the same concepts, let us say preliminarily that a language encodes the propositional contents of beliefs speakers express when they use language. If we restrict ourselves to the declarative sentences of a language, the idea is that a language is a system of sentences, and a semantic theory assigns to each sentence a proposition. To use a declarative sentence is to express a belief that has the proposition assigned to it by the semantics as content. On this simplified model, the audience to which such an utterance is addressed will (under certain conditions) acquire a belief with the same propositional content.

One way in which this is a simplification is that even if we restrict ourselves to declarative sentences, utterances of declarative sentences do not always in any straightforward sense “express” beliefs. Thus, I might utter the declarative sentence “Sam smokes.” without believing that Sam smokes. If I don’t *have* the belief that Sam smokes, I can hardly “express” such a belief! Nevertheless, even if I don’t believe that Sam smokes, under the right conditions, I might still *present myself as believing* that Sam smokes by uttering the sentence “Sam smokes.”. My utterance might nevertheless be an *assertion* that Sam smokes. We might have ideas about what it is to assert something: for example, that to assert something *sincerely* is to believe the proposition asserted; or that an assertion is true if and only if the proposition asserted is true. A theory that tells us which sentences are declarative and then tells us for each declarative sentence which content would be asserted if the sentence were uttered (under certain normal conditions), would therefore allow us to make predictions as to the conditions under which such an utterance would be sincere or the conditions under which such an utterance would be a true assertion.

Complementary theory of illocutionary force

A general description of the meaning properties of the sentences of a natural language can accordingly be divided into a theory of illocutionary force and a theory of content. The theory of force divides sentences into declarative, interrogative, imperative sentences, and perhaps more. This tells us for the performance of which type of illocutionary act (e.g. assertion, question or command) the sentence can be used under certain normal conditions. The theory of content tells us for each sentence what its propositional content is, and this content will be the content that is asserted if the sentence uttered is a declarative sentence (and conditions are

normal). Similarly, this content will be the content of the question or command issued, if the sentence uttered is imperative or interrogative.[21]

Model only contents of declarative sentences

The task of the theory of force is comparatively simple, for each sentence marks just one illocutionary force.[22] For our purposes, moreover, it will be sufficient to concentrate on declarative sentences only. To specify the *propositional contents* of the sentences of a natural language (or of an artificial language serving as an interesting model for natural languages), is a major theoretical task and the main occupation of many natural language semanticists. So I will be assuming that we are able to identify the declarative sentences of natural languages, and that our efforts will be directed primarily at giving an account of the content of declarative sentences in such languages. A semantic theory will describe the compositional structure of sentences and the propositional contents they express, that is, the contents of the assertions one would make when uttering these sentences.

Which aspects of contents are modelled?

Such a description will predict certain logical properties and relations such as logical truth, logical falsehood and logical consequence amongst the propositional contents of sentences. That the propositions assigned to the sentences of the language have these properties and stand in these relations will have further consequences for corresponding properties of, and relationships amongst, the assertions potentially made with these senten-

---

21 What I am saying here about interrogative utterances works straightforwardly for yes/no questions only: just as one can assert that *p*, one can ask whether *p*. Wh-questions, like “Who did it?” or “Where are we going?”, do not seem to have straightforward propositional contents, see for example Larson & Segal 1995, ch. ?. Interrogative utterances do seem to have propositional contents that can also serve as the contents of assertion – the command “Fetch me that book.” arguably has the same content as the assertion “You will fetch me that book.” made in the same context. However, conversely there is not, for every assertion, a natural command with the same content: what would be a command that shares a content with the assertion “The moon is a planet of the earth.”? Even if “May the moon be a planet of the earth.” or “Moon, be a planet of the earth.” are acceptable grammatically, they seem to violate certain conditions of felicity required for successful commands. See Searle 1969, ch. 3.

22 Illocutionary force can be occasionally embedded, as in conditional commands: “If John calls, tell him that I am not at home.”. Some have even argued that the best account of natural language conditionals takes the subordinate if-clause to be a modifier of illocutionary force, so that a conditional sentence such as “If John calls, then he is alive.” is not an assertion of the conditional proposition that John is alive if he calls, but rather a conditional assertion that John is alive on the condition that he calls, see Edgington 1995. Nevertheless, even if there is some embedding of force indicators, the task of the theory of force is still easy compared to the main compositional task of compositionally specifying the semantic contents of the sentences of a language, see below. This does not mean that the theory of illocutionary force is uncontroversial. On the contrary, see for example Stalnaker 1978, Brandom 1984, Williamson 1996, MacFarlane 2011.

ces and of the beliefs expressed when the assertions are sincere. If the semantics merely aims to examine these properties and relations, it need not describe the language beyond a certain level of specificity, i.e. it needs to provide a more detailed account only of the logical vocabulary (just like an architectural model meant to represent only the outside appearance of a building need not represent any of the inside structure that has no effect on outside appearance). Thus the semantics need not, in it's assignment of propositions to sentences go beyond the logical structure of the propositions assigned. This is the focus of attention formal semantics has inherited from its roots in formal logic. However, in a second step, the semantics can go beyond modelling logical structure and represent the complete conditions of correctness of the propositions it assigns to sentences.

No language-independent evidence for the norms to which users are subject

Let us consider once again, what exactly such a theory would predict, and what empirical data these predictions would be answerable to. As I said, the main task of the semantics is to assign propositions to sentences, for each sentence the proposition it "expresses". A simplistic, and as we will see misguided approach would be to say that the semantics predicts that the beliefs expressed by sincere utterances of declarative sentences (and the beliefs acquired by audiences of such utterances) have as their propositional content precisely the proposition expressed by the sentence. We might now think that the way to check whether the semantics makes the right predictions is to find out by *independent* means whether the beliefs expressed and acquired do indeed have the contents thus predicted, i.e. whether believers are indeed subject to the corresponding norms. However, as I explained in the previous chapter, the propositional contents we are ascribing to beliefs are idealizations of the calibration procedures of communities. In our framework, the propositional contents of beliefs as well as their conceptual constituents are supposed to be an idealisation, namely they are supposed to be the contents corresponding to the norms of correct application to which a thinker would be sensitive if she was an ideal product of the calibration procedures. This means that the norms to which thinkers are subject, and consequently the propositional contents of their beliefs are not independent of their speech behaviour. For the norms to which they are subject are the norms to which an ideal product of the calibration processes would be sensitive, and the calibration process is mediated by language. It is plausible that it is primarily through the use of language that thinkers correct or emulate one another. A training scheme for thinkers will at the same time be a training scheme for language users.

User-assessments of correctness as data to be predicted

A more promising procedure is therefore to rely on different data, namely data concerning the conditions of correct usage of sentences. Competent speakers will be able to respond to utterances of sentences and say whether, or under what conditions, these utterances would be correct. To be a competent speaker just is to have the ability to tell what would make

an utterance correct or incorrect. Competent users won't be 100% accurate in these judgements, but if their sensitivity tracks the norms to which they are subject, then the assessments that they make of the correctness of utterances will be a good indication of the norms to which they *are* subject. They will be an even better indication if these assessments are uncontroversial among a representative group of competent users.

The data that a semantic theory is supposed to predict, and which are fairly unproblematically accessible to us as theorists, are data as to the correctness of utterances. A semantic theory will therefore assign propositions to sentences, and (where these sentences are declarative) this assignment will amount to a prediction as to which proposition is asserted in using the sentence, and of the conditions under which the resulting assertion would be correct.

Difficulty: many dimensions of correctness

The proposal clearly needs refinement. One problem is that we might be assessing a variety of things when we assess an utterance for correctness –"correctness" is just the most general term used to indicate compliance with a norm, and there are many different norms to which language users are subject, not all of them directly related to the norms governing the concepts they express. For example, we can assess the literal truth of an utterance, the metaphorical truth of an utterance, the truth of implicit messages conveyed by an utterance, the politeness of an utterance, the relevance, usefulness, prudential wisdom and even poetic value of an utterance. These assessments will sometimes conflict. Thus, an utterance of "I could eat an ox." might be literally untrue, while the message conveyed by the utterance is true. An utterance of "You smell like a fishmonger at the end of a long working day" will, even if literally true and not devoid of poetic value, be impolite and often lack prudential wisdom. Each of these types of assessment represents a different sense in which an utterance can be said to be correct. In some sense all of these types of evaluation are assessments of correctness. One might say that a weighing up of these various different dimensions of correctness produces an *overall* judgement of correctness. However, it seems unrealistic to expect the semantic theory to assign propositions to sentences that reflect exactly these overall assessments of the conditions of correctness.

Filtering out correctness data relevant for modelling compositionality: literal truth

A more promising approach would be to concentrate on one central dimension of assessment that may perhaps derivatively help with the other dimensions. Natural language semanticists usually concentrate on correctness in the sense of the *literal truth* of sentences, and I'll say more about this in a moment. Why concentrate on literal truth rather than some other dimension of correctness? Ultimately, semanticists should choose a dimension of correctness that helps them construct an *overall* theory that correctly predicts certain unproblematically accessible data like competent

users' judgements of the conditions under which utterances are true. One central task, as we will see, is to explain how this competence extends to novel sentences that speakers have never used before. Speakers seem to be able to derive this competence with new sentences from their competence with the component parts of these sentences, and with the ways in which they can be used in new combinations. In order to provide a good model, the semantic theory therefore has to show how the propositions expressed by novel sentences depend on the semantic properties of their parts. The semantics of the model language should be compositional. The choice of literal truth as the central notion that semantics aims to predict reflects the judgement of theorists that this task of modelling compositionality is facilitated when the compositional part of the theory predicts literal truth. Thus, the idea is that we can provide a compositional semantics that predicts the conditions for the literal truth of "I could eat an ox.", and that a further element of the overall theory will then explain how this sentence can be used to communicate a further, non-literal content.[23]

Difficulty in identifying data concerning literal truth

We do seem to be able to distinguish quite clearly the question of whether an utterance is *true* from the question of whether the utterance is polite, relevant, useful, prudentially wise etc. There is also a pre-theoretical sense in which we distinguish between *literal* truth on the one hand and metaphorical truth, or the truth of indirectly conveyed messages, on the other. The latter distinction, between literal and the non-literal is harder to draw precisely. It is a notion of literal truth that semanticists are usually interested in as providing the data their semantic theories are supposed to predict directly. As we will see below, making this distinction in a principled way is not at all easy, and quite a lot of controversy has arisen precisely around identifying the relevant judgements of correctness or truth that competent speakers make, and which we are trying to model directly by a semantic theory.

For example, it is an interesting question whether we should say that an utterance of "some of your shares have gone down" can be true when *all*

23 H.P.Grice is the pioneer of this idea. See his "Logic and Conversation". There is lively controversy about how exactly Grice's basic idea is best implemented. The idea is that the semantics predicts the literal content of sentences (in context), i.e. "what is said" by an utterance of the sentence, and that communication of any other contents is then explained by a further theory, which in Grice's original proposal is based on general assumptions about rational co-operative behaviour. No doubt Grice's way of drawing the line between the literal and the non-literal is influenced by his using standard first-order languages of formal logic as the default model language. More sophisticated model languages, like for example ones that introduce contextual domain restriction of quantifiers, allow a treatment of some of those phenomena that counted as non-literal in Grice's scheme as aspects of the literal meaning of utterances (for discussion, see for example Neale 1990 and Stanley & Szabó 2000).

of your shares have gone down. Some might insist that the literal truth of the utterance requires that not all shares have gone down. Others will say that even though an utterance of the sentence carries a non-literal message that not all shares have gone down, which in that case is false, the utterance is nevertheless true literally.

Shelve the difficulty

In order to be able to proceed with our outline of what a semantic theory does, let us shelve the problem of identifying and classifying the data of correct usage exactly. Let us pretend, for the moment, that we can discern amongst all the assessments competent speakers make of utterances, those that concern the literal truth of utterances. This means we are supposing that we have, in many cases, access to data about the conditions under which an utterance would be literally true. We can revisit this assumption at a later stage. A number of further simplifications are also still in place, such as our focus on declarative sentences only, our failure to take into account phenomena of context sensitivity and our failure to consider the deep structure as opposed to the surface structure of sentences. All these will make an appearance at various points below.

Limitations of the project

Perhaps I ought to scale down expectations in yet another respect. In saying in the last chapter that propositions are individuated in terms of their constituent concepts, and that concepts are individuated in terms of their norms of correct application, I might have created the impression that a semantic theory, in assigning propositional contents to sentences, will spell out in detail the norms of correct application that govern the conceptual building blocks of the propositions that it assigns. In fact, semantic theories usually provide a detailed account only of certain syntactic elements, such as logical constants, quantifier phrases or operators, as well as certain frequent modes of composition, such as putting together a quantifier phrase with a verb phrase. Ordinary lexical elements, such as particular verbs, are usually elucidated only by relying on natural language expressions in a meta-language (we will soon see examples of this). This focus on certain syntactic elements is the result of a particular traditional interest modelling the logical properties of natural languages, such as the notions of logical consequence and logical truth. But more importantly from our point of view it is the result of a primary interest in the *compositionality* of natural languages, which explains an important and very useful feature of language use, namely the fact that humans can successfully communicate with the help of novel sentences, i.e. sentences the use of which has not previously been rehearsed by the participants in question.

3.3 Compositionality and Semantic Values: the Basic Idea

Let us now move on to some of the more detailed questions of how a semantic theory can describe a language in such a way that predictions of this sort can be made, i.e. how the semantic theory can assign propositional contents to the sentences of a language in such a way that these contents correspond to the judgements of literal truth of competent speakers. As I said, one central task in devising such a theory lies in modelling the compositionality of natural languages. Natural languages quite obviously contain semantically complex sentences, i.e. sentences that are formed by concatenating (=chaining up) several independently meaningful expressions. To say that such sentences are *semantically* complex is just to say that the meaning of these sentences somehow results from the meanings of their constituent parts. Roughly speaking, that's what compositionality means: the meaning of complex expressions is determined by the meanings of their constituent parts. Meanings here will of course be the semantic properties that the semantic theory describes – in the case of sentences we were operating with the preliminary idea that these meanings are centrally propositions. We will soon be concerned also with the semantic properties of subsentential parts.

Compositionality

The assumption that natural languages are compositional is in fact a somewhat controversial assumption.[24] However, compositionality provides a commonsensical explanation for an obvious fact, namely the fact that speakers of natural languages are able to construct and understand completely novel sentences. Many, and probably most, of the sentences you have been reading in this text, are sentences you, the reader, have ever encountered before. Nor have I, the author, ever encountered them before writing them down. Nevertheless, we understand these sentences, we are competent at using them. The obvious and commonsensical explanation of this fact is that competent speakers know (in some sense) the semantic properties of the familiar parts from which these sentences are constructed, and they also know how these properties, together with the exact manner of composition, determine the meaning of these novel sentences. Thus, the parts of novel sentences are not novel, nor are the ways of putting the parts together and the impact this has on the resulting novel sentence. Language users (or perhaps their sub-personal language faculty) can work out the meanings of the novel sentences from what they already know. So let us assume, for now, that compositionality is a genuine phenomenon that our semantic theory ought to account for.

24 See, for example, Travis 1997, Lahav 1989, Recanati 2004; but see also Pagin & Westerståhl 2010a, 2010b.

Modelling the compositional structure of languages is an area where formal semantics excels, and it is one of the primary motivations for natural language semanticists to turn to the methods of formal logic in the first place. Already the grandfather of modern logic, Frege, adhered to a methodological principle of compositionality. He did so not only in his formal system, but when he did discuss natural language, he brought the principle to bear on it (see especially Frege 1892). Relatedly, Frege is also the originator of one key fundamental idea in modern semantics, namely the idea of construing the meanings of expressions as *functions*.

Functions

How does this fundamental idea work? Functions are studied in several areas of mathematics, in particular in set theory, but for our purposes it is enough to say that a function is something that assigns a unique *value* from a certain *range* to each entity (*argument*) from a certain *domain*. Thus, consider the function f such that $f(x)=x^2$.[25] Let's take as its domain the natural numbers, and as its range also the natural numbers. The value of this function for the argument 2 is 4, i.e. f(2)=4. Similarly f(3)=9, f(1)=1, f(25)=625 and so on. One can think of functions as sets, namely sets of ordered pairs. So for example, our function f can be thought of as the set of ordered pairs such that the first element of each pair is a natural number and the second element is that number squared.

Applying this idea to semantics, we can think of complex expressions as being the concatenation of one expression *e*1, standing for a function, with another expression *e*2, standing for a possible argument of that function, i.e. a member of its domain. We take the complex expression *e*1^*e*2[26] to function rather like "the value of *e*1 for the argument *e*2" or "*e*1(*e*2)". If we think of the meanings of expressions generally as appropriate functions, then we can say that the meaning of any complex *e*1^*e*2 is always the value of the meaning of *e*1 for the meaning of *e*2 as argument. This would give us a very simple general rule for how the meaning of a complex expression is determined by the meanings of its constituent parts and their manner of composition. The only manner of composition that is in play here is that of concatenation.

A single compositional rule: Concatenation as function application

Let me repeat this last point, because it is important: the idea is that we have a general compositional rule, i.e. one rule that says how the meaning of a complex expression is determined by the meanings of its parts. This rule says that if we write one expression after another, then the meaning of the resulting complex will be the value of the function that is the mean-

25 For "$f(x)=x^2$" read: "the value of f for an argument *x* is x^2", or for short: "f of *x* is x^2".

26 "*e*1^*e*2" is short for: "the string formed by prefixing *e*2 with *e*1" or "the expression formed by concatenating *e*1 and *e*2 in that order".

ing of the first expression for the the meaning of the second expression taken as argument. So, if we generally write '*m*(*e*)' for the meaning of *e*, then the idea is that the meaning of a complex expression *e*1^*e*2 is always *m*(*e*1)(*m*(*e*2)). We could build a more complex expression by concatenating *e*1^*e*2 with another expression *e*3: *e*3^*e*1^*e*2.[27] The meaning of '*e*3^*e*1^*e*2' is now simply the value of the meaning of *e*3 for the meaning of '*e*1^*e*2' as argument. This is the fundamental idea on which semantics (in the tradition of Frege, Carnap, Montague) builds.[28] As long as we find appropriate functions that can figure as the meanings of simple expressions and deliver the right functions as the meanings of complex expressions, we will have a compositional semantic theory, for clearly a function determines for each argument in its domain a unique value. The meanings (i.e. functions) that the semantics assigns to expressions in this manner are often called "semantic values". This nomenclature reflects the idea that those meaning features described by assigning "semantic values" to expressions are compositional, i.e. are determined by the semantic values of their constitutive parts and their manner of composition.

The compositional rule I just considered was only one possible very simple way of construing a compositional rule: SV(*e*1^*e*2) = SV(*e*1)(SV(*e*2))[29]. It could be reversed: SV(*e*1^*e*2) = SV(*e*2)(SV(*e*1)), for example. Thus, to give a rough example, we could construe the semantic value of "Sam smokes" as the value of the semantic value of "Sam" when applied to the semantic value of "smokes" as argument. Or we could construe the semantic value of "Sam smokes" as the semantic value of "smokes" when applied to the semantic value of "Sam" as argument. In each case we may have to assign different semantic values to "Sam" and "smokes", so that the semantic values in question are functions of the right kind, i.e. are defined for the relevant arguments and deliver the right kinds of semantic values as values for those arguments. These are decisions that need to be made by looking at the actual syntactic and semantic properties of natural languages, and sometimes there may be different compositional rules that work equally well. The test whether an assignment of SVs works will always be the testable predictions that the theory ultimately makes. I have

27 I am assuming that concatenation is "left-associative", i.e. that '*e*1^*e*2 … ^*en*' is the result of applying the left-most expression *e*1 to the remainder '*e*2 … ^*en*', and so on. We can use brackets to make this more perspicuous: in '*e*3^(*e*1^*e*2)' the brackets indicate that *e*1 and *e*2 were first concatenated in that order, and then *e*3 was applied to the complex expression *e*1^*e*2 (rather than *e*3 first being concatenated with *e*1, and *e*3^*e*1 then being prefixed to *e*2, which would be the expression (*e*3^*e*1)^*e*2.

28 See Lewis 1970 for an excellent discussion of the fundamentals of this approach.

29 Instead of a meaning function m, I am now speaking of a semantic value function SV.

proposed that these predictions should be construed as predictions about the literal truth of utterances. Thus we know that the SVs we ultimate choose must deliver such predictions for sentences, namely predictions as to the conditions under which an utterance of any given sentence is literally true.[30]

Alternative composition rules

There could be more complex compositional rules. For example, there might be rules that concern the concatenation of three expressions of specified syntactic categories, e.g. SV(*e*1^*e*2^*e*3) = SV(*e*2)(SV(*e*1), SV(*e*2)). According to this rule, when *e*2 is in a certain syntactic category and it is preceded by an expression *e*1 of a certain category and succeeded by an expression *e*3 of a certain category, then the semantic value of the resulting expression is the value of the semantic value of *e*2 for the argument that is the ordered pair of the semantic value of *e*1 and the semantic value of *e*3. This might be the intuitively correct type of rule for natural language transitive verbs such as “loves”: the semantic value of “Sam loves Pam” is the value of the semantic value of “loves” for the ordered pair <SV(“Sam”), SV(“Pam”)> as argument. Again, whether our semantic theory should use a compositional rule of this form, or whether it should rather view *e*1^*e*2^*e*3 as involving two compositional operations, e.g. *e*2^*e*3 and then *e*1^(*e*2^*e*3), will be a question that needs to be decided on the basis of general considerations of predictive adequacy, theoretical elegance and fruitfulness. The *basic* functional approach originated by Frege, i.e. that SVs are viewed as appropriate functions and arguments, leaves all these details open.

Interaction between syntax and semantic value assignment

What functions should we use as the semantic values of expressions, and which compositional operations should we account for? This is where the real work begins, and it is also what most discussions in semantics are ultimately concerned with. Let us look at some of the basic moves. First, it seems clear that not every expression can be concatenated with every other expression. Some strings of expressions are syntactically malformed. While “George is dancing a waltz.” is a sentence and is well-formed, “is waltz a George dancing” is not. So, the functions that serve as the semantic values of our expressions need only have in their domains the functions that can figure as the semantic values of expressions with which they can be concatenated as argument expressions. Similarly, the semantic values of expressions need to be in the domain of the functions that are the semantic values of those expressions with which they can be syntactically conbined as functional expressions. Thus we will have to have a syn-

30 As Tugendhat 1970 argues, this is exactly Frege’s idea of “Bedeutung”, i.e. reference: the Bedeutungen of subsentential expressions are exactly their potential to contribute to the truth-values of the sentences in which they can occur, i.e. their “truth-value potential”.

tax that tells us which expressions can be combined with which other expressions, and our assignments of meanings (i.e. semantic values) will need to make provision only for these permitted combinations.

Categorial grammar

One approach that makes the semantics quite easy and systematic is that pursued in the description of *categorial* or *typed* languages. On this approach, there are a number of basic syntactic categories, for example the category of names, N, and the category of sentences, S. Expressions in these categories will then have certain semantic values, e.g. sentences might have truth-values as SVs, and names might have individual objects as SVs. All other categories are derived: for example the category of those expressions that combine with names to form sentences would be called "S/N", the category of expressions that form S/N expressions from names would be called "(S/N)/N", and the category of expressions that combine with sentences to form sentences would be called "S/S". We can now describe systematically what SVs the expressions in the derived categories would need to have: an expression in S/N needs to have a SV that is a function from SVs of names, i.e. objects, to the SVs of sentences, i.e. truth-values. An expression from (S/N)/N needs to have a SV that is a function from SVs of names, i.e. objects, to SVs of expressions in S/N – which, as we just saw, are functions from objects to truth-values. The general principle is that an expression of category α/β is a function from the SVs of expressions in β to the SVs of expressions in α. Thus, there are categories of SVs that correspond exactly to the syntactic categories.

A left-assocaitive composition rule is too restrictive

A simple single compositional rule like the one I considered first: SV(*e*1^*e*2) = SV(*e*1)(SV(*e*2)), is left-associative: it allows only complexes where the a functional expression is *prefixed* to a corresponding argument expression, so that the semantic value of the whole is the value of the semantic value of the prefixed expression for the semantic value of the other expression taken as argument. This is very restrictive. A language that works according to these principles will have complex expressions that always begin with the principal functional expression. Thus, if a natural language was treated in this way, we would be forced to say, that the sentence "John kissed Mary", has as its SV the SV of "John" for the SV of "kissed Mary" as argument, and that the SV of "kissed Mary" is the SV of 'kissed" for the SV of "Mary" as argument. So far, this is OK: if "Mary" and "John" are in N, "kissed" is in (S/N)/N, then "kissed Mary" is in S/N, and the whole sentence is in S. The SVs would be accordingly. But now consider the sentence "John kissed Mary passionately". We would have to start by saying that its SV is the value of the SV of "John" for the SV of "kissed Mary passionately" as argument etc. In the end we would have to say that "Mary" has as its SV some function that can take the SV of "passionately" as argument. If we continue to say that "kissed" is in (S/N)/N, then "Mary passionately" would need to be a name. This would mean that "Mary" can't

be in the category of names. Rather, “Mary” would be in the category of N/*x*, where *x* is “passionately”’s category. Presumably, this argument could be repeated for any bona fide name. It could also be repeated for our newly discovered name “Mary passionately”. For there are sentences like “John kissed Mary passionately in Paris”.

Exercise 4*: if the SVs of expressions in N are objects and the SVs of expressions in S are truth-values, then what kind of function would be the SV of “kissed Mary” and of “kissed”?*

While there are formal languages (e.g. Polish notation in logic) for which the simple compositional rule is appropriate, it is not appropriate for natural languages. So there must be syntactic operations beyond the mere stringing up of expressions, so that we can allow, for example, that “passionately” is a (S/N)/(S/N) and that in “John kissed Mary passionately” it operates on the S/N “kissed Mary”. This means that before assigning a semantic value to a natural language sentence, we first need to provide an analysis of the underlying syntactic structure. Along these lines, we might think of the surface sentence “John kissed Mary passionately” as having the *logical form*

John(passionately(kissed (Mary)))

On this conception, a semantic theory operates on the logical form of a sentence, and this logical form first needs to be recovered from the surface form, which is thought of as being generated via certain transformational rules from the logical form. The semantics is made much easier if it operates only on syntactically perspicuous logical forms.[31] Of course this means that we have to postulate an implicit extra level of syntactic representation, which is a disadvantage. On the other hand, there are also certain phenomena of (structural) ambiguity, which can be explained easily on this model. For ambiguity arises, on this model, when the transformation of two different logical forms generates the same surface form, as for example in this case:

(1) bancrupt(business(consultant))

(2) bancrupt(business)(consultant)

Logical form vs surface form

Both logical forms will be transformed into the same surface compound noun: “bancrupt business consultant”, but the first will apply to business

[31] Those who argue for a need for underlying logical forms will often invoke phenomena of binding, for which it is convenient to postulate variables that are not visible on the surface of sentences. More on this below.

consultants who are bancrupt, while in the second to consultants of bancrupt businesses.

Thus, we now have the following picture: A semantics assigns a semantic value to each sentence in a disambiguated base language, i.e. to each logical form. These logical form sentences are transformed via transformational rules into the surface sentences that are actually used by language users, and can be recovered, if necessary by disambiguation, from this surface sentence.

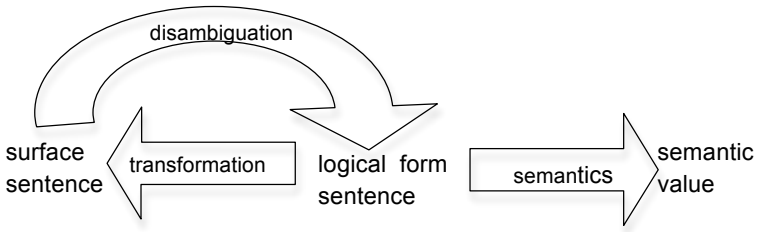

For the moment, we will not worry about the transformational component of a semantic theory, nor will we be very ambitious when it comes to finding the right syntactic categories. Rather we will begin by considering a very simple type of language that can serve as the base language on which the semantics operates. When we ask ourselves what kind of functions can serve as the semantic values of expressions, we can start with one constraint that I have introduced a while back. I said we are looking for a semantic theory that predicts the conditions under which utterances of sentences are literally true. Thus we already know that at the very least the semantic value of a sentence needs to determine conditions under which an utterance of the sentence is literally true. Earlier, I spoke of propositions, but I didn't specify further, what propositions are, except that they are constituted by concepts. Let us therefore begin with the simplest way in which a semantics might predict the truth values of (utterances of) sentences, namely by saying that the semantic values of sentences simply *are* truth-values.

3.4 Extensional Languages

A simple language of first order logic will do just this. I shall here introduce a simple language with restricted quantifiers, introducing variables at a later point.[32] We will have three underived categories: a category of N of

32 I am not starting with a standard first-order language with universal and existential quantifier simply in order to be a little bit closer to natural language grammar.

names, a category CN of common nouns and a category S of sentences. The SVs of names are objects from a domain O of objects, the SVs of common names are sets of objects from O, and the SVs of sentences are truth-values from the set TV={0, 1}. The semantic values of other, derived categories are functions that can deliver the SVs of these basic categories. The semantics will consist in providing a specification of all atomic expressions of the language, together with their category, and assigning a semantic value to each.

I shall mention expressions of L1 by putting them in bold. Another convenient abbreviation I shall be using is this.[33] The SVs I'll be assigning to expressions will correspond to their syntactic category. An expression of syntactic category S/N will need to have a SV which is a function from the SVs of expressions in N to the SVs of expressions in S. The SVs of expressions in S/(S/N) will need to be functions from the domain of SVs of expressions in S/N (as just described) to values that are SVs of expressions in S. Generally an expression of category β/α will need to have a SV which takes SVs of expressions in α as arguments and delivers values that are SVs of expressions in β. In order to be able to talk about sets of functions in an easier way, we can adopt the following definition of "semantic category":

(D) (a) The set O of objects is a semantic category
(b) The set TV of truth-values {1, 0} is a semantic category
(c) The set E of subsets of O, i.e. $\{x: x \subseteq O\}$, is a semantic category
(d) For all semantic categories β, α: $<\alpha\rightarrow\beta>$ is a semantic category
(e) Nothing else is a semantic category
(f) For all semantic categories β, α: $<\alpha\rightarrow\beta>$ is a set of functions from the domain α to the range β.

Thus, we can now conveniently say that the SVs of expressions in S/N are in the semantic category <O→TV>, i.e. they are functions from objects to truth-values. And we can say that the SVs of expressions in S/(S/N) are in the semantic category <<O→T>→TV>, i.e. they are functions from functions from objects to truth-values to truth-values. And so on.

The following is a description of a simple extensional language L1.

33 The abbreviation is adapted from Heim and Kratzer 1998, p. 28.

Lexicon of L1:

N:

Sue: SV(**Sue**)=Sue (Sue ∈ O)

Peter: SV(**Peter**)=Peter (Peter ∈ O)

…

S/N:

dances: SV(**dances**) = the function $f \in$ <O→TV>, such that $f(x) = 1$ iff x dances and $f(x) = 0$ otherwise. [34]

frowns: SV(**frowns**) = the function $f \in$ <O→TV>, such that $f(x) = 1$ iff x frowns and $f(x) = 0$ otherwise.

…

CN:

man: SV(**man**) = the set of men

woman: SV(**woman**) = the set of women

…

(S/(S/N))/CN:

every: SV(**every**) = the function $f \in$ <E→<<O→TV>→TV>>, such that $f(x)$ = the function $g \in$ <<O→TV>→TV>, such that $g(y) = 1$ iff for all $z \in$ SV(x), $y(z) = 1$.

some: SV(**some**) = the function $f \in$ <E→<<O→TV>→TV>>, such that $f(x)$ = the function $g \in$ <<O→TV>→TV>, such that $g(y) = 1$ iff for some $z \in$ SV(x), $y(z) = 1$.

…

S/S:

not: SV(**not**) = the function $f \in$ <TV→TV>, such that $f(x) = 1$ iff SV(x) = 0.

…

(S/S)/S:

and: SV(**and**) = the function $f \in$ <TV→<TV→TV>>, such that $f(x)$ = the function $g \in$ <TV→TV>, such that $g(y) = 1$ iff SV(x) = SV(y) = 1.

L1 allows us to derive the semantic values of sentences like **not**(**frowns**(**Sue**)) or **some**(**man**)(**dances**), at least if we have information about whether Sue frowns and about whether there is a member of O who is a man and also dances.

[34] This notation is short for: "the function f from the domain O to the range {0, 1}, such that $f(x) = 1$ iff x dances and $f(x) = 0$ otherwise".

3.5 Adding Variables and Binding

Variables

Even though L1 contains quantifiers, it contains no variables, so as things stand, L1 does not permit the binding of variables. It is not obvious that natural languages contain variables. However, some inferential properties of natural language sentences have been taken to be manifestations of the implicit presence of variables at the level of deep structure or logical form. In order to introduce binding, we will need to do two things: first, we need to introduce variables into the category of names. We will introduce variables with numerical subscripts $\boldsymbol{x_1}$, $\boldsymbol{x_2}$, … $\boldsymbol{x_n}$. The point of variables is to enable us to consider the SVs of expressions containing variables under various different possible assignments of SVs to the variables. So we will have to relativize the function assigning SVs to assignments A of values to variables. An assignment is simply a sequence of members of O, i.e. $<o_1, o_2, \ldots o_n>$, such that $o_1, o_2, \ldots o_n \in$ O. We will then say that the SV of a variable $\boldsymbol{x}_i$ relative to to an assignment $a \in$ A is the object in ith position in a, i.e. $a(i)$. As the SVs of other expressions will be affected by this relativity to assigments of variables, we'll generally assign SVs relative to assignments. Thus our semantic value function will now be two-place: SV(e, a) will be the SV of expression e for the assignment a. For those e that do not contain any variable, the assignment argument a will be inert, and SV(e, a) will be the same as SV(e, b) for any assignments a and b.

Lambda operators

The second modification we need to make has to do with the fact that our quantifier phrases, e.g. **some**(**woman**) are complex expressions in the category S/(S/N). We do now have formulae containing variables, e.g. **frowns**($\boldsymbol{x_2}$), but these are in the category S, so a quantifier phrase from category S/(S/N) cannot be applied to them. To remedy this, we will introduce new expressions in category (S/N)/S, which can turn sentences into predicates. They are called "lambda operators". For example, $\boldsymbol{\lambda x_1}$ can combine with any sentence ϕ to form the predicate $\boldsymbol{\lambda x_1}(\phi)$; $\boldsymbol{\lambda x_2}$ can combine with any sentence ϕ to form the predicate $\boldsymbol{\lambda x_2}(\phi)$; and so on. For each variable $\boldsymbol{x}_i$, there is a lambda operator $\boldsymbol{\lambda x}_i$ in category (S/N)/S. One can read "is an $\boldsymbol{x_1}$ such that ϕ" for $\boldsymbol{\lambda x_1}(\phi)$. For example, read "is an $\boldsymbol{x_1}$ such that **frowns**($\boldsymbol{x_1}$)" for $\boldsymbol{\lambda x_1}$(**frowns**($\boldsymbol{x_1}$)). The SV of a lambda expression $\boldsymbol{\lambda x_2}(\phi)$ will be bound up with the SV of the sentence ϕ in question. The case we are interested in is the case where ϕ contains $\boldsymbol{x_2}$, so that SV(ϕ, a) might vary with the assignment parameter a. At an assignment a, $\boldsymbol{\lambda x_2}(\phi)$ will have as its SV the function from objects to truth-values that assigns 1 to all objects x such that, for all assignments b that differ from a at most in 2nd place, if $x = b(2)$ then SV(ϕ, b) = 1. This ensures that for any assignment a and any α in S/N, SV(α, a) = SV($\boldsymbol{\lambda x_2}(\alpha(\boldsymbol{x_2}))$, a).

***Exercise 5**: Actually, this suggests that we should treat these operators as complex, so that λx_1 results from applying λ to x_1, that λx_2 results from applying λ to x_2. Question: in which category is λ?*

A third modification will add transitive verbs, such as **loves** and **admires** to the lexicon. These will be in category (S/N)/N.

Giving up the simple compositional rule

Until now I have religiously followed the Fregean idea that there is just one recursive compositional rule: for all expressions *e*1 and *e*2, SV(*e*1(*e*2)) = SV(*e*1)(SV(*e*2) (i.e. that the semantic value of a complex expression is always the value of the semantic value of the one for the semantic value of the other taken as argument). This had the advantage that the semantics consisted merely in assigning SVs to all the atomic expressions, everything else followed from the compositional rule. If, as I am now proposing, we add variables, the price for sticking to a simple lexicon of atomic expressions and one recursive compositional rule is considerable presentational complication. Any SV in $<\alpha\rightarrow\beta>$, we have had so far, will now need to be a SV in $<A\rightarrow<\alpha\rightarrow\beta>>$, where A is the set of assignments. For the sake of presentational ease, I will therefore simplify from now on, by allowing several recursive compositional rules.

Giving up assigning SVs to every expression

This means that I'll start by listing expressions in categories N, S/N, (S/N)/N and CN together with their syntactic category and an assignment of a semantic value to each of them. After that, I shall articulate a number of recursive clauses that specify for various syntactic combinations, what their effect on semantic value is. These recursive clauses will also in part introduce new lexical expressions whose semantic effect is given only by the effect they have on complex expressions in which they occur.

<u>Syntax of L2:</u>

Vocabulary of atomic expressions of L2:

Names:

constants:	**Sue**, **Peter**, …
variables:	x_1, x_2, … x_n
n-place predicates:	
1-place predicates:	**dances**, **frowns**, …
2-place predicates:	**loves**, **admires**, …
…	
Common nouns:	**man**, **woman**, …
Determiners:	**every**, **some**, …

1-place connectives: **not**, …

2-place connectives: **and**, …

Binders: λx_1, λx_2, … , λx_n

Definition of complex expressions of L2:

Determiner phrases:

1. If α is a determiner and β is common noun, then '$\alpha(\beta)$' is a determiner phrase. Nothing else is a determiner phrase.

Sentences:

2. If α is an *n*-place predicate and β_1, β_2, … , β_n, are names or quantifier phrases, then '$\alpha(\beta_1, \beta_2, \ldots, \beta_n)$' is a sentence.
3. If α and β are sentences, then '**not**(α)' and 'α **and** β' are sentences.
4. Nothing else is a sentence.

Complex 1-place predicates:

5. If α is a binder and β is a sentence, then '$\alpha(\beta)$' is a 1-place predicate. The only 1=place predicates are the ones mentioned under the atomic ex pressions of L2 and the complex 1-place predicates constructed according to
6. There are no further expressions in L2.

<u>Semantics of L2:</u>

The set of objects O={o_1, o_2, … o_n}

The set of assignments A={<a_1, a_2, … a_n>: a_1, a_2, … $a_n \in$ O} (the set of sequences made up from members of O)

The set of truth-values TV={0, 1}

We define a function SV(*e*, *a*):
Domain of SV: *e* is either a name, an *n*-place predicate, common noun or sentence of L2. $a \in$ A. [35]
Range of SV: O $\cup$ $\wp$(O) $\cup$ O×O $\wp$(O×O) $\cup$ TV

Semantic values of atomic expressions:

Constant names:

1. **Sue**: For all $a \in$ A, SV(**Sue**, *a*)= o_1 ($\in$ O)
2. **Peter**: For all $a \in$ A, SV(**Peter**, *a*)= o_2 ($\in$ O)

[35] Note that the semantic value function does not assign semantic values to all expressions. The contribution of other expressions to the SVs of complex expressions that do have semantic values is specified in some of the recursive clauses. This is part of the price for expository convenience.

Variable names:

3. $\boldsymbol{x_1}$, $\boldsymbol{x_2}$, … $\boldsymbol{x_n}$: For all i such that $1 \leq i \leq n$, and for all $a \in$ A, SV($\boldsymbol{x_i}$, a)= $a(i)$ [i.e. the object in ith position in assignment a]

1-place predicates:

4. **dances**: For all $a \in$ A, SV(**dances**, a) = {o_1, o_{12}, o_{17}, o_{18}, … } ($\subset$ O)

5. **frowns**: For all $a \in$ A, SV(**frowns**, a) = { o_2, o_3, o_4, o_{18}, o_{279}… } ($\subset$ O)

…

2-place predicates

6. **loves**: For all $a \in$ A, SV(**loves**, a) = {< o_2, o_1>, <o_{17}, o_1>, … } ($\subset$ O×O)

7. **admires**: For all $a \in$ A, SV(**admires**, a) = {< o_2, o_1>, <o_1, o_2>, … } ($\subset$ O×O)

…

Common nouns:

8. **man**: For all $a \in$ A, SV(**man**, a) = {o_2, o_7, o_9, o_{77}, … } ($\subset$ O)

9. **woman**: For all $a \in$ A, SV(**woman**, a) = {o_1, o_5, o_8, o_{77}, … } ($\subset$ O)

…

Semantics for complex expressions of L2 (recursive part):

Sentences:

10. If α is a 1-place predicate and β is a name, then for all $a \in$ A: SV('α (β)', a) = 1 iff SV(β, a) $\in$ SV(α, a).

11. If α is a 2-place predicate and β and γ are names, then for all $a \in$ A: SV('α (β, γ)', a) = 1 iff <SV(β, a), SV(γ, a)> $\in$ SV(α, a).

12. If α is a common noun and β is a 1-place predicate, then for all $a \in$ A: SV('**every (**α) (β)', a) = 1 iff for every $x \in$ O: if $x \in$ SV**(**α, a), then $x \in$ SV**(**β, a).

13. If α is a common noun and β is a 1-place predicate, then for all $a \in$ A: SV('**some (**α) (β)', a) = 1 iff for some $x \in$ O: $x \in$ SV**(**α, a), and $x \in$ SV**(**β,a).

14. If α is a sentence, then for all $a \in$ A: SV('**not**(α)', a) = 1 iff SV(α, a) = 0.

15. If α and β are in S, then for all $a \in$ A: SV('α **and** β', a) = 1 iff SV(α, a) = 1 and SV(β, a) = 1.

…

Complex 1-place predicates:

16. $\boldsymbol{\lambda x_i}$: If α is a sentence, then for all $a \in$ A and for all i: SV('$\boldsymbol{\lambda x_i}$(α)', a) = {x: for all $b \in$ A that differ from a at most in i-th place: if $x = b(i)$ then SV(α, b) = 1}

In L2, semantic values are assigned relative to assignments of objects to variables. This was newly introduced in order to be able to introduce bound variables. We are hoping to predict the truth-values of sentences, so there is a question as to what our assignments of the value 1 to sentence-assignment pairs mean for our predictive ambitions. The answer is

simple: any *closed* sentence, i.e. any sentence that contains no free variable, will have the same truth-value for all assignments. So we can define:

(T) For all sentences α of L2: α is True iff for all $a \in$ A, SV(α, a) = 1.

(F) For all sentences α of L2: α is False iff for all $a \in$ A, SV(α, a) = 0.

Definition of Truth and Falsity in L2

According to these definitions, sentences containing unbound variables will come out as neither True nor False. However, for the moment the idea is that our proposed empirical data, i.e. users' assessments for literal truth, concern utterances of closed sentences only – sentences containing free variables are simply incomplete and not usable in communication. We will reconsider this point later.

3.6 The Need for Intensional Semantics

L2 is extensional

Even though the notion of an "extension" has not figured in the description of the semantic value of expressions above, L2 is in effect "extensional", in the sense that the truth-value of a sentence will not change if one expression is substituted for another expression that has the same extension. The "extension" of a name is the object that is its SV for all assignments (its "referent"), and the extension of an *n*-place predicate *F* is the set of *n*-tuples that are members of the semantic value of *F* for all assignments (intuitively the tuples to which the predicate applies). The extension of a complete sentence (a sentence that contains no unbound variables) is a truth-value. Modelling natural languages on extensional first order languages delivers our main desideratum, namely the prediction of truth-values for sentences on the basis of the SVs of the sentences' constituents, but it also faces a number of basic problems, i.e. it cannot model certain natural language phenomena.

Problem: all truth-vales can be deduced from semantic information

One problem is that the truth-values of all sentences can be *deduced* from the semantics alone. If we know the semantic values of all the expressions, all that separates us from knowing the truth-values of sentences is a deduction. This would reduce the point of testimony amongst fully competent language users to the sharing of a priori deductive knowledge. Now, one way to alleviate this point is to articulate the semantic values in a non-transparent way. As it is, the semantics for L2 lists the *n*-tuples of objects to which a predicate applies (the predicate's extension) in such a way that it can be deduced whether, for example, the SV of a certain name is an element in any of these *n*-tuples – in this sense, it is transparent. We might instead specify membership of extensions in a non-transparent way by articulating (in the language of the semantic theory) a *condition* that is necessary and sufficient for an *n*-tuple to be in the extension of a predicate.

Thus, our specification of the extension of “loves” could just be: “the set of ordered pairs <o_1, o_2> such that o_1 loves o_2. We cannot *deduce* the truth-value of sentences containing “loves” from this knowledge.

An analogous problem arises for the treatment of logical truth. A sentence is logically true just if its logical form alone – i.e. the quantifiers and sentence connectives as they occur in it, independently of the semantic values of any non-logical vocabulary in it – determines that the sentence is true. On the current semantics, the predictions as to literal truth of sentences will not distinguish between any two logically true sentences. Similarly for logical falsehoods. However, is it really true that competent users tend to accept all logically true sentences as literally true (and tend to reject all logically false ones)? Presumably, some logical truths (falsehoods) are not easily recognized as such, so the relevant data – judgements of literal truth – will vary with respect to them. We could deal with this by further idealizing our data set and utilize only the judgements of literal truth of competent users *after ideal reflection* or something of the sort. But it seems clear that this would abandon certain predictive ambitions from the start.[36]

Problem: competence is more than knowledge of literal truth-value

Another, arguably more serious, shortcoming may have struck you from the start. Initially, we were looking for a semantics that delivers the conditions under which sentences are literally true, for we thought that competent language users would be reasonably good at recognizing these conditions due to the fact that they have acquired a sensitivity to norms of correct use that track the norms to which they are subject. But this ability of users to recognize the conditions under which a sentence is literally true seems to involve the ability to tell situations or states of affairs in which an utterance of the sentence is literally true from those where it is not. Or perhaps more modestly: they know which features of the world would be relevant for assessing the utterance of a sentence for literal truth. Thus, what the semantics should deliver for sentences is not literal truth-values, but *conditions* for literal truth. So a more promising sort of semantic value for sentences would be functions from ways things are or states of the world to truth values. Similarly for subsentential expressions like predicates: it seems that what competence involves is some knowledge of what it takes to be in the extension of the predicate, what conditions an object (or *n*-tuple of objects) needs fulfil to be in the extension.

Problem: intensional contexts

It seems that what competent speakers can recognize is not simply whether an utterance of a sentence is literally true. Rather, they can recognize

36 See Stalnaker 1987, especially chapter 5, for a defence of the view that all logically true sentences and beliefs have the same content.

whether the utterance would have been true if circumstances had been a certain way. They need not know whether an utterance of "Jack admires Sue." is true, but they know what it would take for it to be true. This dependence of the truth-value of a sentence on the way things are, i.e. on the circumstances, also seems to be needed to account for the behaviour of certain kinds of expressions, namely so-called "intensional contexts". The best known intensional contexts are belief ascriptions and modal contexts. If we wanted to add to our language expressions that are to model natural language expressions like "believes (that)" or "It is possible that" (in the sense of "It could have been the case that"), we would have the following problem. Both "believes (that)" and "It is possible that" seem to be of a syntactic category that allows them to be concatenated with *sentences*. "It is possible that" prefixed to a sentence, yields another sentence, as in "It is possible that Jack admires Sue". "Believes that" prefixed to a sentence yields a predicate, which again prefixed with a singular term yields a sentence, as in "John believes that Jack admires Sue.". However, if the semantic values of sentences are merely truth-values (or functions from assignments to truth-values, but *constant* such functions in the case of sentences containing no free variables), then it would seem that the semantic values of these complex sentences should be sensitive only to the truth-value of the embedded sentence (i.e. the one to which we have prefixed the intensional expression). Thus if the sentence "Jack admires Sue" receives the value True, and the sentence "John believes that Jack admires Sue" also receives the value True, then any sentence of the form 'John believes that *s*.', where *s* is another true sentence, should also be true. And if the sentence "Sue admires Jack" is false, and "It is possible that Sue admires Jack." is true, then any sentence of the form 'It is possible that *s*', where *s* is a false sentence, should also be true. But this is not so: The sentence "Jack admires Sue and Jack does not admire Sue." is false, yet the sentence "It is possible that (Jack admires Sue and Jack does not admire Sue)" is not (literally) true. Similarly, it is not difficult to find two true sentences *s* and *r*, such that 'John believes that *s*' is true, yet 'John believes that *r*' is false. Thus we either have to give up the idea that these expressions concatenate with sentences as semantic constituents,[37] or we have to look for alternative semantic values for sentences.

Intensions, possible worlds

Intensional languages remedy the situation by viewing the semantic values of sentences as "intensions" – a term introduced by Carnap (1956) for his successor notion to Frege's notion of Sense. One way of construing a sentence intension is as a function from possible states of the world to truth-values. But any other kind of entity that determines such a function would also do as a semantic value for sentences. Thus some theorists will

37 As does Davidson in his "On saying 'that'".

operate with *propositions* that are structured entities constituted by concepts, as I did in chapter 2. The important feature of the reformed semantic values for sentences (on this construal) will be that they *determine* functions from possible states of the world to truth-values.[38] I will speak simply of "possible worlds" or just "worlds", where each world is a way things might have been. The actual world is the way the world really is. I shall remain neutral on the question whether non-actual worlds are concrete entities (Lewis 1986), whether they are just properties the one and only real and concrete world might have had (Stalnaker 1976, 1984), or whether they are simply (and minimally) maximally consistent sets of sentences (Carnap 1956). The important feature of worlds in our model is that they each represent a maximally specific way one might believe things to be. Or perhaps less ambitiously: anything one might believe corresponds to a partition of the set of possible worlds into the ones which are as that belief represents things as being and those that are not.

Equivalence between an assignment of intensions to sentences and an assignment of extension to sentence-world pairs

Below is what an intensional semantics for a language L3, similar to L2 might look like. One feature of the particular notation I am using may be confusing, even though it is a completely insubstantial presentational matter. I have been speaking of semantics assigning semantic values to sentences, and I have now been speaking of an intensional semantics that assigns intensions to sentences, i.e. functions from possible states of the world to truth-values. Now, on the face of it, this might lead one to expect that the semantics is basically a function SV, which has as it's domain all the expressions of L3 and which has values from a range that contains intensions. So, if I is the set of intensions, and E is the set of expressions, then one would expect a semantic value function SV: E→I, so that for example, the value of SV for "Sue frowns", SV("Sue frowns"), would be the intension that assigns the value True to any world at which Sue frowns. However, in the formal semantics for L3 below, I have instead defined a function SV from sentence-assignment-world triples to truth-values (or other types of extension for expressions that are not sentences). Thus, the semantics is assigning extensions to expression-assignment-world triples <*s*, *a*, *w*>. This makes no substantial difference. From a function *f* that assigns extensions *e* to triples < *s*, *a*, *w*>, we can easily reconstruct another function *f** that assigns triples <*a*, *w*, *e*> to expressions *s*.[39]

38 There are also related approaches, which construe sentence intensions as functions not from possible states of the world but from other entities, such as for example epistemically possible states of the world, as in Chalmers 2002.

39 See Lewis 1980 for some dicussion of this. Lewis speaks of variable but simple vs constant but complex semantic values. While Lewis is already one step ahead of us, because he is considering languages that are not only intensional but also exhibit context dependence, the same considerations apply already for us here.

Below is a characterization of L3:[40]

Syntax of L3:

Vocabulary of atomic expressions of L3:

As for L2 with three additions:

Modal operators: **possibly**, **necessarily**, …

Propositional attitude verb: **believes that**

Definition of complex expressions of L3:

As for L2, with two additions:

Sentences:

3a. If α is a sentence, then '**necessarily**(α)' and '**possibly**(α)' are sentences.

3b. If α is a name and β is a sentence, then 'α **believes that** β' is a sentence.

Semantics of L3:

The sets O, A and TV are as in the semantics for L2.

The set of possible worlds W={ w_1, w_2, … w_n}

We define a function SV(*e*, *a*, *w*):

Domain of SV: *e* is either a name, an *n*-place predicate, common noun or sentence of L3. $a \in$ A. $w \in$ W.

Range of SV: O $\cup$ $\wp$(O) $\cup$ O×O $\wp$(O×O) $\cup$ TV

Most clauses are just modified from L2 by adding the new world dependence of SV. But 14a., 14b. and 14c. are additions.

Semantic values of atomic expressions:

Constant names:

1. **Sue**: For all $a \in$ A, for all $w \in$ W, SV(**Sue**, *a*, *w*)= o_1 ($\in$ O)

2. **Peter**: For all $a \in$ A, for all $w \in$ W, SV(**Peter**, *a*, *w*)= o_2 ($\in$ O)

Variable names:

3. $\boldsymbol{x_1}$, $\boldsymbol{x_2}$,… $\boldsymbol{x_n}$: For all *i* such that $1 \leq i \leq n$, for all $a \in$ A and for all $w \in$ W: SV($\boldsymbol{x_i}$, *a*, *w*)= *a*(*i*) [i.e. the object in *i*th position in assignment *a*]

1-place predicates:

4. **dances**: For all $a \in$ A, for all $w \in$ W: SV(**dances**, *a*, *w*) = {*x*: $x \in$ O and *x* dances at *w*}

5. **frowns**: For all $a \in$ A, for all $w \in$ W: SV(**frowns**, *a*, *w*) = {*x*: $x \in$ O and *x* frowns at *w*}

---

40 Note to the reader: you should study and try to understand the new additions to L2 that make up L3. I will comment on the more important additions below, and explain their rationale.

…

2-place predicates

6. **loves**: For all $a \in$ A, for all $w \in$ W:
SV(**loves**, *a*, *w*) = {<*x*, *y*>: *x*, *y* $\in$ O and *x* loves *y* at *w*}

7. **admires**: For all $a \in$ A, for all $w \in$ W:
SV(**admires**, *a*, *w*) = {<*x*, *y*>: *x*, *y* $\in$ O and *x* admires *y* at *w*}

…

Common nouns:

8. **man**: For all $a \in$ A, for all $w \in$ W:
SV(**man**, *a*, *w*) = {*x*: *x* $\in$ O and *x* is a man at *w*}

9. **woman**: For all $a \in$ A, for all $w \in$ W:
SV(**woman**, *a*, *w*) = {*x*: *x* $\in$ O and *x* is a woman at *w*}

…

<u>Semantics for complex expressions of L2 (recursive part):</u>

Sentences:

10. If α is a 1-place predicate and β is a name, then for all $a \in$ A and for all $w \in$ W:
SV('α (β)', *a*, *w*) = 1 iff SV(β, *a*, *w*) $\in$ SV(α, *a*, *w*).

11. If α is a 2-place predicate and β and γ are names, then for all $a \in$ A and for all $w \in$ W:
SV('α (β, γ)', *a*, *w*) = 1 iff <SV(β, *a*, *w*), SV(γ, *a*, *w*)> $\in$ SV(α, *a*, *w*).

12. If α is a common noun and β is a 1-place predicate, then for all $a \in$ A and for all $w \in$ W:
SV('**every** (α) (β)', *a*, *w*) = 1 iff for every $x \in$ O: if $x \in$ SV**(**α, *a*, *w*), then $x \in$ SV**(**β, *a*, *w*).

13. If α is a common noun and β is a 1-place predicate, then for all $a \in$ A and for all $w \in$ W:
SV('**some** (α) (β)', *a*, *w*) = 1 iff for some $x \in$ O: $x \in$ SV**(**α, *a*, *w*), and $x \in$ SV**(**β, *a*, *w*).

14. I f α is a sentence, then for all $a \in$ A and for all $w \in$ W:
SV('**not**(α)', *a*, *w*) = 1 iff SV(α, *a*, *w*) = 0.

15. If α and β are in S, then for all $a \in$ A and for all $w \in$ W:
SV('α **and** β', *a*, *w*) = 1 iff SV(α, *a*, *w*) = 1 and SV(β, *a*, *w*) = 1.

15a. If α is a sentence, then for all $a \in$ A and for all $w \in$ W:
SV('**possibly**(α)', *a*, *w*) = 1 iff there is a $w^* \in$ W such that SV(α, *a*, w^*)=1.

15b. If α is a sentence, then for all $a \in$ A and for all $w \in$ W:
SV('**necessarily**(α)', *a*, *w*) = 1 iff for all $w^* \in$ W, SV(α, *a*, w^*) = 1.[41]

---

[41] In modal logic, possibility/necessity at a world *w* are usually defined as truth at some/every world *w** *accessible* from *w*. Thus, 14a. might then read:

… SV('**possibly**(α)', *a*, *w*) = 1 iff there is a $w^* \in$ W such that R(*w*, w^*) and SV(α, *a*, w^*) = 1.

15c. If α is a name and β is a sentence, then for all $a \in$ A and for all $w \in$ W: SV('α **believes that** β', a, w) = 1 iff for all $w^* \in$ W such that w^* is compatible with what SV(α, a, w) believes at w, SV(β, a, w^*) = 1.

Complex 1-place predicates:

16. $\lambda \boldsymbol{x}_i$: If α is a sentence, then for all $a \in$ A, for all $w \in$ W and for all i: SV('$\lambda \boldsymbol{x}_i(\alpha)$', a, w) = {x: for all $b \in$ A that differ from a at most in i-th place, if $x = b(i)$ then SV(α, b, w) = 1}

As before, we need to consider what our assignment of semantic values 1 and 0 to sentence-assignment-world triples means for our predictions of user assessments of literal truth. As before, we are at this point only interested in sentences that do not contain unbound variables, so in the interesting cases, our semantic values will not vary with the assignment parameter. But what about the world parameter? To be sure, sentences with dominant **possibly** or **necessarily** will have semantic values that are constant across worlds (just like expressions containing bound variables do not vary in semantic value from assignment to assignment). But many complete sentences (that do not contain unbound variables) will vary in semantic value from one world to another.

Now, we might say that there is one world, namely the actual world @, the way things really are, which is privileged. It is privileged in the sense that we ought to believe only what is really so, so that correspondingly we are especially interested in the semantic value 1 at @. Accordingly, let me redefine Truth and Falsity for L3 as follows:

($T^{@}$) For all sentences α of L3: α is True iff for all $a \in$ A, SV(α, a, @) = 1.

($F^{@}$) For all sentences α of L3: α is False iff for all $a \in$ A, SV(α, a, @) = 0.

Definition of Truth and Falsity in L3.

These definitions can serve for an evaluation of utterances as true or false. However, from the point of view of testing our predictions, we continue to be interested in the relativized semantic values. For language users will assess utterances according to what they believe to be the case, and this, as we saw, may diverge from what is actually the case. Moreover, the linguistic competence of users will consist in them being able to judge the *conditions*, whether they actually obtain or not, under which an utterance *would* be true. Thus, there is a good sense in which it is the relativized assignments of semantic values that represent the linguistic competence of users. Two language users can agree in their assessments as to the conditions under which, i.e. the possible worlds in which, an utterance would receive the value 1, without agreeing on whether the utterance is in fact True in the sense of ($T^{@}$).

Varying the accessibility relation R will then generate varying notions of possibility and necessity. I have omitted this feature for expository convenience.

This intensional language L3 will allow fairly straightforward solutions to two of the problems we noted. First, we can now treat the data we hope to predict with our model in a more realistic way. Competent users will not just make judgements as to which utterances are (literally) true, but they will do so as a result, in part, of their view of the actual state of the world, i.e. their beliefs. Suppose a speaker, Jack, is under the mistaken impression that Sue admires Peter. Then he might regard an utterance of the sentence "Sue admires Peter" as literally true, even though in fact it is false. Now intuitively, the datum that Jack regards the utterance as literally true together with the mentioned datum about his beliefs, should be taken as confirmation that he uses a language in which the sentence uttered is true just if Sue admires Peter. The extensional approach does not allow this. The fact that Jack regards the utterance as true disconfirms the theory that Jack is using a language of which L2 is a fragment, for L2 predicts that the utterance is literally False. We might disregard data of this sort and admit as evidence only assessments as to literal truth that are made under ideal conditions, where the ideal conditions somehow preclude that the competent speaker in question has a false belief. But it would seem a pity to throw out the easily accessible data and restrict ourselves to data that are difficult to obtain. By contrast, Jack's acceptance of the utterance would confirm the hypothesis that Jack speaks a language of which L3 is a fragment. For even though L3, just like L2, predicts that the sentence is False (given that Sue does not, actually admire Peter), it predicts that it would be true if the world were as Jack believes it to be.

Empirical confirmation of hypotheses regarding natural language intensions

This advantage extends further. What we want to predict, more generally, is under what conditions – hypothetical or actual – users would classify an utterance as literally true. This is, more strictly speaking what competence amounts to. For users' competence it is not directly relevant what they believe: they might be unreliable believers (at least in some areas) yet fully competent language users. So the data we need to predict is user assessments as to the circumstances in which a given utterance would be literally true. So, in gathering data concerning a given sentence, we need to match possible circumstances with the truth-value an utterance would have given those circumstances. We can gather these data by asking users for their assessments of literal truth in this or that circumstance. We can depict the circumstance in a number of different ways: describing the situation, presenting them with a picture, or a sound or smell, or a combination of these. This is, roughly, the methodology for confirming hypotheses regarding natural language intensions proposed by Carnap (1956).[42]

The intensional operators in L3

[42] It is not part of this picture that competent users must always know some definition or analytic truth in order to manifest their linguistic competence with a term. For possible situations can be referred in (partly) non-descriptive ways. Thus the proposal

Intensions also allow us a fairly straightforward treatment of the expressions **believes that**, **possibly** and **necessarily**. As specified in 15a., **possibly**, when applied to a sentence α, will produce a new sentence with a new intension (function from worlds to truth-values). **Possibly** is therefore an *intensional operator*. The semantic value of the new sentence '**possibly** α' at a world *w* will be 1 just if there is some possible state of the world *w**, such that semantic value of α at *w** is 1. 15b. provides a similar treatment for **necessarily**. If we had had to assign a function as a SV to **possibly** and **necessarily**, it would have been functions from one sentence intension to another (i.e. a SV of type <W→TV>→<W→TV>). The new expression **believes that** combines with a name and a sentence. Together with a name α, it forms an expression of the same category as **possibly**, i.e. an intensional operator. This is the complex intensional operator 'α **believes that'** which in turn combines with a sentence β, to form a sentence 'α **believes that** β' whose semantic value is 1 at a world *w* just if the semantic value of β is 1 at all worlds *w** that are not ruled out by what the SV of α at *w* (i.e. the "referent" of α at *w*) believes. This is specified in clause 15c. This corresponds to the idea, mentioned above, that to believe something is to rule out certain possibilities.[43] There are many other intensional expressions in natural languages that can be treated in a similar fashion. For example, "it is contingent whether", "says that", "desires that", "hopes that", "assumes that", "wonders whether" etc.

***Exercise 6**: articulate an appropriate semantic clause for **it is contingent whether** on the model of 15a. and 15b*

Treatment of names

It is worth noting that the only subsentential expressions that vary in SV with the world parameter are the *n*-place predicates and common nouns (see clauses 4.–9.). Thus, it will be different objects who frown and dance in different possible states of the world, objects may love or admire different objects in different possible state of the world, and it might be different objects who are men and women in different possible worlds.[44] The names in L3 are not subject to this sort of variation. Thus, as clauses 1.-3. of L3's semantics show, the semantic values of names do not vary with the possible world parameter, be they constant with respect to assigment,

does not immediately fall foul of Quinean worries about analyticity. Recently, Chalmers, in his Locke lectures, has pursued a similar strategy (http://www.philosophy.ox.ac.uk/lectures/john_locke_lectures). Chalmers, with his "cosmoscope" takes Carnap's idea of non-descriptive acquaintance with circumstances much further. See also Chalmers 2002.

43 See Stalnaker 1987 for such an account.

44 So the framework allows for manhood and womanhood to be a non-essential properties of objects, which is good news for those who regard sex-change as a possibility!

like **Sue** and **Peter**, or be they variable with respect to assignment, like x_1, x_2, … x_n.

Treatment of definite descriptions

This may make some reader wonder what role expressions like “the first logician” or “the capital of Norway” would occupy in a language like L3. Suppose Aristotle was the first logician, as he himself thought. Surely, there are possible worlds where Aristotle was preceded by another logician, say Plato, or possible worlds where Aristotle never got into logic, so that in those worlds someone else was the first logician. Similarly for “the capital of Norway”, for Tromsö could have been the Norwegian capital. This suggests that the semantic value of “the first logician” and “the capital of Norway” should vary from one world to another. For we want the semantic value of “the first logician dances” at a possible world w at which Plato invented logic, to depend on whether Plato dances at w. Similarly, if at w, Tromsö is Norwegian capital, then we want the semantic value of “The capital of Norway is south of the polar circle.” at world w to depend on the location of Tromsö, not Oslo. This means that in L3, we should cannot treat “the first logician” and “the capital of Norway” as names. What kind of expressions could they be? Both expressions are complex and contain common nouns: “capital of Norway” and “first logician”. These are the kind of expression that vary their semantic value with the possible world parameter. What we need to do is to add the determiner the to the determiners we already have. Then the two terms “the first logician” and “the capital of Norway” would count as quantifier phrases, as specified in clause 1 of the syntax of L3.

Logically necessary or contradictory sentences

One of the problems I mentioned for extensional semantic theories has not yet been addressed: the semantic values of any two logically necessary or logically false sentences will be the same at all possible worlds, namely a constant function from possible worlds to 1 or 0 respectively. This problem could be solved by introducing semantic values that are more fine-grained than intensions, such as structured intensions or structured propositions.45 Alternatively, one could bite the bullet and maintain that contrary to appearances, all logically necessary (contradictory) sentences make the same contribution to the contents of sentences in which they can be embedded, explaining away pragmatically the appearance that our assessments of literal truth do seem to be sensitive to the differences between different logically necessary (or contradictory) sentences.46 As this issue will not play any role for the main theme of this course, I shall leave this matter open.

45 See, for example, Lewis 1970, or King 2007.

46 This is the strategy pursued in Stalnaker 1984.

3.7 Summary

We have now arrived at a basic description of the type of language that will serve as a model for natural languages. Let me summarize what I have done in this chapter. I have begun by considering the kind of predictions we would like a semantic theory for a language to make, namely predictions as to the judgements competent users make about the correctness of utterances. However, I pointed out that there are many ways in which the correctness of utterances can be assessed, and that these ways may diverge. I conjectured that in order to arrive at a systematic theory of meaning, it is necessary to focus on some central dimension of assessment, a dimension with respect to which other dimensions would be derivative. I therefore focussed on assessments of the truth of utterances. However, there were again several ways in which the truth of an utterance can be assessed. I suggested that it would be promising to focus on assessment of literal truth, rather than the truth of non-literally conveyed messages. While it is a controversial matter how exactly we should draw the line between the literal and the non-literal, I assumed for the time being that there would be some way in which we can distinguish competent users' judgements as to literal truth from their judgements as to non-literal truth. The starting point was then to look for a description of a language that allows us to make predictions as to the literal truth of utterances of its sentences.

With the data to be predicted identified in this way, I proceeded to examine how formal model languages familiar from formal logic might be used to make such predictions. I pointed out that in principle, the languages of formal logic are good candidates to figure as models of natural languages. For they are compositional languages and compositionality seems to be, at least prima facie, the right way to explain how language users are able to be competent in the use of novel sentences that they haven't encountered before. If the semantic properties of any semantically complex expression is a function of the semantic properties of its constituent parts, then such an explanation is available. However, I showed how a language of first order logic failed to justice to our requirements because it was extensional. One problem with the extensionality of languages of this type was that they did not lend themselves to a compositional treatment of certain "intensional" contexts, such as "believes that" or "possibly".

Intensional languages offered a fix: they assign intensions to sentences, i.e. functions from possible states of the world to truth-values. This allowed a straightforward treatment of intensional expressions as well. It also permitted an easier approach to the data we are trying to predict, namely judgements as to the literal truth of utterances. Language users' judgements as to whether an utterance is literally true will depend in part on

their beliefs about the way things are. Thus if we take the data to be constituted by competent users' judgements as to which utterances would be correct given certain assumptions about what the world is like, we don't need to restrict ourselves to the judgements of ideal users.

A further complication was the fact that the compositional treatment offered by compositional formal languages may require certain syntactic structures that diverge from natural languages. I proposed that the semantic theory could operate at the level not of surface sentences, but rather at the level of a base language (logical form) from which the surface forms of natural language sentences are generated via a transformational grammar. The resulting framework is summarized in this diagramme:

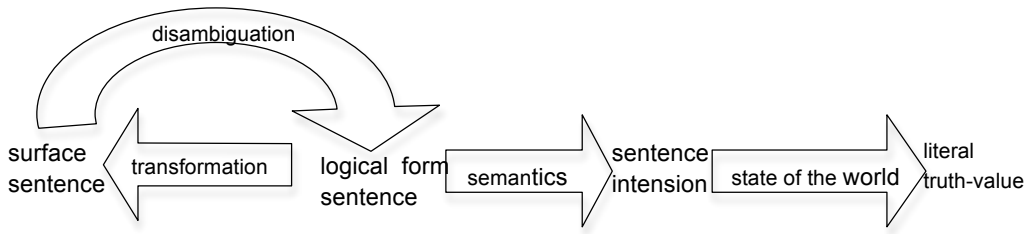

Our theory, aspiring to explain/predict competent users' judgements as to literal truth, starts with a transformational grammar which will tell us of which sentences in the base language the surface sentences of a natural language are transformation. Base sentences then receive a semantic interpretation that assigns to each sentence an intension, i.e. a function from possible states of the world to truth-values. This is the point at which the framework interacts directly with the data concerning competent language users' judgements of literal truth. They will judge an utterance of a sentence to be literally true just if its intension delivers the value 1 for all the possible states of the world that are compatible with what they believe.

Despite the considerable complications we have already considered, the resulting framework still suffers from many limitations. It provides no way of modelling non-declarative sentences. Nor is there room for context sensitive sentences, such as "I am hungry.", which will not express the same intension whenever they are used. When I say "I am hungry.", the utterance might be true, while if you say "I am hungry" at the same time, the utterance might be false. If we assume that the same possible state of the world is relevant for the evaluation of both utterances, then our current framework does not allow us to predict that one utterance is true while the other is false. But it should. For competent users will sometimes judge utterances of the same sentence to have different truth values. The next chapter will outline one way of expanding this kind of framework in order to allow for various kinds of context sensitivity.

4. Kaplan's Framework

4.1 Context dependence and temporal intensions

In the framework as so far developed for modelling natural languages, the semantics assigns to each (disambiguated, closed) sentence a sentence intension, which will be, or at least determine, a function from a possible state of the world to truth-values. This accorded well with our initial idea that what the semantics needs to predict is the conditions under which an utterance of a sentence is literally true, i.e. what the state of the world must be like in order for the utterance to be true. However, our framework requires that the same sentence will always have the same sentence intension, and consequently that the conditions for the literal truth of an utterance will depend exclusively on the sentence uttered.

But natural languages seem to contain sentences that do not easily fit this restriction. Thus, it seems that the conditions for the literal truth of an utterance of the sentence

(1) I am hungry now.

Context dependence

can vary depending on who utters it at what time. If I utter the sentence on 13th of March 2010 at 12 noon, then the truth-value of the utterance will depend on whether I am hungry at that time. The possible states of the world in which *I* am hungry at noon on 13th March 2010 are the ones with respect to which the utterance should count as true. However, if *you* utter sentence (1) at the same time, then the literal truth of the utterance depends on whether *you* are hungry at that time. Similarly, if I utter sentence (1) at a different time, say 2pm, on the same day, then the literal truth of the utterance will again depend on a different aspect of the world, namely my state of hunger at 2pm, not at noon. As we saw in chapter 1, there are many expressions that produce this dependence on the context in which a sentence is used.

Indexicals in Kaplan's framework

In this chapter, I want to outline the best known framework for modelling context dependence, namely David Kaplan's (Kaplan 1977/1989).[47] This

[47] In fact, it seems that a number of theorists were working on context dependence in the early seventies. What I am mainly interested in introducing in this chapter is the idea of "double-indexing", which seems to come out clearest in Kaplan's paper (which was circulated widely from 1971 on and was presented at a conference in 1977). A pioneer who preceeded Kaplan is Kamp (1971), at least as the time-index is concerned. Other pioneers include Stalnaker 1970 and Lewis 1980. Lewis 1970 does not yet take into account double-indexing.

framework primarily addresses itself to the context sensitivity produced by *indexical* expressions, i.e. expressions like "I", "you", "she", "now", "yesterday", "here" and so on, as well as demonstratives such as "this", "that" or "he", "she" (in their demonstrative uses). The sentence "I am hungry now." Exhibits this indexicality as a result of containing the indexicals "I" and "now". As a result, the literal truth of utterances of the sentence will depend on how things stand with respect to the person uttering the sentence and with respect to the time at which they are uttering it. According to Kaplan, this means that the content (intension) of the same sentence can vary with the context in which it is uttered. Thus, the meaning of the sentence should be construed not as an intension (function from possible world to truth-value) but as a function from a *context of use* to an intension. Or, put differently, it is sentences *in context*, or sentence-context pairs, that are assigned an intension by the semantics, not just the sentences. This is due to certain expressions, like the ones mentioned above, the intensions of which will vary with the context of use.[48]

There are also other forms of context dependence in Kaplan's system. Consider the sentence

(2) I am hungry.

Temporal intensions

Like sentence (1), sentence (2) seems to exhibit context dependence with respect to time. For the truth of an utterance of (2) seems to depend just as much on the time at which the utterance takes place as an utterance of (1) would. But this can no longer be attributed to the context-dependence of "now", because that word has been dropped (the context dependence produced by "I" remains the same). According to Kaplan, however, unlike utterances of sentence (1), an utterance of sentence (2) expresses a content or intension that is unspecific as to time. This means that the truth-value of a content (intension) expressed by (2) will vary not just with a possible state of the world, as is familiar from the last chapter, but it will also vary with a time. In Kaplan's framework, the truth-value of a propositional content will vary from world to world and from time to time. Kaplan's intensions are functions from world-time pairs to truth-values.

Varying intensions vs constant intensions that vary in truth-value

48 Apart from the question of modelling the context-dependence of intensions, another central theme of Kaplan 1977/1989 is Kaplan's claim that the intensions of indexical expressions in contexts are *constant*, i.e. that they assign the same value to every possible world, or in other words, are "rigid". In the discursive part of 1977/1989, Kaplan also defends what he calls his "metaphysical" doctrine, namely that the contribution an indexical in context makes to the content of sentences that contain them is not a constant function, but rather the individual itself that would otherwise have been the constant value of the function. This is Kaplan's doctrine of "Direct Reference". I will leave these issues aside, as they do not concern the central theme of this course, namely the construal of context-dependence.

This is an important point, so let me expand a little. When someone *a* utters sentence (1) at a time *t*, *a* thereby asserts a propositional content that is specifically about *a* and *t*. This content will be true at a time and world just if *a* is hungry at *t* at that world. Thus, this content will be true at the world-time pair <*t*1, *w*1> just if *a* is hungry at time *t* in world *w1*, and it will be true at the world-time pair <*t*2, w2> just if *a* is hungry at time *t* in world *w*2, etc. The content is not world-invariant (i.e. it is contingent), but it is time-invariant. By contrast, when someone *a* utters sentence (2), at time *t*, then the content expressed is different and it is no longer time-invariant. This second content will be true at a world-time pair <*t*1, w1> just if *a* is hungry at time *t*1 in world *w*1, and it is true at <*t*2, *w*2> just if *a* is hungry at time *t*2 in world *w*2, etc.

Why should we treat (1) and (2) so differently? There is an observable difference between (1) and (2): while (2) can easily be embedded in temporal operators,[49] (1) can't. (2) can be qualified by temporal expressions, as in "Sometimes I am hungry." or "Every morning I am hungry.". Not so with (1): "Sometimes I am hungry now." or "Every morning I am hungry now." sound odd. Perhaps these can be used, with some kind of poetic license, to express "Sometimes I am hungry at that very moment." and "Every morning I think to myself 'I am hungry now'.".[50] But many competent users sense a kind of semantic oddness here. This can be explained, if we think of the job of "sometimes" and "every morning" as making specific some temporal unspecificity in the sentence to which they attach. (1) is already temporally specific, so there is no job to do for the temporal qualifier. (2) lacks this temporal specificity, so that's why it can be embedded without oddity in the ways suggested. In fact, "now" seems to be an expression that does the same kind of job that "sometimes" and "every morning" do, namely the job of removing temporal unspecificity.

Temporalism

Kaplan (along with other temporalists) says that utterances of a sentence like (2) express a propositional content that is temporally unspecific, i.e. a function from world-time pairs to truth-values that is not constant with respect to time, while contents expressed by utterances of (1) differ only in that they are constant over the time argument, taking always the value that

49 By calling these expressions "temporal operators" I do not mean to presuppose a temporalist treatment like Kaplan's and Prior's. For now, read "temporal operator" wide enough to include quantifiers that bind temporal variables. Perhaps "temporal qualifier" would be sufficiently neutral. The matter will be discussed extensively below.

50 Consider this line from a song: "Heute ist der Tag, an dem ich dich ein Leben lang ertrag." ("Today is the day when I put up with you all my life."), which contains similar play with semantic oddness: how can one restrict putting up with someone all one's life to just one day? If that restriction is necessary, then it undermines the claim of putting up all one's life.

the first, unspecific content would take for the time of utterance. Similarly, "sometimes" makes the content expressed by (2) temporally constant: "Sometimes I am hungry" constantly has the value true if the content expressed by (2) has the value true for some time. This explains the oddity I observed as a form of redundancy: if the job of "sometimes" is to remove temporal unspecificity, then attaching it to a sentence that already expresses temporally specific contents is redundant. It is similar to the redundancy of using a binder on a formula that contains no variable.

Parallel treatment of tense and modality

As we will see below, this treatment of "sometimes" and "now" is not the only alternative. But for now let's observe the parallelism of Kaplan's treatment of tense and modality. A sentence like (2) expresses a propositional content, or intension, that varies in truth-value with time and world: someone hungry at one time and world need not be hungry at the same time in another world, and similarly need not be hungry at a different time in the same world. Hunger is a temporary and contingent property of people. So just as the extension of "is hungry" will vary with the possible state of the world, so will it vary with time. This means that "sometimes" and "always" are treated as intensional operators that are analogous with the modal operators "possibly" and "necessarily". The semantic value of a sentence 'possibly α' at a context *c*, time *t* and world *w* will be 1 just if there is a world *w** such that the semantic value of α is 1 at *c*, *t* and *w**. Exactly analogously, the semantic value of a sentence 'sometimes α' at a context *c*, time *t* and world *w* will be 1 just if there is a time *t** such that the semantic value of α is 1 at *c*, *t** and *w*. "sometimes" and "always" are intensional operators that make sentence intensions invariant with respect to time, just as "possibly" and "necessarily" are intensional operators that make sentence intensions invariant with respect to possible world.

"Now" and "actually"

What, then, is the role of "now", as in (1)? "Now" can also be treated as an intensional operator, namely one that sets the time of evaluation to the time of the context. So a sentence 'now α' will be evaluated as true at a context *c*, time *t* and world *w*, just if the sentence α is evaluated as true at *c*, the time of *c* and *w*. Thus it is a context-sensitive intensional operator which sets the time of evaluation to the time of the context. The operator "actually" is often treated in an exactly analogous way, namely as a context-sensitive intensional operator that sets the world of evaluation to the world of the context. So a sentence 'actually α' will be evaluated as true at a context *c*, time *t* and world *w*, just if the sentence α is evaluated as true at *c*, *t* and the world of *c*. Thus each context determines a world of the context, just as it determines a speaker, a time, a place etc of the context.

Circumstances of evaluation

As we saw, in Kaplan's system, intensions are evaluated at world-time pairs. We might call these world-time pairs "points of evaluation" (as does Predelli 2004), or with Kaplan "circumstances of evaluation". It will be useful to have a general term for the entity with respect to which intensions are evaluated because we will soon consider suggestions for additional parameters with respect to which intensions should be evaluated. So let us say that in Kaplan's system, circumstances of evaluation, i.e. the points at which intensions are evaluated, are ordered pairs <*t*, *w*> consisting of a time *t* and a possible world *w*.

This gives us the following picture of how the semantics evaluates a disambiguated closed sentence as true or false:

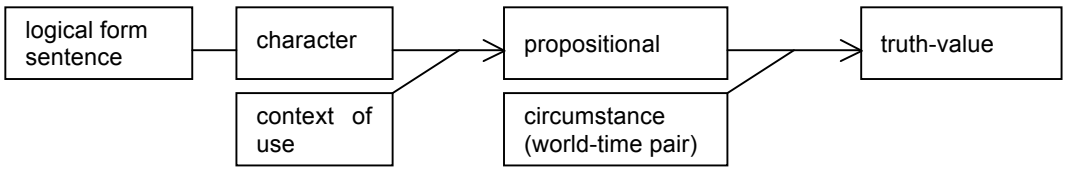

At the end of the last chapter, I asked how the predictions of the semantics as to the intension of a sentence bear on our data concerning the literal truth of sentences. For these intensions are in many cases non-constant functions from worlds to truth-values, and it is not immediately obvious the truth-value with respect to which possible world is relevant for the prediction the semantics will make as to the literal truth of utterances. If sentence contents have their truth-values only relative to a world at which they are evaluated, then how do we predict which utterances competent users will regard as literally true?

The world relevant for user assessments

I pointed out that there is one possible world that we are especially interested in, namely the actual world: the way things really are. Thus we can define an absolute notion of truth in terms of the relative notion the semantics operates with: truth at the actual world (see ($T^{@}$) and ($F^{@}$) above). However, I also pointed out that our predictions should take into account what the competent users whose assessments we are trying to predict believe: naturally, they will assess an utterance as literally true when its intension is true at all the worlds that are compatible with what they believe. But their beliefs might be mistaken (false at @), in which case the actual world is not among the worlds compatible with what they believe. So their assessments of literal truth do not need to coincide with actual truth as predicted by the semantics. Rather they need to coincide with what they judge to be actually true, i.e. with truth at all the worlds compati-

ble with what they believe. I also envisaged further predictions as to users' assessments of literal truth given certain circumstances, hypothetical or actual. These assessments are the assessments users make when they suppose the hypothetical circumstances to be actual.

The world-time pair relevant for user assessments

Our new circumstances of evaluation contain an additional time parameter, and there will again be sentences that have intensions at a context of use that are non-constant with respect to this time parameter of evaluation. Thus we can raise an analogous question: truth with respect to which circumstance of evaluation is relevant for the predictions of competent users' assessments of literal truth? Again, the actual world seems to be of special interest, as everyone is aspiring to beliefs that are true at the actual world. But which time is especially interesting when we are evaluating the content expressed by a sentence-context pair? The obvious candidate is the time of the context. For we will judge an utterance of (2) to be literally true just if the speaker is hungry at the time of utterance. Thus the circumstance consisting of the time of the context and the actual world seems especially interesting. Similarly, we can ask users to evaluate purely hypothetical utterances, which occur at worlds that are not actual or times that are not the time of evaluation.

As we saw, just as a context of use determines a time, it will also determine a world. For each utterance is made at a possible world, just as it is made at a certain time. So there is another way in which we can specify the world that is especially relevant for the evaluation of a sentence-context pair: the world *of* the context. Thus another circumstance that is relevant for the evaluation of a sentence-context pair is the circumstance consisting of the time and world of the context. This will be interesting, for example, if we are assessing the truth of a non-actual utterance. Such an utterance will be the utterance of a sentence in a non-actual context, i.e. a context the world of which is not the actual world. If we ask ourselves whether such an utterance is true, we can be interested in different kinds of assessment. First we might be interested in whether the utterance would be true as evaluated against the time and world at which it is made, i.e. from the perspective of the non-actual world at which the utterance is made. This will be truth as non-actually aspired to by people in that world, including the speaker and audience of the non-actual utterance. Secondly, we might be interested in whether the utterance is true as evaluated against the time of the utterance and the actual world. This might be what we are interested in if we ask whether someone would have said something true if they had made a certain utterance.

World of evaluation vs world of utterance

For example, consider a non-actual utterance of "I am speaking now", which is made at a certain context *c*. Suppose further that in the actual world, the speaker of *c* is silent at the time of *c*. Then the answer to the

question "Is the utterance true?" will be different depending on which of the two senses we have in mind. In the first sense, it is true, while in the second sense it is false. The speaker, on the one hand, would have been speaking at the relevant time in the non-actual world of context *c*. However, he is not actually speaking at that time. Thus, we can assess the intension expressed at the non-actual context with respect to the world of the non-actual context, or with respect to the actual world (or the world we take ourselves to be in). It does not make sense to ask which is *the* right way of assessing an utterance, for both might play a role. What is important is not to confuse them.

In using our semantics to predict which utterances competent users classify as literally true, we will therefore have to be sophisticated. When considering their assessments of actual utterances, we predict that they judge them to be literally true just if the semantics assigns to the sentence used in the context in which it was used an intension that is true at the worlds that are compatible with what the competent users believe. But we could also ask them to assess the content expressed by that actual utterance with respect to non-actual circumstances. If they assess whether a merely possible utterance would have been true, we have to distinguish their assessment of truth at the non-actual world of utterance from their assessment of truth at the actual world as they see it. In general, we have to keep in mind, that user assessments are now sensitive to two factors: the context in which the sentence is used and the circumstance with respect to which it is evaluated.

In the next two sections, I shall look in more detail at how L3 can be modified to implement the new context dependence of contents and the new time-dependence of truth values.

4.2 Context Sensitivity

Does L3 have room for context dependent intensions?

Let us consider again, why the current semantics for L3 is not adequate for indexical sentences like (1). The theory will assign to each closed sentence of L3, i.e. to each sentence that does not contain unbound variables, a sentence intension (propositional content), which is a function from possible states of the world to truth values. We saw that in the very same state of the world, two utterances of the sentence (1), "I am hungry now.", could differ in literal truth. For example, my utterance of it at noon could be true and my utterance at 2pm false. Similarly, given the very same state of the world, my utterance of the sentence could be literally true and your utterance of the same sentence false. Given that our semantics is supposed to predict the literal truth of utterances, we have to find a way in which the semantics assigns two different propositional contents to the two

utterances of the same sentence, i.e. two different functions from world states to truth values.

Only if we deny, with Frege, that the same complete sentence was uttered

Given a framework like that of L3, there is only one way of doing this, namely to deny that the two utterances of the phonetic type (1) are utterances of the same sentence. In other words, the phonetic type (1) does not yet give us the complete sentence used, which then gets assigned an intension. Rather, the phonetic type (1) together with some aspect of the respective events of utterance gives us two different complete sentences with different propositional contents. This is the solution that Frege adopted (Frege 1918). Frege thought that in these cases, we should regard the mere form of words "I am hungry now" as not by itself determining the complete sentence, i.e. the complete thought-expression. The sentence is complete only in conjunction with some concrete aspects of the particular token of the sentence that is produced in each of the utterances. Thus, the semantics assigns intensions not to phonetic sentence types, but to sentence types together with certain properties of the particular utterance of the sentence. Thus, the expression type "I am hungry now." by itself needs completion with certain features of actual utterances of the type, and it is to these "hybrids" made up of expression types and features of concrete utterances that the semantics assigns propositional contents.[51]

Rationale for not going Frege's way

However, there also seems to be a very good sense in which on each of the utterances, the sentence "I am hungry now." is used with exactly the same meaning, and exactly the same semantic properties are exploited. It also seems that it is knowledge of the semantic properties that are shared between the different utterances of the same type which constitutes speakers' competence with the sentence, that knowledge of the same semantic properties is brought to bear each time. Of course, mere knowledge that the expression *type* "I am hungry now" has been used in an utterance may not give one *full* understanding, in some sense of "full". One might think that for *full* understanding I also need to know who uttered the sentence and at what time.[52] Nevertheless, there is a less demanding sense of "understanding" in which I can understand utterances of the sentence even without knowing much about who made the utterance and

51 See Künne 1992 and Kripke 2008 for detailed discussion.

52 It may be quite difficult to specify what exactly full understanding requires. For there is a sense in which a competent speaker always knows about *whom* and about *what time* a particular use of (1) is, no matter how little she knows about the event of using the sentence. Arguably, she will always know that the utterance was about the user of the sentence and about the time of use. I am, cautiously, speaking of "uses" and "users" of the sentence, rather than utterances, because of special cases, such as recorded messages, where the user and the time of use might be distinct from the utterer and the time of utterance, at least in the most straightforward sense of "utterer" and "utterance". See Predelli 1998 and Perry 1997.

when. This less demanding understanding seems to play an important role, for it is the kind of knowledge of language that speakers possess continuously, and which they bring to bear on new communicative situations. Kaplan's semantic framework therefore introduces a further stage in the determination of literal truth values: sentences get assigned contents only relative to contexts, and these in turn have truth-values only relative to circumstances or points of evaluation.

Character

In Kaplan's framework, the semantics assigns to each sentence a *character*, where a character is a function from a *context of use* to a propositional content. Equivalently, we might say, that in Kaplan's framework, propositional contents are not assigned to sentences, but to sentence-context pairs. This is not in the end so very different from Frege's solution: Frege's sentences (his "complete expressions of thought") are much like Kaplan's sentence-context pairs.

Constant and non-constant characters

If a sentence is indexical, its character is a non-constant function from context to content. For illustration, compare the sentence "I am hungry now." with the sentence "MK is hungry at noon on 13th March 2010.". The character of the first sentence will be such that for a context where the speaker is MK and the time is noon on 13th March 2010, it determines the propositional content that MK is hungry at noon on 13th March 2010. The character of the second sentence will have exactly the same value for *that* particular context. However, the character of the second sentence will have that value in all contexts of use, while the first one will not.

Let me begin with the treatment of "I". For simplicity, let us consider a sentence that only contains one context-dependent expression:

(3) I am hungry at noon on March 13th 2010.

"I"

And let us treat "is hungry at noon on March 13th 2010" as if it was a simple predicate, that has an extension that varies with a possible world. At each world, its extension contains exactly those individuals that are hungry at noon on March 13th 2010 at that world. How should we implement the context-dependence of "I", "you" and similar indexicals in a semantic theory? The assignment of a truth-value, at a world *w*, to an utterance of sentence (3) should depend on whether the speaker of the utterance is in the extension at *w* of the predicate. For an utterance of sentence (3) will be judged to be true at a world just in case the speaker of that utterance is judged to be hungry at the time in question at that world.

Thus it looks like we should say that the SV of "I" at a context should be the speaker of that context. Similarly, the SV of "you" at a context should be the audience at that context, and the reference of "here" at a context

should be the place of the context. So a context needs to determine at least a speaker, an audience and a place. I already said that Kaplan's contents (intensions) are no longer functions from worlds to truth-values (or other extensions), but rather functions from *circumstances of evaluation* to truth-values, where a circumstance is an ordered pair of a time and a possible world. So our semantics will now define the semantic value function SV(*e*, *c*, *a*, *p*), which delivers extensions for a quadruple of an expression *s*, a context *c*, an assignment *a* and a circumstance (or "point") of evaluation $p \in$ P (P is the set of ordered pairs <*t*, *w*>, with *t* a time and *w* a world). In a moment, I shall say more about what a context is. For the moment all we need is that a context determines a unique speaker, so that we can speak of A(*c*), i.e. the speaker of context *c* ("A" for "agent" of the context). It seems plausible to treat "I" syntactically like a name, as it can combine with 1-place predicates to form sentences (though more on this in a moment). So, if our semantic clause for an ordinary non-context-sensitive name **Sue** is now something like:

1. **Sue**: For all contexts *c*, for all $a \in$ A, for all $p \in$ P:
   SV(**Sue**, *c*, *a*, *p*)= o_1

then our semantic clause for the indexical **I** should be:

2a. **I**: For all contexts *c*, for all $a \in$ A, for all $p \in$ P:
   SV(**I**, *c*, *a*, *p*)= A(*c*) (i.e. the speaker of *c*)

Sample clause for "I"

The difference between **Sue** and **I** is that **Sue** has the same SV for all contexts of use, while the SV of **I** varies with the context.

It is important to notice that the SV of **I** does not vary with the world parameter. Thus, even if we evaluate **I** at a non-actual world where the actual speaker does not speak, it is still the speaker of the context who is relevant. Thus, at a given context of use, sentence (3) will express a certain propositional content which is about the speaker of that context, and it is the proposition that that person is hungry at noon March 13th 2010. Suppose that that proposition is true at the actual world. If we consider a different world in which I eat an enormous fried breakfast at 11am, then the proposition is false at that world. It does not matter that perhaps at that other world I do not utter the sentence, so am not a speaker. This follows from 2a. and is as it should be. For consider the sentence

(4) The person speaking here now is hungry at noon on 13th March 2010.

Indexicals make a rigid contribution to content

Now compare utterances of (3) and (4) in the same context of use *c*1. Consider a possible world *w*1 where the speaker of *c*1 has an enormous breakfast at 11am 13th March 2010, and where she is not speaking, but where someone else is speaking at the time and place of *c*1, someone

who is hungry at noon 13th March 2010. The content of the utterance of (3) in *c*1 would be false at *w*1, because the speaker of *c*1 is not hungry at *w*1. But the content of the utterance of (4) in *c*1 would be true at *w*1. For at *w*1, the person speaking at the time and place of *c*1 is indeed hungry. The speaker of *c*1 is irrelevant for an evaluation of the content expressed by the utterance of (4) in *c*1.

"I" as name

I said a moment ago that "I" is plausibly put into the category of names. The example I just gave shows why we can't easily put it into the category of quantifier expressions, even though these combine equally well with 1-place predicates. Suppose we said that "I" functioned just like "the speaker". Then its semantic value would vary with the possible world parameter, for worlds can differ with respect to who is speaking

Jobs the context needs to do

How should the new context of use parameter be construed? Consideration of "I". "you" and "here" shows that a context must at least determine a speaker, an audience and a place. But it must do more. If we want to allow indexicals like "now", the context must also determine a time. If we want to have indexicals like the demonstrative "that", as in "that man", or "this" in "this car", the context also needs to determine a demonstrated object. For it seems clear that the truth of utterances containing these demonstratives depend on which object is being demonstrated by the speaker. Anaphoric pronouns seem to be another example of expressions whose SV depends on context. When I say: "So does Peter.", the SV will depend on what has been said before. If what has been said before is "I admire Sue.", then the SV will be the same as that of "Peter admires Sue." If what has been said before was "You dance.", then the SV of "So does Peter." will the same as that of "Peter dances.". As mentioned, "actually" is also often construed as indexical, and so is "actual", as in "the actual winner". These will pick out the world of the context. Thus the context will need to determine not only a speaker, an audience, a time, a place and a demonstrated object but also a previous discourse and a world.

Definition of "context of use"

There are many other context sensitive expressions which depend on yet further aspects of a context of use. The context will need to determine relevant values for each of them. How can we do justice to that in the semantics? Luckily, we can reduce the number of parameters to four: an agent (the speaker), a place, a time and a possible world. This is enough if we stipulate that the agent of a context is at the place of that context at the time of that context at the world of that context. We might say, with Lewis (1980), that contexts are "locations" in space-time and logical space. These four factors will then determine anything else that we might need. Once we have an agent, a time, a place and a possible world, then there can only be one object that is being demonstrated by the agent at that time, place and world. And there is only one preceding discourse and

audience. With a different preceding discourse or audience, the world or the point in space and time would have had to be different, for possible worlds are maximally specific ways the world might have been. Thus we can construe a context simply as an ordered set <*a*, *l*, *t*, *w*>, where *a* is an agent, *l* a location in space, *t* a time and *w* a possible world, such that *a* is at *l* at *t* in *w*. We will call the set of these quadruples "C".

This stipulation has certain desirable effects. Consider the sentences

(5) I am speaking.

(6) I am here.

(7) I exist.

A moment's reflection shows that it is quite difficult to utter any of these sentences without expressing a propositional content that is true at the world and time of the context of utterance. Our stipulation that "I" has as its SV at a context the speaker of that context explains why the content of (5) expressed at a context will always be true, at least at the world of that context. Similarly with (6) and (7): since "I" has as its SV at a context the speaker of that context, and the speaker of a context is always at the location of the context, we only need to say that "here" has as its SV at a context the place of the context in order to ensure that the content expressed by (6) will be true at least at the world and time of the context. Similarly, in order for a speaker to be at a place and time in a world, the speaker needs to exist at that world at that time. This ensures that (7) will always express a content that is true at the world and time at which it is uttered.

Effects of the definition

This does not mean that the contents expressed by these sentences do not receive different truth-values at different worlds. Whoever is using (5) at whatever place and whatever time in whichever possible state of the world: that person might not have been speaking at that place and time, so there will be possibilities where that person is not speaking at that place and time, so that the content expressed by that person using (5) at the first-mentioned place and time and world will be false at those other possibilities.

4.3 The Semantics for L4

We can now attempt a modification of our model language L3 that makes it conform to Kaplan's framework. Have a look at the syntax and seman-

tics for L4, paying special attention to the clauses 2a-c., 13a-b. and 15c-i. of the semantics, as these explain the newly added vocabulary:

Syntax of L4:

Vocabulary of atomic expressions of L4:

As for L3 with some additions:

Names: **I**, **you**, **she** [demonstrative]

Determiners: **the**, **that**

Sentence operators:

indexical: **now**, **it was the case that** (past tense), **it will be the case that** (future tense), **yesterday** …

iterable: **the day before**, **at a later time**, …

Definition of complex expressions of L4:

As for L3, with one addition:

Sentences:

3c. If α is a sentence operator and β is a sentence, then '$\alpha(\beta)$' is a sentence.

Semantics of L4:

The sets O, A, TV and W are as in the semantics for L3.

The set of times T={ t_1, t_2, … t_n}

The set of locations L={ l_1, l_2, … l_n}

The set of contexts C={<*a*, *l*, *t*, *w*>: *a* ∈ O, *l* ∈ L, *t* ∈ T, *w* ∈ W and *a* is located at *l* at *t* in *w*}

Abbreviation: For all *c* ∈ C if *c* = <*a*, *l*, *t*, *w*>, then:
A(*c*) := *a*, L(*c*) := *l*, T(*c*) := *t*, W(*c*) := *w*.

The set of circumstances (points) of evaluation P ={<*t*, *w*>: *t* ∈ T, *w* ∈ W}

Abbreviation: For all *p* ∈ P, if *p* = <*t*, *w*>, then:
T(*p*) := *t*, W(*p*) := *w*.

We define a function SV(*e*, *c*, *a*, *p*):
Domain of SV: *e* is either a name, an *n*-place predicate, common noun or sentence of L4. *c* ∈ C. *a* ∈ A. *p* ∈ P.
Range of SV: O ∪ $\wp$(O) ∪ O×O $\wp$(O×O) ∪ TV

Most clauses are just modified from L3 by adjusting them to the new domain of the semantic value function SV. But 2a–c., 13a–b. and 15d–i. are additions.

Semantic values of atomic expressions:

For all *c* ∈ C, for all *a* ∈ A, for all *p* ∈ P: …

Constant names:

1. **Sue**: … SV(**Sue**, *c*, *a*, *p*)= o_1 (∈ O)

2. **Peter**: … SV(**Peter**, *c*, *a*, *p*)= o_2 (∈ O)

2a. **I** … SV(**I**, *c*, *a*, *p*)= A(*c*)

2b. **you** … SV(**you**, *c*, *a*, *p*) = A(*c*)'s audience

2c. **she** … SV(**she**, *c*, *a*, *p*) = the female demonstrated in *c* [i.e. de monstrated by A(*c*) at L(*c*) at T(*c*) at W(*c*)], if any, otherwise unde fined.

Variable names:

3. $\boldsymbol{x_1}$, $\boldsymbol{x_2}$, … $\boldsymbol{x_n}$: … for all *i* such that $1 \leq i \leq n$, SV($\boldsymbol{x_i}$, *c, a, p*)= *a*(*i*)
[i.e. the object in *i*th position in assignment *a*]

1-place predicates:

4. **dances**: … SV(**dances**, *c, a, p*) = {*x*: *x* ∈ O and *x* dances at T(*p*) at W(*p*)}

5. **frowns**: … SV(**frowns**, *c, a, p*) = {*x*: *x* ∈ O and *x* frowns at T(*p*) at W(*p*)}

…

2-place predicates

6. **loves**: … SV(**loves**, *c, a, p*) = {<*x*, *y*>: *x*, *y* ∈ O and *x* loves *y* at T(*p*) at W(*p*)}

7. **admires**: … SV(**admires**, *c, a, p*) = {<*x*, *y*>: *x*, *y* ∈ O and *x* admires *y* at T(*p*) at W(*p*)}

…

Common nouns:

8. **man**: … SV(**man**, *c, a, p*) = {*x*: *x* ∈ O and *x* is a man at T(*p*) at W(*p*)}

9. **woman**: … SV(**woman**, *c, a, p*) = {*x*: *x* ∈ O and *x* is a woman at T(*p*) at W(*p*)}

…

<u>Semantics for complex expressions of L4 (recursive part):</u>

For all *c* ∈ C, for all *a* ∈ A, for all *p* ∈ P: …

Sentences:

10. … if α is a 1-place predicate and β is a name, then SV('α (β)', *c, a, p*) = 1 iff SV(β, *c, a, p*) ∈ SV(α, *c, a, p*).

11. … if α is a 2-place predicate and β and γ are names, then SV('α (β, γ)', *c, a, p*) = 1 iff <SV(β, *c, a, p*), SV(γ, *c, a, p*)> ∈ SV (α, *c, a, p*).

12. **every** … if α is a common noun and β is a 1-place predicate, then SV('**every** (α) (β)', *c, a, p*) = 1 iff for every *x* ∈ O: if *x* ∈ SV(α, *c, a, p*), then *x* ∈ SV(β, *c, a, p*).

13. **some** … if α is a common noun and β is a 1-place predicate, then SV('**some** (α) (β)', *c, a, p*) = 1 iff for some *x* ∈ O: *x* ∈ SV(α, *c, a, p*), and *x* ∈ SV(β, *c, a, p*).

13a. **the** … if α is a common noun and β is a 1-place predicate, then SV('**the** (α) (β)', *c, a, p*) = 1 iff there is exactly one *x* ∈ O, such that *x* ∈ SV(α, *c, a, p*), and for all *y* ∈ SV(α, *c, a, p*), *y* ∈ SV(β, *c, a, p*).

13b. **that** … if α is a common noun and β is a 1-place predicate, then SV('**that** (α) (β)', *c, a, p*) = 1 if there is exactly one *x* ∈ SV(α, *c, a, p*), such that *x* is the demonstratum of *c*, and *x* ∈ SV(β, *c, a, p*);

SV('**that (**α) (β)', *c, a, p*) = 0 if there is exactly one $x \in$ SV**(**α, *c, a, p*), such that *x* is the demonstratum of *c*, and $x \notin$ SV**(**β, *c, a, p*); and undefined otherwise.

14. **not** … if α is a sentence, then SV('**not**(α)', *c, a, p*) = 1 iff SV(α, *c, a, p*) = 0.

15. **and** … if α and β are sentences, then SV('α **and** β', *c, a, p*) = 1 iff SV(α, *c, a, p*) = 1 and SV(β, *c, a, p*) = 1.

15a. **poss** … if α is a sentence, then SV('**possibly**(α)', *c, a, p*) = 1 iff for some $p^* \in$ P such that T(p^*) = T(*p*), SV(α, *c*, *a*, p^*) = 1

15b **nec** … if α is a sentence, then SV('**necessarily**(α)', *c, a, p*) = 1 iff for all $p^* \in$ P such that T(p^*) = T(*p*), SV(α, *c*, *a*, p^*) = 1.

15c. **believes** … if α is a name and β is a sentence, then SV('α **believes that** β', *c, a, p*) = 1 iff for all $w^* \in$ W such that w^* is compatible with what SV(α, *c, a, p*) believes at T(*c*) and W(*c*), SV(β, *c*, *a*, <T(*p*), w^*>) = 1.

15d. **now** … if α is a sentence, then SV('**now**(α)', *c, a, p*) = 1 iff SV(α, *c*, *a*, <T(*c*), W(*p*)>) = 1.

15e. **was** … if α is a sentence, then SV('**it was the case that** (α)', *c, a, p*) = 1 iff for some $p^* \in$ P such that T(p^*) <T(*c*) and W(p^*) = W(*p*), SV(α, *c*, *a*, p^*) = 1.

15f. **will** … if α is a sentence, then SV('**it will be the case that** (α)', *c, a, p*) = 1 iff for some $p^* \in$ P such that T(p^*) > T(*c*) and W(p^*) = W(*p*), SV(α, *c*, *a*, p^*) = 1.

15g. **yesterday**… if α is a sentence, then SV('**yesterday** (α)', *c, a, p*) = 1 iff for some $p^* \in$ P such that T(p^*) is on the day preceding T(*c*) and W(p^*) =W(*p*), SV(α, *c*, *a*, p^*) =1.

15h. **always** … if α is a sentence, then SV('**always** (α)', *c, a, p*) = 1 iff for all $p^* \in$ P such that W(p^*) = W(*p*), SV(α, *c*, *a*, p^*) = 1.

15i. **day bef** … if α is a sentence, then SV('**the day before** (α)', *c, a, p*) = 1 iff for some $p^* \in$ P such that T(p^*) is on the day preceding T(*p*) and W(p^*) = W(*p*), SV(α, *c*, *a*, p^*) = 1.

15j. **later** … if α is a sentence, then SV('**at a later time** (α)', *c, a, p*) = 1 iff for some $p^* \in$ P, such that T(p^*) > T(*p*) and W(p^*) = W(*p*), SV(α, *c*, *a*, p^*) = 1.

Complex 1-place predicates:

16. λx_i: … if α is a sentence, then for all *i*: SV('λx_i(α)', *c, a, p*) = {*x*: for all $b \in$ A that differ from *a* at most in *i*-th place, if *x* = *b*(*i*) then SV(α, *c*, *b*, *p*) = 1}

The two "especially interesting" notions of absolute truth for (closed) sentence-context pairs that I discussed in section 4.1, are these:

(T^K) For all sentences α of L4 and contexts *c*: <α, *c*> is True iff for all *a* $\in$ A, SV(α, *c*, *a*, <T(*c*), @>) = 1.

(F^K) For all sentences α of L4 and contexts *c*: <α, *c*> is False iff for all *a* $\in$ A, SV(α, *c*, *a*, <T(*c*), @>) = 0.

(T^C) For all sentences α of L4 and contexts *c*: <α, *c*> is True iff for all *a* $\in$ A, SV(α, *c*, *a*, <T(*c*), W(*c*)>) = 1.

(F^C) For all sentences α of L4 and contexts *c*: <α, *c*> is False iff for all *a* $\in$ A, SV(α, *c*, *a*, <T(*c*), W(*c*)>) = 0.

The demonstrative "she"

Let me discuss some of the novelties in L4 that I have not discussed yet. In addition to **I** and **you**, I introduced a demonstrative name **she** in 3b. The idea is that in a context of use *c*, **she** picks out the female demonstrated by the speaker of *c* at the time, place and world of *c*. A demonstration is typically a gesture of pointing or some other kind of gesture. If there is no female demonstratum at *c*, then SV defines no value. This picks up the idea that if a speaker uses a demonstrative, but fails to demonstrate anything, then the utterance is unsuccessful at least in that respect. In the case of **she**, the demonstratum of the context needs to be in addition female. Thus if I utter '**dances** (**she**)' and the object I demonstrate is a man, or I fail to demonstrate anything at all, then the semantics of L4 says that I fail to express a sentence intension, because **she** does not have an intension defined for that context. This seems to be a good prediction, for utterances involving a demonstrative that are accompanied by no appropriate demonstration do create an impression of failed communication. Sometimes, of course, we will make pragmatic adjustments. For example, when an utterance of "she" is accompanied by a demonstration of an object that is not female, the audience might guess the speaker's intention, so that communication succeeds after all. There is also some room for tweaking the notion of a demonstrated object: the jury is out on whether speaker intentions play a role in determining which object is demonstrated. There is corresponding room for some tweaking in deciding which expressions count as demonstratives. Thus, if a candidate demonstrative seems to be used regularly and with success to communicate something about an object that is not demonstrated (say: gestured at), then this might be reason to re-classify and construe the context dependence of the candidate demonstrative in some other way. Thus some theorists distinguish the demonstrative sense of "she" and other personal pronouns from an anaphoric sense. Others might prefer a general and uniform account of personal pronouns such as "she", according to which in every appropriate context it is a melange of factors that makes an object the suitable semantic value for "she". These discussions need not detain us here. It will be enough to know that if we did think that a language we are modelling con-

tained a demonstrative whose semantic value in a context is the object demonstrated, then we know how to describe that demonstrative in the semantics.

Not a complete novelty is the semantic clause given for the determiner **the** in 13a. This treatment was already mentioned at the end of section §3.5. **The** combines with a common noun to form a quantifier phrase, just like the determiners **every** and **some** that we already know. The semantics here given (see 13a. and 9.) entails that '**the** (**woman**) (**frowns**)' is true in a context *c* and circumstance *p* just if there is exactly one thing that is a woman at the time and world of *p*, and that woman frowns at the time and world of *p*. This treatment is equivalent to Russell's famous treatment of definite descriptions (see Russell 1905, 1918/9). There are, of course, well-known problems with this sort of treatment. The most glaring inadequacy is, at first sight, that an utterance of "The woman frowns." in English is not normally regarded as false just because at the time and world of utterance there is more than one woman. There are various possible remedies. One is to insist along Gricean lines that such an utterance is indeed false, but that it carries a conversational implicature which may be true (namely an implicature that there is only one relevant or salient woman, who frowns, see Grice 1989 and Neale 1990 for discussion). Another remedy, which conserves the treatment of "the" as a determiner, and of "the woman" as a quantifier phrase, along the lines indicated in 13a., is to say that the domain of the quantifier is restricted contextually. In other words, the domain of objects mentioned in 13a. (the domain O) is restricted in a way determined by the context to the set of objects that are relevant or salient in the context (see Stanley and Szabó 2000 for dicussion). We might then get a slightly modified clause 13a. along the following lines:

13a*. **the** … if α is a common noun and β is a 1-place predicate, then SV('**the** (α) (β)', *c, a, p*) = 1 iff there is exactly one $x \in \{x: x \in O$ and x is salient in $c\}$, such that $x \in$ SV**(**α, *c, a, p*), and for all for all $y \in$ SV**(**α, *c, a, p*), $y \in$ SV**(**β, *c, a, p*).

The determiner "the"

Such a treatment makes **the** in effect an indexical expression. A third treatment would do away with considering "the woman" as a quantifier and treat it instead as a name, albeit, presumably a name that is again in some way context-sensitive. Again, we do not need to detain ourselves any further with these issues. Our main aim here is to survey the possibilities of Kaplan's framework.

The demonstrative determiner "that"

13b. is also worth brief consideration. This clause is so long because I treated **that** as a demonstrative determiner, so that it requires a context with an appropriate demonstratum on pains of the semantic value not be-

ing defined for that context. In effect it follows from 13b. that '**that woman**' has a very similar character as **she** has according to 2a (with the difference between "female" and "woman" accounting for all the differences between the two). Again, **that** creates sentences with non-constant character, it is indexical.

Kinds of temporal intensional operator

15d – j are clauses for intensional operators that are in the widest sense temporal. However, 15d – g. are indexical, while 15h – j. are not. Moreover, 15d – h. are intensional operators that remove sensitivity to the time parameter in the circumstance of evaluation, while 15i – j. are not. These are important differences, so let me explain.

Let us start with 15d – j. The propositional contents in Kaplan's framework are often "tensed" or "temporal" in the sense that they vary in truth-value over time. Contingency is the analogous phenomenon for the world parameter: we call contents that vary in truth-vale from world to world "contingent". The opposite of temporal or tensed may be called "temporally fixed". Thus a content that does not vary in truth-value over time is temporally fixed (the analogue in the modal case would be the property of being either necessarily true or necessarily false). Now, the operators in 15d – j., i.e. **now**, **it was the case that**, **it will be the case that**, **yesterday** and **always**, are all loosely speaking operators that turn temporal contents into temporally fixed contents, but they do so in different ways. Another way of saying this is that these operators turn functions from circumstances to truth-values (intensions) that are non-constant with respect to the time of circumstance into functions that are constant with respect to the time of the circumstance. Suppose the content expressed by some sentence α at a context *c* is a temporal intension. Then the content expressed by '**now** α' at that context is a temporally fixed intension, namely the one that takes for any circumstance the value that α takes when we shift the time of the circumstance to the time of *c*. This is what 15d. stipulates.

Operators that remove temporal variability

Similarly, but more complicatedly, for **it was the case that**. Again, loosely speaking, it turns a temporal content into a temporally fixed one. The truth-value of the content of '**it was the case that** α' at a context *c*, is 1 just if there is a time prior to the time of *c* at which α has the value 1. Or, more accurately, the content of '**it was the case that** α' at context *c* is a function from circumstance to truth-value which assigns the value 1 to a circumstance *p* just if the content of α at *c* has the value 1 for some circumstance that differs from *p* only in that its time is some time prior to the time of *c*. This is complicated, but there is no simpler way of saying it if we want to state it accurately. Once you have got your mind around the clause for **it was the case that**, it will be easy also to understand the clauses for **it will be the case that** and **yesterday**. For these are only simple variations. To get the hang of it, try to work out the truth-values at given contexts and

circumstances of a few sentences like '**now dances** (**Peter**)' or '**yesterday frowns** (**Sue**)'.

"Always"

Now, **it will be the case that**, **it was the case that**, and **yesterday** are all indexical expressions in that the intension expressed by sentences containing them will always depend on the context in which they are uttered. **Always** is not indexical. It also turns temporal contents into temporally fixed ones, but the way it does this is independent of context. If we ignore the world-parameter, the content of a sentence '**always** α' at a context is an intension that takes the value 1 just if the intension of α at that context takes the value 1 at all times. Or, more accurately: SV('**always** (α)', *c, a, p*) = 1 just if for all circumstances *p** that have the same world as *p*, SV(α, *c*, *a*, *p**) = 1. Since **always** does not introduce any context sensitivity, a sentence '**always** (α)' will only be context-sensitive if α already is. Otherwise '**always** (α)' has the same content in all contexts.

Iterability of temporal operators

The way I have defined these five temporal operators, namely as operators that produce sentences that express temporally fixed contents, applying two of these operators consecutively to a sentence will not make much sense. For the application of any of these operators only has any effect on the content expressed if the sentence to which they are applied express temporal contents. Thus '**now always frowns** (**Peter**)' will not differ in its semantic value from '**always frowns** (**Peter**)'. Similarly with '**always will dance** (**Sue**)'. This consequence suggests, I believe, that the future and past tense in English, as well as the expressions "now" and "always" are not accurately modelled by the operators I defined in L4. For consider the sentences "Peter always frowns now." and "Sue will always dance." in English. Clearly, the expression "now" in the first sentence and the expression "always" in the second are not redundant. Rather, they clearly contribute to the message conveyed. If the operators I defined in L4 were indeed the best model for the corresponding English operators, then quite a lot of work would remain to explain these cases pragmatically. So ideally, we should model these English expressions in a different way. This is a complicated task, and one that need not concern us here. The important point here is to understand operators of the sort stipulated in L4 in order to understand the possibilities of the framework so that we can bring them to bear on the phenomena that will concern us in coming chapters.

Temporal operators that do not remove temporal variability

Finally, let us consider the operators **the day before** and **later**, as defined in 15i–j. I introduced these primarily to contrast them with **yesterday** and **it will be the case that**. **The day before** and **later** are again intensional operators, but unlike the previous five, they do not turn a temporal content into a temporally fixed one, and unlike the first four temporal operators, but like **always**, they are not indexical. Suppose we have a non-indexical sen-

tence that expresses a temporal content, like **frowns** (**Sue**). When we evaluate the content expressed by '**the day before frowns** (**Sue**)' at a circumstance *p*, we have to look to the value **frowns** (**Sue**) gets for a circumstance that is like *p* except that its time is on the day before the time of *p*. In other words (ignoring worlds), the value of the complex '**the day before** α' at a certain time *t* is just the value α has a day before *t*. This shifts the time parameter, but it does not fix it. So, the resulting content is still temporal. The difference between the clauses for **the day before** and for **yesterday** is easily overlooked: where **yesterday** shifts the time of evaluation to the day before the time of the *context*, **the day before** shifts the time of evaluation to the day before the time at which we are evaluating the complex. This one-letter difference between 15g and 15i makes a huge difference. While **yesterday** makes contents temporally fixed, **the day before** does not. This means that **the day before** can be iterated, and can be embedded within further temporal intensional operators. '**The day before the day before** α' simply gets the value, at a time *t*, that α gets two days before *t*. Moreover, a sentence '**the day before** α' can be further temporally embedded, as in '**it will be the case that the day before** α'. This reflects the fact that in English we can say things like "I will have spoken to him the day before." but not "I will have spoken to him yesterday."

Fixing and shifting temporal operators

Studying 15f. and 15j. will reveal that **it will be the case that** and **later** are related to one another in exactly the way **yesterday** and **the day before** are. We can give a name to the difference between the first five temporal operators and the last two: the first five are operators that "fix" the time parameter of the circumstance, while the last two are operators that merely "shift" the time parameter. So we can speak of "fixing" and "shifting" temporal operators.

Again, the point is to make clear the subtly different ways in which the temporal variability of contents can be exploited by intensional operators. The difference between fixing and shifting operators can also be made in the case of modal operators. The ones we have defined here, **possibly** and **necessarily** as defined in L4, are *fixing* modal operators, that is, any content expressed by a sentence introduced by **possibly** or **necessarily** will not vary in truth-value over possible worlds. Similarly, an operator **actually** could be defined in exact analogy with **now**, which would make for another fixing modal operator, though this time an indexical one. We don't have them in L4, but we can easily imagine shifting modal operators. In fact, if instead of simplifying I had defined the modal operators in the standard way, then they would have been shifting operators. Standardly, necessary truth is defined as truth not in all worlds, but as truth in all *accessible* worlds. Here is how this would go in detail:

15b* **nec** … if α is a sentence, then V('**necessarily**(α)', *c, a, p*) = 1 iff for all $p^* \in$ P such that T(p^*) = T(p) <u>and W(p^*) is accessible from W(p)</u>, SV(α, *c*, *a*, p^*) = 1.

Defined in this way, **necessarily** is a shifting operator (at least as long as not all worlds are accessible from all worlds). Accordingly, sentences introduced by it can be non-redundantly further embedded in modal operators.

In order to provide a more precise model of the various candidate intensional operators in English we would have to observe carefully whether and how they are indexical, as well as whether and how they are fixing or shifting. For example, in order to account for the fact that "Peter always frowns now." can be used to convey that Peter frowns at all times in a range somehow indicated by "now", we could construe "always" as an indexical shifting operator of some sort. But modelling these candidate intensional operators is not our main concern, so we will now move on to examine what alternatives there are to Kaplan's view that many contents are temporal, i.e. his temporalism.

***Exercise 7**: Look back to the last paragraph of chapter 4.3 and attempt to construe „always" as a shifting operator as indicated!*

4.4 Temporalism and its Alternatives

I remarked in §4.1 that sentence (2): "I am hungry." seems context-dependent in the same ways in which sentence (1): "I am hungry now." does. That is, competent speakers' assessments of the literal truth of utterances of (1) and (2) will not differ. Utterances of both sentences will depend for their truth on who utters the sentence and when they utter it. However, while sentence (1) contains an explicit indexical expression, "now", which accounts for the dependence on time of utterance, sentence (2) does not. In Kaplan's framework this gives rise to a difference between the respective contents of (1) and (2) in a given context. While the content of (2) in any context is a temporal content, the content of (1) will be a temporally fixed content. Thus, the contents asserted by each of the two in a given context will be different.

Variability of content vs variability of truth-value

Let us look at two slightly different examples, so that we only have to think about time-dependence:

(8) MK is hungry now.

(9) MK is hungry.

Again, (8) will express a temporally fixed content while (9) will express a tensed or temporal content at any context. We will not easily notice this difference in users' assessments of corrrectness, because of the way we assess assertions of temporal contents. We grant the time of utterance special importance in the assessment of assertions of temporal contents, just as we grant special significance to the world of utterance (see the discussion at the end of §4.1 and (T^C) in §4.3). Thus, when someone utters (9), assessments of literal truth of their utterance will usually correspond to the truth of the temporal content asserted *at the time and world of utterance*. Of course, users can be mistaken about the time of an utterance and this can lead to a divergence of literal truth as adjudicated by them and Truth as defined by (T^C). However, this will lead to exactly the same divergences when the sentence in question is (8).

Contents as objects of thought …

Kaplan says at some points that his contents are the "objects" of thought. So perhaps we can use this idea to justify that the contents of (8) and (9) are different? Suppose you have two thoughts, one before lunch and one after lunch, and these are thoughts you would express sincerely by uttering (9) at the respective times. Is it plausible to say that the object of thought is the same? There clearly is a sense in which it is: each time it is a thought that MK is hungry. In that sense, the object of thought has not changed over lunch, but (supposing MK does eat a satisfying lunch) the truth-value of that object of thought has changed. But there is also a sense in which the object of thought is not the same: since the truth-value of the two thoughts is different, their objects must be different. Many follow Frege in construing the objects of thought (propositions) as having absolute truth-values. That's just what they mean by "object of thought". On that conception, clearly, the objects of thought are different.

… or assertion do not provide the motivation

The same can be said about the object of assertion: we may well have a conception of assertion and its objects according to which if the object asserted is the same, the assessment of literal truth of the assertion must also be the same. But there also seems to be a conception of assertion which allows assertions with the same object to differ in truth-value. There seems nothing prima facie incoherent about the way sketched in §4.1 of assessing the truth of assertions of temporal contents, according to which I can assert the same temporal proposition before and after lunch (by uttering (9)), and one comes out as literally true, the other as literally false. So, there does not here seem to be a clear motivation for Kaplan's particular approach to (9), according to which the two utterances of it express the same content.

Conditional motivation

But there is a certain motivation for this treatment which depends on a premiss: *if* "sometimes", "always", "yesterday", "on May the 3rd", "it will be the case that" etc are intensional operators, then we ought to treat (9) as

expressing a temporal content. For if it didn't, then applying "always" to (9) would be redundant. But it isn't. So (9) does express a temporal intension. This also explains why "always" or "on May 3rd" can be non-redundantly applied to (9), but not to (8). "On May the 3rd, MK is hungry now." is clearly odd, while "On May the 3rd, MK is hungry." is not. We do need to explain this, and a semantic explanation seems to be required. The joint assumptions that (9) expresses a temporal content while (8) doesn't and that "on May the 3rd" is an intensional operator provide this semantic explanation.

Content as "what is said"

However, the assumption that "on May the 3rd" and the other operators are intensional operators is resistable – and has been massively resisted (see King 2003 for an excellent discussion of this). Is there any other motivation for temporal contents in Kaplan 1977/1989? Kaplan sometimes says that his contents correspond to "what is said" by an utterance. But again, this does not seem to differ fundamentally from the question of what is asserted by an utterance. The object of assertion, what is said by an utterance of (9) can, I believe, with equal right be construed as a temporal or a temporally fixed content. There does not seem to be any conclusive reason that prevents us from saying that the object asserted or said by utterances of (8) and (9), when the utterances take place at the same time, are the same. In some sense at least, such utterances say the same thing. This is not the sense in which one might say that what was said by the utterance of (9) is usually true before lunch and false after. But it is a perfectly legitimate sense. I do not believe that there is any clear pre-theoretical sense, or any particularly theoretically privileged sense, of "what is said" which requires us to treat the objects of sayings and assertions as potentially temporal contents.

Alternatives to temporalism

In what follows I will explore how we might construe "on May the 3rd" and the other temporal operators not as intensional operators but as quantifiers or singular terms. This would allow us to abandon Kaplan's idea that (9) expresses a temporal content. The various ways in which this can be done will be instructive because the phenomena discussed in coming chapters will force us to face analogous choices: between an intensional treatment and a quantifier treatment of certain operators, and between construing certain sentences as expressing contents that vary in their truth-value with (some parameter in) the circumstance of evaluation or as expressing contents that are fixed with respect to that parameter.

The indexical view of tense

There are several ways of avoiding temporally unspecific propositional contents. Let us start with the idea that (8) and (9) do in fact have the same *character*, i.e. they express the same content at any context of use. This would mean that sentence (9), despite lacking any overtly indexical expression, is nevertheless an indexical sentence. We might flesh this out

by saying that it is the tense of the copula ("is") that introduces the indexicality. We would say analogous things about all the other sentences that are claimed to express temporal contents in Kaplan's system. (This should be OK even for the claim that it is the tense of the verb that introduces the indexicality, for sentences do in general contain a main verb, and any definite verb form has a tense.) In this case we no longer have sentences that express temporal contents. Just like (8), (9) expresses a temporally fixed content in every context, but which temporally fixed content that is varies with the time of the context. So let us call this first view the "indexical view" of tense.

Syntactic questions for the indexical view

The first problem for the indexical view of tense is that it needs to explain why (9) does and (8) does not admit embedding within temporal constructions like the ones we have mentioned, e.g. "always" or "on May the 3rd", or even "now". The second problem is that it needs to give a positive account of what constructions like these actually do to tensed sentences like (9). "Always", being equivalent to "At all times", seems quite clearly to involve some kind of quantification over times (as it does in a Kaplanian language like L4), while "on May the 3rd" and "now" seem somehow to involve reference to times. But if (9) is already a complete indexical sentence like (8), then what could these expressions contribute? What is their syntax?

A proposed answer

It seems that we would need to say that the indexical contribution made by tense in (9) is somehow de-activated when (9) is embedded in these ways. Suppose we think of the logical form of (9) as

(9F) be hungry (MK, now)

conceiving of "be hungry" as a two-place predicate one argument place of which is taken by "MK", the other argument place is taken by the present tense of "is hungry", here construed as a singular term "now" referring indexically to the time of utterance. What would then be the logical form of (8) or of "MK is always hungry"?, or of "MK is hungry on May 3rd"? The obvious suggestion is this:

(8F) be hungry (MK, now)

(10) MK is hungry on May 3rd

(10F) be hungry (MK, on May 3rd)

(11) MK is always hungry.

(11F) For all times *t*: be hungry (MK, *t*)

(11F) is also supported by the fact that "At all times, MK is hungry" seems to be equivalent to (11). How would we explain that singular terms like "now" and "on May 3rd" can combine with a sentence? Moreover, how would these expressions interact semantically with (9)? It seems that the only option is to say that the indexical element that makes reference to the time of utterance in unembedded occurrences of (9) is somehow inert when (9) occurs embedded. This inertia in effect means that (9), when embedded, functions syntactically and semantically like a one-place predicate. In the case of (8), this predicative expression is combined with a singular term indexically referring to the time of utterance. Hence (8) and (9) express the same content. In the case of (10) the singular term is not indexical. In the case of (11), the predicate is combined with a quantifier phrase.

Not compositional …

The problem with this is that now the semantic value of the various sentences properly containing (9) have contents at a context that are not determined by the contents of their constituents at that context. For it is not the content of (9) at a context that contributes to the content of, for example (10) at that context, but rather the content of a proper part of (9). This would mean that we have to give up on our idea that the content of a sentence in context is a function of the contents of its immediate constituents in that context.

… or unnecessary complication

If we want to avoid this, we could construe "now", "on May 3rd" and "always", not as names or quantifier phrases, but rather as operators that shift the time parameter of the context. Kaplan claimed that in English, there are no operators that shift contextual parameters – he called such operators "monsters". But he had in mind indexicals like "you", "today", "here" and "now", not the tenses of verbs. In any case, Kaplan did not give any argument why there couldn't be such operators, and in fact some theorists argue that natural languages do contain monsters (see e.g. Schlenker 2003). However, the point of introducing contexts of use and character seems in part to have been precisely this feature of indexicals, that they are unshiftable, which requires us to treat the time and world of the context separately from the time and world of the circumstance of evaluation. So, this would cause a loss of elegance. But there seems no reason in principle to doubt that one can treat temporal operators as context-shifting operators. The indexical theorist of tense would have to admit, though, that (8) and (9) are not *exactly* alike in their indexical character. For (8) would be unshiftably temporally indexical, while (9) would be shiftably temporally indexical. This would have the disadvantage that we would now be distinguishing between shiftable and unshiftable contextual parameters, which was just the contrast for which we introduced the distinction between context and circumstance. It would be much more straight-

forward and elegant to treat the shiftable content of (9) as temporal in the way Kaplan does.

So the indexical view will either give up the compositionality of contents, or it will introduce a further distinction between shiftable and unshiftable contextual parameters. Both seem to be disadvantages vis-à-vis Kaplan's temporalist account.

What alternatives does an anti-temporalist have to the indexical view? In other words, can we avoid saying that (9) expresses a Kaplanian temporal content without claiming that (8) and (9) have the same content in the same context? There are two possibilities. According to the first, (9) is not a sentence, but rather in the syntactic category of a complex 1-place predicate. (8), (10), and (11) are then just the normal result of combining a one-place predicate with singular terms or quantifier phrases. These would be the logical forms:

(8F') be hungry (MK, now)

(9F') be hungry (MK,)

(10F') be hungry (MK, on May 3rd)

(11F') For all times *t*: be hungry (MK, *t*)

The predicate view of tense

On this view, the problem is that (9) seems to be a full-fledged sentence, and is perfectly usable, on its own, for communicative purposes. But one-place predicates cannot usually be used to say something. So how can that be explained? The view might be that while (9)'s content is merely the content of a predicate, expressing that content can nevertheless serve to convey proper sentence contents. Perhaps the audience can reconstruct, and the speaker intends them to reconstruct, a full propositional content from the predicate content semantically expressed. Let us call this the "predicate view of tense".[53]

The second view of this kind does say that (9) is syntactically a *sentence*, but it is a sentence that contains an implicit unbound variable. The temporal operators are then expressions that can bind variables. They might be syntactically like a binder together with a name in L4. Or they might be like a quantifier phrase together with a binder. Thus they will be syntactically a sentence operator, but they can non-redundantly operate only on senten-

53 The predicate view is probably the truest successor of detensing approaches, such as Quine 1960. See Burgess 1984.

ces that contain unbound variables. These might be the corresponding logical forms:

(8F'') λt(be hungry (MK, t)) (now)

(9F'') be hungry (MK, t)

(10F'') λt(be hungry (MK, t)) (3rd of May)

(11F'') All times (λt(be hungry (MK, t)))

The free variable view of tense

Let us call this view the "free-variable view of tense". Again, this view has to explain, how a sentence with an unbound variable, which it says (9) is, can be used to say anything, for it will have truth-values only relative to an "assignment". The answer will have to be that the context will somehow help identify assignments that are especially interesting for an evaluation of utterances of the sentence. Does that sound familiar? It should, for we answered the corresponding question for Kaplan's temporal contents already: how can one assess an assertion of a temporal content for literal truth? The answer then was that there are various ways of assessing the assertion by way of looking at the content, and that usually the truth-value of that content at the time of utterance was particularly interesting. The free variable theorist can make analogous moves: when using an open sentence and thereby expressing a content that has truth-values only relative to an assignment, we are particularly interested in its truth-value relative to the assignment that assigns the time of utterance as value to the variable.

Thus, the free-variable view uses the apparatus of assignment dependence that we introduced to handle the binding of variables, to handle the time dependence of sentences like (9). In my view, this is the most attractive version of all the non-temporalist views, as it exploits this already existing structure in the semantics. It is in fact analogous with Kaplan's view, the only disanalogy being that in Kaplan's view, tense operators do not involve variables and binding.

It seems, then, that we could remove the time-component in the circumstance of evaluation, as the indexical view, the predicate view and the free variable view each provide an alternative that does without it. One might argue about which is the more elegant view. King 2003 forcefully argues that something much like the free variable view is more elegant.

Could the assignment parameter handle all intensional phenomena?

But the more fascinating question, surely, is whether we could have done without intensionality altogether, and handle even modality and other intensional phenomena with the apparatus of assignment dependence. Would that work equally well? On this view, the only closed sentences

would be the ones we are used to think of as expressing necessary truths, and the ones expressing contradictions. All other sentences would now be treated as involving some unbound variable ranging over possible worlds. Thus contingency would be an attribute that could be said of the functions from assignments to truth-values that the semantics assigns to each sentence-context pair. The differences between world dependence and time dependence would just be a difference in the ranges of certain variables.

The reason I have been exploring in detail the alternatives to adding a time parameter to the circumstances of evaluation is simply that the next chapter will examine arguments that purport to show that we should add even more circumstantial parameters. It will be useful in examining these arguments, to be familiar with the alternative moves that one can in general make to avoid adding further circumstantial parameters.

4.5 Summary

In this chapter, I began by showing that the simple intensional framework from chapter 3 is not adequate to model the context-dependence of many natural language expressions. It seems clear that sentences like "Your uncle phoned yesterday." or "I dance here." depend for their truth on certain aspects of the context in which they are used. This dependence of semantic value on utterance context needed to be allowed for in the semantics. I explained how Kaplan in fact modifies simple intensional semantics in two major ways: first, he adds a context of use parameter, so that his semantics assigns intensions no longer to sentences but to sentence-context pairs. He also introduces another truth-determining factor that functions just like the world-dependence familiar from simple intensional semantics. Thus, what sentences express at contexts is not simple intensions – functions from possible worlds to truth-values – but potentially temporal intensions, i.e. functions from world-time pairs, i.e. circumstances of evaluation to truth-values. I explained some of the aspects of the sometimes complex interaction between the circumstance dependence and the context dependence of sentences, and I illustrated some of the possibilities of the Kaplanian framework by modifying L3 and adding a range of new expressions, to create L4. One of the features of Kaplan's framework that is rejected by most semanticists today is his temporalism, i.e. his treatment of the phenomena of tense and time-dependence within the intensional apparatus. I therefore explored in some detail what the alternatives to Kaplan's temporalism are, so that we will be able to consider similar options in later chapters when considering the merits of analogous intensional proposals.

5. Extending Kaplan's Framework: Relativism

5.1 Introduction: Assessing Claims on Non-objective Matters

In the last chapter we saw how a Kaplan-style double-index semantics makes room for indexicals, i.e. expressions the intension or content of which varies with the context in which they are used. In this framework, the semantics assigns to each sentence a *character*. A character is a function from contexts of use to *contents*, which in turn are functions from circumstances of evaluation to one of two truth-values. We also saw that Kaplan's contents can be temporal, so that the truth-value of a content can vary with time as well as possible world. Hence Kaplan's circumstances of evaluation consist of two parameters, times and worlds. This chapter will consider some recent proposals to add further parameters to the circumstances of evaluation, such as standards of taste or states of information.

Patterns of assessment data so far

Broadly speaking, the motivation for adding further parameters – be they of context or of circumstance – always concerns certain perceived contextual variations, which, it is then argued, are best accommodated by adding the further parameter. In the case of ordinary indexicals, such as the sentence "Today is Monday.", the contextual variation can be clearly read off competent users' assessments of literal truth: competent users will largely concur in their assessments of some utterances of the sentence as (literally) true and others as false. Which ones they evaluate as true will clearly depend on the day of utterance. Even if there are divergences in assessment amongst users, there will hardly be anyone who assesses all utterances of the sentence, no matter on which day they occur, to have the same truth-value. If anyone did, then this would be taken as a clear sign of incompetence. Thus, all clearly competent users will concur in the variability of the truth-value of the sentence from utterance to utterance. So we have very good evidence that utterances of the same sentence vary in truth-value. Considerations like the ones in the last chapter will then lead us to treat this variation either as due to the sentence having a non-constant *character* (different content in different contexts of use) or non-constant *content* at a given context of use (e.g. a temporally variable content).

The cases I wish to discuss in this chapter exhibit a different and more complex pattern of contextual variation. When I say on Monday "Today is Monday", and you assess my utterance on Tuesday, you will presumably say that my utterance was true. It will not matter for your assessment that you are assessing me on a Tuesday. When you say "I am Italian.", I will assess your utterance as true if I think you are Italian, and it will not matter

for my assessment that I, the assessor, am not Italian. Now suppose you utter the following two sentences:

(1) Hip-hop is better than swing.

(2) Pickled herring is tasty.

One type of assessment

If I am asked to assess the truth of your utterances, I will probably be unsure as to what sort of assessment I am asked to provide. If given no further instructions, I will most likely compare hip-hop and swing and answer accordingly, without taking into account the context in which (1) was uttered: in other words: I will say the utterance is true just if I think hip-hop is better. Similarly, in assessing the truth of the utterance of (2), I will ask myself whether pickled herring is tasty, and answer accordingly, without taking into account who uttered the sentence where and when. If I am careful, I might clarify my assessments and say, about the utterance of (1): "*What you said* is (is not) true.", or "Hip-hop is (is not) better.". Similarly, I might express my assessment of your utterance of (2) with "That's true." or "It is indeed tasty.".

No variation across contexts

Let us, for the time being, interpret assessments of this type as the assessments of literal truth that we are trying ultimately to predict with our semantics. Then we should first notice that these assessments are insensitive to the context in which the utterances were made. No matter, where, when and by whom the two sentences are uttered, any given competent users' assessments of this type will remain the same. Thus it looks like these sentences are not context-dependent.

Variation across assessors

However, there will be significant divergences in the assessments made by different users. If we suppose, realistically, that a significant proportion of language users prefer hip-hop to swing, while the remainder, also significant, does not, then we should expect the assessments of competent users to be highly heterogeneous: many will assess the utterance of (1) as true, but many others will assess it as not true. Similarly with the utterance of (2): suppose, realistically, that there is a significant number of people who like the taste of pickled herring, while there is also a significant number of people who don't, then the assessments of truth of the utterance of (2) will be heterogeneous.

In our framework, there seem to be two ways of explaining this data pattern: we could either say that one significant group is committing some kind of error, either by getting the facts wrong or by being incompetent with some of the expressions. Or we could say that some of the expressions in question do not have definite public meanings, that users are not co-ordinated in their use of them.

To illustrate this, the situation is similar to the data we might collect regarding these sentences:

(3) Human activity has caused global warming.

(4) Whenever one rejects a claim, one refutes it.

Massive error

Let us assume, realistically, that there are significant numbers of people who believe that human activity has caused climate change, but also significant numbers who disbelieve this. Moreover, let us assume that significant numbers of people use the verb "to refute" as a synonym of "to deny", and also significant numbers of people who use this verb as synonymous with "to disprove". Clearly, our pattern of assessments by competent users of utterances of (3) and (4) will be similar to that concerning (1) and (2): while individual users will not vary their assessments with the context of use, they will diverge in their assessments amongst each other. Now, in the case of (3), this data pattern is explicable by the fact that different people have been exposed to different evidence regarding the causes of global warming, or that different people have different levels of competence in assessing the evidence. Thus, presumably one of the two groups will be just have mistaken beliefs about the causes of global warming.[54] Call this an "explanation by massive error".

Linguistic divergence

In the case of (4), an explanation by massive error is implausible, for it is highly unlikely that significant numbers of people erroneously believe that every denial amounts to a successful disproof. Thus in the case of (4), the obvious explanation is that different users use the verb "to refute" differently, i.e. that there is a divergence of competence. Call this an "explanation by linguistic divergence". In the case of (3), we can exclude an explanation by linguistic divergence by asking users how they would evaluate (3) under the assumption of various scenarios that describe what exactly humans did, how climate is influenced by various factors etc. If users mostly concur in these conditional assessments, then that confirms that there is not linguistic divergence, but rather substantial disagreement between them.

Neither linguistic divergence nor massive error

In the case of (1) and (2), neither an explanation by massive error, nor an explanation by linguistic divergence seems satisfactory. It seems intuitively wrong to say that the different users diverge in the meanings they attach to the sentences, i.e. that they in fact use different languages. For it is precisely *because* of their competence with the same expressions and com-

54 A note of caution: it has been argued (e.g. Menzies 2007), that causal claims are context sensitive. I am currently assuming for the sake of argument that this view is wrong.

pliance with the same rules that they diverge in their assessments. If the observations in chapter 2 were correct, then speakers are calibrated to adjust their judgments of what is tasty, or what is better than what aesthetically, to their own relevant personal features, in this case their aesthetic preferences and tastes. It is because speakers employ the relevant concepts competently, and interpret the relevant expressions as expressing these concepts, that they diverge in their assessments. Thus, it is implausible to say that the divergent assessments are a symptom of massive error or of massively divergent language use.

Assessor-centred assessments

If we want to avoid a conclusion of massive error or failure of linguistic calibration, we must look at our data in a more differentiated way. When asking users to assess the truth of an utterance, what are we asking them to do? Pre-theoretically, there are at least two clear forms of assessment we might have in mind. First, we might be asking an assessor to judge whether what was said by the utterance is true, in the sense in which I have been interpreting the assessments above. Assessing what someone said as true in this sense signals agreement in an intuitive sense. In this sense a user should assess an utterance as true just if she could sincerely respond to the utterance by saying, for example, "That's true.". When assessing what was said by an utterance of "Today is Monday." in this sense, I will consider the day on which the utterance was made, and whether that day was Monday. However, when assessing an utterance of (1) in this sense, I will ask myself whether hip-hop is better than swing, and I will judge it to be true if I think that hip-hop is better than swing. Let's call this form of assessment "assessor-centered". Assessor-centered assessments of (1) and (2) by a given assessor will not vary from one utterance to another. However, there will be significant divergences from one assessor to another.

Speaker-centred assessment

Secondly, when asking for an assessment of an utterance, we might be asking whether an utterance is correct in the sense of manifesting a correct use of language and a correct employment of concepts. This will not make a difference for the assessments of utterances of "Today is Monday.". But it will lead to different assessments of utterances of (1) and (2). Now assessors will consider whether *the speaker* of (1) prefers hip-hop to swing, rather than whether hip-hop is better than swing, period. If the first is the case, they will say that the utterance is correct. Given our realistic assumption that there are significant numbers of people who prefer hip-hop to swing, but also significant numbers that don't, we will now get assessments of correctness that vary from utterance to utterance of (1). Given our assumption that significant numbers like pickled herring, and significant numbers don't, the assessment of (2) will now vary from utterance to utterance. Let's call this form of assessment "speaker-centered".

Divergence of assessor- and speaker-centred assessments

If what I am saying is correct, then assessor-centered and speaker-centered assessments will coincide in many cases, but not in assessing utterances of (1) and (2). When someone asserts that pickled herring is tasty by uttering (2), and also believes that pickled herring is tasty, we should be able to say that the *asserter* has fully complied with all the norms of assertion and belief without thereby being committed to saying that pickled herring is tasty, that what was asserted is true and thereby to agreeing with what was asserted. It seems that in some cases assessor-centered and speaker-centered assessments can diverge. If we want to avoid explanations from massive error or from linguistic divergence then we have to take into account speaker-centered assessments when modelling linguistic competence.

Non-indexical relativism about taste

It is not clear how the framework, as currently outlined, can make room for this. For if we treat (1) and (2) as having constant character, i.e. expressing the same content in all contexts of use, then the only parameters with which the truth-value of these contents can vary are the time and world parameters in the circumstance of evaluation. Thus, if speaker-centered assessments are to vary from utterance to utterance, then in the formal model, these assessments would have to correspond to the value of the content at different circumstances. Given that time does not seem to be relevant, this would mean that speaker-centered assessments track the value of the content at different possible worlds. But why should the speaker's compliance with the norms of assertion and belief be measured in terms of truth at one possible world, while the assessor herself measures the truth of the asserted content in terms of truth at another possible world? This would seem to violate against one of our basic assumptions in chapter 2: that we all share the same world.[55] So, if we are saying that (1) and (2) have constant character, then we need to add a new circumstantial parameter. Thus, in addition to contingent and temporally variable contents, we would need *gustatorily variable* contents, i.e. contents that vary with something like a standard of taste parameter. Let us call this kind of view "non-indexical relativism about taste".

According to non-indexical relativism about taste, the contents of assertions made by uttering (1) or (2), which are also the contents of the beliefs expressed by these assertions, receive different truth-values relative to

55 One way to respond to this would be to retract the assumption that we all share the same world, and replace it, perhaps, with the claim that while we don't share the same world, our different worlds at least overlap substantially. This is the approach proposed by Einheuser 2008, who calls the aspects of a world that are shared by all the "substratum" of that world. I see no reason to think that this is more than a terminological choice and will continue to take it as given that we share the same world with anyone we ever talk to.

different standards of taste. The standard of taste relative to which we have to assess the correctness of these assertions and beliefs can be different for different asserters or believers. This explains the variation in speaker-centered assessments of different utterances of (1) and (2). It also explains the variation of assessor-centered assessments from one assessor to another.

Indexical relativism

But there is also a second option. Instead of adopting non-indexical relativism, we might adopt "indexical relativism about taste", namely the view that it is the content of utterances of (1) and (2) that varies from utterance to utterance and explains the variation in speaker-centered assessments.[56] The suggestion is that (1) and (2) are indexical in that they make hidden reference to the speaker's standard of taste or preferences. So (1) and (2) are somewhat similar to "On my standard of taste hip-hop is better than swing." and "According to my taste, pickled herring is tasty.". Or perhaps: "I prefer hip-hop to swing." and "I like pickled herring.". Now, clearly, this form of indexicality would differ from the forms we have so far admitted into the framework. First, the indexicality is not as obviously visible in the surface form of the sentences, i.e. there is no expression that is obviously responsible for the indexicality. Secondly, as we have already seen, speaker-centered and assessor-centered assessments do not usually diverge when we are dealing with ordinary indexical sentences. This is connected to the fact that ordinary indexicals do not seem to admit of "monsters" (see §4.4), i.e. are not shiftable by embedding. More about this below.

The sorts of issues I have raised in connection with (1) and (2) also arise in a range of other cases. In each of these cases, there are reasons to consider indexical and non-indexical relativism ways to explain certain contextual variations. The range includes sentences ascribing knowledge, sentences attributing probability or epistemic modality, conditionals, sentences that are otherwise evaluative, for example morally evaluative, sentences describing causal interactions, and more. There is a growing literature about how these cases should be dealt with.

---

56 The term "indexical relativism" is from Wright 2001. Since similar problems have to some extent been discussed independently in different areas of philosophy, the terminology is not consistent. Indexical relativism has also been called "contextualism", especially in the literature about the context dependence of knowledge, or more recently "indexical contextualism" (e.g. Cohen 1986, DeRose 1992). In Metaethics, indexical relativism has also just been called "moral relativism" (Harman 1975). Non-indexical relativism has also been called "genuine relativism" (Kölbel 2005), "non-indexical contextualism" (MacFarlane 2008). Some (Stanley 2005, Hawthorne 2004) use the term "invariantism" for positions that claim that the semantic content remains invariant in the relevant sense (thus non-indexical relativists are invariantists and indexical relativists deny invariantism). Some distinguish "radical" relativism from "moderate relativism" or "non-indexical contextualism" (Recanati 2007).

The purpose of this chapter is to examine systematically the issues that are at stake in these debates, and to do so in relation to the general Kaplanian framework outlined in previous chapters. I shall begin this discussion by going back to the framework and re-examining how exactly context-dependence broadly conceived can arise in it. In §5.3, I shall then examine indexical relativist proposals for matters of taste, and move on in §5.4 to the alternative, non-indexical relativist proposal. In §5.5, I shall examine further proposals along the lines of the predicate view and the free-variable view that I explored in §4.4 above. These discussions of one case – that of matters of taste – will be taken to be largely representative of many of the other cases in which analogous proposals have been discussed. I shall survey a range of these other cases in §5.6.

5.2 Varieties of Context Dependence

In the model language L4, the semantics defines a function SV from quadruples <*e*, *c*, *a*, *p*> to appropriate extensions, where *e* is an expression, *c* is a context of use, *a* is an assignment of objects to variables, and *p* is a circumstance (point) of evaluation. If we restrict ourselves to sentences that do not contain free variables, then this function SV determines for each sentence-context pair the function from circumstances to truth-values that is its content.

Review of the predictive role of semantics

In order to place this in the wider context of the purpose of semantic theorizing, let's remind ourselves that this function is supposed to generate assignments of contents that are testable against competent speakers' assessments of the correctness of utterances under given conditions. If the semantics predicts that a certain sentence-context pair has a certain content, i.e. function from circumstances to truth-values, then this should somehow correspond to the assessment language users make of whether utterances of (surface) sentences would be correct. But the semantics does not make any content assignments directly to *utterances* or *surface* sentences. Rather it assigns contents to *logical form* sentences paired with *contexts in the technical sense of agent-location-time-world quadruples*. Thus, before we can bring to bear the semantics on an actual utterance event, we need to decide which pair of a deep structure level sentence and context in the technical sense adequately represents the utterance event.[57] This involves three decisions: decide which surface sentence was uttered, which deep structure sentence it represents (disambiguation), and which context in the technical sense corresponds to the utterance event. The diagramme above shows this.

57 This important point is made in Predelli 2004.

Extra-semantic factors in prediction of truth-value

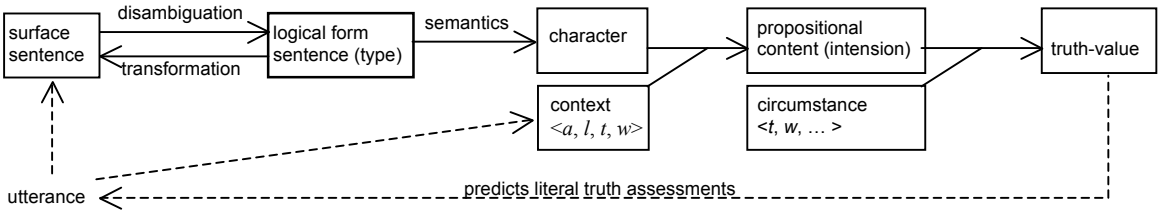

Let us review the aspects of this overall theory that are not semantic in the strict sense, i.e. do not concern the function SV, i.e. the compositional assignment of semantic values. Let us assume that it is unproblematic to decide which surface sentence was uttered in a given utterance event, i.e. that the theorist can parse the sounds produced and recognise the expressions of L4 that were used. Let us further assume that we are able, with the help of a transformational component to determine which logical form sentence or sentences represent the utterance. If this does not yield a unique logical form sentence, then we have a case of ambiguity.

Ambiguity resolution

We will now need to consider, presumably with the help of contextual information in a wide sense, which of these logical form sentences is the most relevant one. In other words, considerations of the previous discourse, of the likely communicative purposes of the speaker, etc will help us disambiguate the surface sentence. Different ways of disambiguating a surface sentence can lead to different assignments of character, and hence to different assignments of content and ultimately truth-values. This is the first way, in which extra-semantic considerations can impact on our predictions of literal truth.

The relevant context of utterance

Next, there is the question of which context in the technical sense is representative of the utterance event. Now, this may seem to be straightforward, for we might just take the context <*a*, *l*, *t*, *w*> which consists of the person *a* who has produced the token of the surface sentence, the location *l* where it was produced, the time *t* at which it was produced and the world *w* in which it was produced (i.e. the actual world for all actual utterances). However, we need to be cautious, for sometimes the person producing the token is not the intuitive referent of the indexical "I", nor is the place and time of producing the token always the intuitive referent, respectively of "here" and "now". For consider an answer-phone message "I am not here now.", which was not recorded by the owner of the answer-phone.[58] Clearly, the person who has made the recording (which is a token of the sentence in one sense) is not the person of which the recording says that she is not there then. Nor is it the person who has triggered the playback of the recording (which is also a token of the surface sentence, in

58 See Predelli 1998.

another sense). Moreover, the time and place of recording are not relevant, nor is the place of the person triggering the playback, or the place of the answer-phone (it might just be a remote server). Presumably in this case we should represent any particular playback of the recording by the context consisting of the user of the phone-line, the time of the playback and something like the intended place of the phone line (whatever that is). There is therefore some room for variation in the choice of the relevant context in the technical sense. Again, different choices might generate different character, hence different content and ultimately different truth-value. So, this is a second way, in which pre-semantic considerations can impact on predictions the semantics might make of the truth of a declarative utterance.

Once pre-theoretical considerations have yielded a unique sentence-context pair, the semantics takes over, and we can now work out the truth-value of the utterance using, for example, one of the definitions of truth for sentence-context-pairs I introduced in section 4.3:

(T^C) For all sentences α of L4 and contexts *c*: $<\alpha, c>$ is True iff for all $a \in A$, $SV(\alpha, c, a, <T(c), W(c)>) = 1$.

Prediction of truth-value of sentence-context pair

According to (T^C), a sentence-context pair is True just if the associated content is true at the circumstance *of* the context, i.e. the circumstance consisting of the time and world of the context.

Imperfect assessors

I suggested earlier that in order for competent users' assessments of utterances to correspond exactly to the predictions L4 makes of Truth in terms of (T^C), the users in questions should not make any mistakes as to the world and time of the utterance. So, if an assessor believes that Sue dances at the time of utterance, he or she will say that an utterance of the sentence "Sue dances" is true. If this belief is false, however, the assessment will not coincide with the semantics' prediction of Truth as defined by (T^C). But we can still utilise the assessment in confirming the semantics because we can take the semantics to predict that a competent user will judge the utterance to be true just if sentence-context-pair is true at the time she takes to be the time of utterance and at any world compatible with what she believes.

There are therefore two ways in which the semantics of a sentence can give rise to assessments that vary from utterance to utterance: either the sentence is indexical, i.e. its content varies from context to context, or it is temporal, i.e. the (otherwise invariant) content is temporally variable, as in the case of the content expressed by "MK is hungry.".

In principle, we thus have four ways in which the truth-value assigned to a (closed) sentence can be sensitive to contextual factors in a broad sense, two pre-semantic and two semantic:

(i) Ambiguity (several logical forms correspond to one surface sentence, e.g. "Every man loves some woman.")

(ii) variability as to which technical context is relevant to an utterance (e.g. "I am not here now.")

(iii) Indexicality (content varies with context: e.g. "I am hungry now.")

(iv) Circumstance sensitivity (truth-value varies with circumstance of evaluation: e.g. "MK is hungry.")

It will be worth mentioning also a source of *apparent* contextual variation in semantic value. Often sentences are used to convey something other than their semantic content, as in the case of metaphor or (other forms of) conversational implicature. This does not give rise to varying contents or contents that vary in truth-value, but rather to a contextual variation in the message conveyed:

(v) Pragmatic variation (while the semantic value does not vary, the message conveyed can vary: e.g. "I could eat an ox.")

Ways in which the semantic value of the same sentence can be variable

So far we have been concerned with competent users' assessments of the correctness of *literal* content, so (v) should not be relevant. But as I noted in chapter 3, it is not always obvious how the line between literal and non-literal meaning is to be drawn. So it will sometimes be controversial whether a given assessment is an assessment of literal content (and therefore directly relevant to the semantics) or whether it is an assessment of some non-literally conveyed content (and therefore only indirectly relevant to the semantics).

5.3 Indexical Relativism about Taste

Assessment data point to either circumstantial or contextual dependence

In §5.1 I argued that speaker-centered and assessor-centered assessments of utterances of sentences about matters of taste like (1) and (2) can diverge, and that they will diverge because the norms governing concepts like the concept of tastiness (to which competent users are subject) require that these concepts be applied in accordance with the applier's personal preferences. I argued that if we want to avoid having to treat the heterogeneity of assessor-centered verdicts across different assessors as a symptom of either massive error or linguistic divergence, then our se-

mantics needed to model speaker-centered assessment data. These data, in turn, suggest that the sentences in question are either indexical (as proposed by the indexical relativist) or that the contents they (invariably) express are gustatorily variable, i.e. vary in truth-value with a new circumstantial parameter whose values are standards of taste.

New argument

In this section, I shall explore the indexical relativist option in more detail. But before I do this, let me present a slightly different kind of argument in favour of postulating a new parameter of dependence (whether contextual or circumstantial). This argument is better in that it does not rely on the kind of "realistic assumptions" I was making earlier. In other words, we do not need to assume that there are significant numbers of users that have opposing views on the relevant question.

Suppose Anna believes that Depp is more handsome than Pitt and Barbara believes that he is not. The relational concept of *being more handsome than* is among the concepts that are taught polyphonically, just like *tasty*, *delicious* or *being (aesthetically) better than*. Given what I said about the norms governing such concepts, they are governed by conceptual norms that require that thinkers apply them in accordance with their own standards and preferences. Thus, it is correct for someone to judge that Depp is more handsome than Pitt only if that person prefers Depp to Pitt in the relevant respect. Since it is possible that Anna prefers Depp to Pitt in the relevant respect, and that Barbara does not, it is possible that each has applied the concept of *being more handsome than* correctly. It is possible furthermore that both Anna and Barbara have access to all the relevant information and have completely adequate justification for their respective beliefs. Now, let's imagine exactly such a case and consider the following two sincere utterances where Anna asserts what she believes and Barbara responds to her by asserting what she believes:

(5a) Anna (in C1): "Depp is more handsome than Pitt."

(5b) Barbara (in C2): "Depp is not more handsome than Pitt."

Hypothesis

These two assertions are clearly not in violation of any linguistic or conceptual norm: each of the two believes exactly what the conceptual norms require and sincerely expresses that belief by making the assertion. Both assertions are free from fault.

Reductio argument

Now suppose (for reductio) that the sentences used are not indexical and the only circumstance-sensitivity their contents exhibit is temporality and contingency. The sentence uttered by Barbara is the negation of that ut-

tered by Anna. So if our treatment of negation is as in L4 (i.e. standard), then the content expressed by Anna's utterance is true at a circumstance just if the content expressed by Barbara is not true at that circumstance. Assertion (5a) should be evaluated with respect to the time and world of C1, i.e. circumstance <T(C1), W(C1)>), while assertion (5b) should be evaluated with respect to the time and world of C2, i.e. circumstance <T(C2), W(C2)>. However, since utterance (5b) is in direct response to (5a), the circumstances with respect to which their assertions should be evaluated do not differ significantly: the world is the same (the world in which both utterances take place), and the time is roughly that of the two utterances (if we want to construe the relevant times to be distinct, they will nevertheless not be *relevantly* different, as the comparative handsomeness of Depp and Pitt cannot change sufficiently quickly). Thus we have to conclude that if (5a) is correct (true at <T(C1), W(C1)>), then (5b) is not correct (true at <T(C2), W(C2)>), and if (5b) is correct then (5a) is not correct. This contradicts our hypothesis that both Anna's and Barbara's assertion are fault-free. So if we want a semantics that makes room for the scenario of our hypothesis then we should reject the supposition. The supposition was that the sentences used are neither indexical nor circumstance sensitive over and above temporality and contingency. So we should construe the two sentences either as indexical or as expressing contents that depend for their truth-value on a factor over and above the time and world of the circumstance.[59]

The general indexical proposal

Let us then begin by exploring the possibility that the two sentences are indexical. They would need to be indexical in such a way that the content expressed by "Depp is more handsome than Pitt." varies between C1 and C2. What kind of indexicality could be responsible for this? On the face of it, the sentences uttered do not contain any indexical elements. So, if we are to accept the conclusion that there is indexicality here, we will have to say that it is beneath the surface. Now, the names "Depp" and "Pitt" do not seem to be context-sensitive in any way that could help here. So the assumption must be that the predicate "is more handsome than" expresses different intensions in different contexts. Anna's use of the predicate ex-

59 In general, there would be another source of contextual variation that could preserve the faultlessness of utterances like (5a) and (5b), namely ambiguity. I tried to preclude this possibility through the way I described the scenario. If this is regarded as an objectionable aspect of the hypothesis, then there are other reasons that speak against the ambiguity of the two sentences. Supposing the ambiguity lies in "is more handsome than", the ambiguity in question would have to be quite systematic and admit as many different meanings as there are relevantly different preference profiles of thinkers. But it is unattractive to treat a variation that is so systematic as a form of ambiguity. Moreover, if, as we will see, there are forms of embedding that exploit the parameter of variation here alluded to, then an ambiguity account would have difficulty in accounting for these embeddings.

presses a function from circumstances (time-world pairs) to extensions that is different from the one Barbara's use expresses. One way (but not the only way) in which all this might be true would be if the predicate had the same character as this predicate: "According to my standard, … is prettier than ---.". We could describe this by saying that the property expressed by Anna's use of "is more handsome than" is different from the property expressed by Barbara's use of the same predicate. An indexical proposal along these lines can save the possibility that neither Anna nor Barbara has violated any norms each is subject to: the content of Barbara's belief and assertion would on this view no longer be the negation of the content of Anna's belief and assertion.

But such a proposal faces a number of problems. Let us examine these. An immediate worry is that the indexical proposal distorts the content of the utterances, for it claims that Anna and Barbara talking and thinking *about* their own standards, even though it seems that they are merely comparing Depp and Pitt without their assertions and thoughts having any reflective content. By itself, this is a fairly vague worry. Moreover, I have been stressing that the only condition of adequacy for a semantics lies in the predictions as to literal correctness (speaker-centered assessments) it makes.[60] But there are a number of concrete phenomena of use that crystallize the issue.

The first problem concerns attitude ascriptions and indirect speech reports. If the sentence "Depp is prettier than Pitt." were indexical in such a way that it expresses different contents in (C1) and (C2), then we should expect speech-reports of such utterances to be sensitive to this difference. Speech reports are subject to the constraint that correct speech reports must adjust indexical elements in the utterance reported to any relevant changes between the context of the report and the context of the original utterance. For example, take our earlier example: if I utter the words "I am hungry now." at 12 noon, then a report by you, using the words: "MK said that I am hungry." Would not correctly report my assertion, because "I" in the report would refer to you. Similarly, if at 2pm I report my own 12 noon utterance with the words "I said that I am hungry now.", then the report would be incorrect, or at least odd: the word "now" in the report suggests that the reported utterance concerned the time of the report, though the presence of the present tense "am" is unusual, and the normal form of words would be "I said that I would be hungry now". A correct report would use the words: "I said at 12 noon that I was hungry then.". The following

60 In particular, I have not been assuming that the semantics needs to represent the structure of the thoughts of thinkers, as do, for example, Recanati 2007 and Carston 2002.

general rule therefore articulates an important constraint on correct speech reporting:

Constraint on reporting indexicals (CR)

(SR) If a sentence *s* is indexical in such a way that an utterance of *s* in context *c*1 expresses a different proposition from an utterance of *s* in a context *c*2, then an utterance by someone *A* of *s* in *c*1 cannot be correctly reported in *c*2 by using the form of words 'α said that *s*.' (where "α" is some term referring to *A*).[61]

First problem: the indexical proposal violates (CR)

Now consider the indexical relativist hypothesis that (5a) and (5b) do not express contradictory propositions because "Depp is more handsome than Pitt." contains an indexical which is sensitive to a change in context that occurs from C1 to C2 (for this is our proposed explanation of how both utterances are fault free). On this hypothesis, Barbara cannot, in the context C2 (or in a context relevantly like it) correctly report (5a) by saying "Anna said that Depp is more handsome than Pitt.". For the sentence "Depp is more handsome than Pitt." contains an indexical element which changes its content with the change from C1 to C2. However, this prediction is incorrect, for such a report by Barbara would clearly be correct. Thus the hypothesis is wrong.

It might be objected that expressions like "local" or "enemy" are indexical without meeting constraint (SR), and thus represent a counterexample. Consider two utterances (6a) and (6b), one reporting the other:

(6a) John (in el Raval, Barcelona): "I met her in a local bar."

(6b) Ben (later, in Digbeth, Birmingham): "John said he had met her in a local bar."

Alleged counterexample to (CR): "local"

The bar mentioned in (6a) is not in the surroundings of the place at which the report, (2b), is made, thus it seems that if (6b) is a correct report of (6a), then (SR) is false.[62]

Not a counterexample to (CR)

However, such cases are not conclusive counterexamples to (SR). The view that (6a) and (6b) constitute a counterexample relies on an underlying assumption about the exact way in which "local" is indexical. Let me explain. (6b) is a counterexample only if the indexicality of "local" is such that it refers to different areas in the contexts of (6a) and (6b). For example, if we assume that "local" always picks out the surroundings of the

61 Single quotes are here used like corner quotes.

62 This type of counterexample is based on an objection raised by François Recanati's to my paper at the relativism workshop in Oslo in 2005. Thanks also to Darragh Byrne for discussion.

place of utterance, then clearly "local" picks out disjoint areas in the two contexts. However, it is not clear that this is how the context-sensitivity of "local" works. A more plausible suggestion is that "local" picks out the surroundings of the place salient in the context of utterance. Thus, if prior to (6b), the area of el Raval in Barcelona has been raised to salience, then that's the area "local" will pick out in it. If not, then not, but then (6b) would not seem a correct report of (6a). Thus on this construal of "local", the case does not present a counterexample to (SR).

First problem not conclusive

Nevertheless, the objection to indexical relativism that is based on (SR) is not conclusive. For there might be a class of indexical expressions that do not conform to (SR), i.e. expressions behave like ordinary indexicals when they occur outside speech-report embeddings in that their content is determined by the context of utterance, but whose content within speech reports is determined not by the context in which the report is made but by some other context (possibly the context of the speech act being reported).[63] Claiming that sentences like the one used in (5a) are exceptions to (SR) would avoid the charge of being ad hoc only if there were further, independent cases. As we shall see below however, there is a surprisingly large range of cases that give rise to analogous problems. So a determined indexical relativist could use this to defend herself against the charge of making ad hoc exceptions.

A second problem …

A second problem for indexical relativism might be articulated in this way. The indexical relativist says that the propositions expressed by (5a) and (5b) respectively are not contradictory – they could both be true in the same circumstance of evaluation. As a consequence indexical relativism predicts that Barbara could come to believe what Anna has asserted (and vice versa) without changing her mind. However, in reality it is clear that Barbara cannot come to accept what Anna has said without changing her mind.[64] So indexical relativism is wrong.

63 Thus, in terms used by Kaplan (1977), speech report contexts are *selective* "monsters", i.e. they are monstrous with regard to some indexicals and not with regard to others. See Schlenker 2003 for extensive discussion of monsters in natural languages.

64 Sometimes a similar sounding point is made by saying "intuitively, Anna and Barbara disagree" or "intuitively, they genuinely disagree". Sometimes my own arguments in Kölbel 2002 and 2003 have been taken to involve this sort of premiss. However, there does not seem to be any sufficiently clear intuitive sense of disagreement on which such an argument against indexical relativism could rely, and at least my own arguments in previous work were not, I think, based on such a premiss. Disagreement might be understood in the sense that two people disagree only if they can't be both right, but that would contradict our hypothesis that Anna and Barbara are both faultless. Disagreement can also be understood in the sense that of two disagreeing parties neither can accept what the other is saying or believing without changing

This argument indirectly relies on the observations on speech-reports I have just been making. Of course, if what Anna has asserted or said is the semantic content of (5a), then indexical relativism does predict that Barbara can happily accept this without changing her mind. However, if the indexical relativist's line on speech reports in the object language is correct then it is probably also correct for us, the theorists. Thus when we are asking whether Barbara can accept what Anna has said, then we seem inadvertently to be considering whether Barbara could accept the proposition that Anna's words would express in Barbara's context. Thus if our earlier reply on behalf of the indexical relativist was ok, then this should explain why we have the impression that Barbara cannot come to believe what Anna has said: in this context we simply do not interpret "what Anna has said" as referring to the semantic content of Anna's utterance, but to the content her words would have in Barbara's mouth. I do not think this is a particularly pretty explanation, but it seems to work as well as the account of indirect speech reports.

… is also inconclusive

However, let me raise a third problem, which perhaps captures in more tractable terms the impression we have of some kind of disagreement between Anna and Barbara. To this end, let us consider in more concrete detail what the indexical relativist is proposing. So far we have merely defined indexical relativism as the view that the sentences used in (5a) and (5b) are indexical in such a way that they express different intensions in (C1) and (C2). But this could take many different forms. Let us consider three concrete proposals of what kind of character the sentence "Depp is more handsome than Pitt" might have:

Three sample implementations of indexical relativism: (IR1), (IR2), (IR3)

(IR1) "Depp is more handsome than Pitt" has the same character as "Depp is more handsome than Pitt according to my aesthetic pre ferences."

(IR2) "Depp is more handsome than Pitt" has the same character as "I prefer Depp's looks to Pitt's."

(IR3) "Depp is more handsome than Pitt" has the same character as "Depp is more handsome than Pitt according to the current aesthetic standards of the community to which you the audience and I the speaker belong."

---

their mind. In that sense the argument would be more promising, but it is not clear that this is what we mean by disagreement. For if I say at noon that MK is hungry, and you say at 2pm that MK is not hungry, then on Kaplan's temporalist view we are each saying things (temporal contents) that the other cannot accept without changing his or her mind. But we do not intuitively disagree (see MacFarlane 2007, which contains an excellent discussion of various notions of disagreement).

It is important to remember that there will be many further forms of indexical relativism, and I am simply considering these as examples. Let us pick out (IR1) as one reasonable proposal and consider two utterances that are also made by Anna and Barbara in the same respective contexts as (5a) and (5b):

(5a*) Anna (in C1): "Depp is more handsome than Pitt according to my aesthetic preferences."

(5b*) Barbara (in C2): "Depp is not more handsome than Pitt according to my aesthetic preferences."

According to (IR1), (5a*) and (5b*) should express exactly the same contents as (5a) and (5b) respectively, for they are utterances of sentences with the same character, made in the same contexts.

Three observations

Now let me make three observations:

A: Anna can coherently and without changing her mind respond to (5b*), but not to (5b) by saying: "I agree.", "I believe you.", "That's true." etc

B: Recall that (5b) is a response to (5a), and (5b*) is a response to (5a*). If Barbara prefixed her words in (5b*) with "No, …", the resulting utterance would be infelicitous. Not so in the case of (5b). In other words, Barbara couldn't felicitously respond to (5a*) by saying "No, Depp is not more handsome than Pitt according to my aesthetic preferences.". But she could felicitously respond to (5b) by saying "No, Depp is not more handsome than Pitt.".

C: At a later time, in context C48, Anna has changed her aesthetic preferences and now finds Pitt to be more handsome. She knows that her preferences have changed in this way. She can, in C48, refer back to her utterance (5a), but not to her utterance (5a*) and correctly say "I was wrong."[65]

Difficulty for (IR1)

These three observations go some way towards capturing the impression we have that Anna and Barbara somehow disagree in (5a) and (5b), but not in (5a*) and (5b*). It seems like the indexical relativist of (IR1) cannot easily explain these data, for she claims that the sentences used express the same contents in (5a) and (5b) as in (5a*) and (5b*). Why should (5a) and (5a*) differ in these ways if they are assertions of the very same content?

Difficulty for (IR2), (IR3) and indexical relativism in general

65 This corresponds to MacFarlane's "retraction data", see MacFarlane forthcoming.

(IR2) does not do any better. (IR3) may do better if Anna and Barbara belong to the same community in C1 and C2. But in that case our initial hypothesis, that they are both faultless, is wrong. The problem for the indexical relativist is that it is not easy to see how she could offer an alternative analysis of the indexical character of "Depp is more handsome than Pitt." that does not suffer from this problem.

The difficulty is independent of any notion of disagreement

So, there seems to be a sense in which an indexical relativist thesis like (IR1) or (IR2) does not capture the sense in which (5a) and (5b) involve an incompatibility or disagreement between the views of Anna and Barbara. We do not need a general analysis of disagreement to notice the differences described in observations A and B. These are clear differences between the original sentence and the sentence offered as a same-character paraphrase by the indexical relativist. They are simply used differently.

Semantic content does not exhaust meaning

I used to think that these differences constituted a clear disadvantage for the indexical relativist.[66] However, I now think that a resourceful indexical relativist, even a defender of (IR1), can explain the observed differences. To see how, we need to remind ourselves that there can be more to the meaning of an expression than its character. To claim that the sentence "Depp is more handsome than Pitt." has *the same character* as "Depp is more handsome than Pitt according to my aesthetic preferences." is not yet to claim that these sentences have exactly the same meaning. For while the sentences express the same content in every context of use, they might still differ in the presuppositions or generalized implicatures to which they give rise.

Indexical relativist explanation in terms of presupposition

Here is how that story goes. The sentence used in (5a) triggers the presupposition that speaker and audience have aesthetic preferences that do not relevantly diverge. The sentence used in (5a*) does not trigger such a presupposition. That explains why Barbara cannot coherently accept Anna's assertion without changing her mind, even though – according to the indexical relativist – she is perfectly happy to believe the content Anna asserted. The contents that are according to the indexical relativist asserted by Anna and Barbara in (5a) and (5b) are compatible. However, they are not compatible on the presupposition that Anna's and Barbara's standards converge in this matter, a presupposition that is carried by (5a) but not (5a*).

Explanation in terms of Gricean generalized implicatures

This difference between the sentences seems to be a permanent feature of this sentence type, as is shown by the fact that observations A and B

66 Kölbel 2002, 2003, 2004. The indexical relativist solution outlined below builds on that outlined in Kölbel 2007.

could be made in other contexts, and even with other sentences containing the same comparative adjective. This might be taken to suggest that they are Gricean conventional implicatures (I suggested this in Kölbel 2008b). But I am now inclined to think they are better explained as Gricean generalized conversational implicatures. For despite being regularly conveyed by the sentence used in (5a), they do arise from co-operative principles of conversation.[67] The sentence "Depp is more handsome than Pitt." might carry the presupposition that there is no relevant difference between the speaker's and the audience's aesthetic preferences simply *because* there is no explicit mention of a particular aesthetic standard. If the speaker had noticed such a relevant difference, then it would have been more co-operative to mention explicitly according to *which* aesthetic standard she is claiming one to be more handsome than the other. But since she is not specifying such a standard it is likely that there are no relevant differences between the speaker's and audience's aesthetic standards. The sentence used in (5a*) does not carry the same presupposition because in this sentence the relevant standard is explicitly mentioned, thereby suggesting, again in a Gricean way, that this *is* relevant.

Indexical relativism off the hook

This explains why Anna cannot respond by saying "I agree." or "That's true.". In order for that to be felicitous, she would first need to cancel the implicature or presupposition according to which her and Barbara's preferences do not diverge relevantly. It also explains observation B: prefixing a remark with "No, …" is appropriate only if there is an incompatibility between what has been said and what is being said after the "No, …". But the fully explicit form of words in (5a*) and (5b*) removes any suggestion that the contents asserted are incompatible, hence inserting "No, …" is inappropriate. Thus the indexical relativist can offer an independently plausible explanation of observations A and B.

Exercise 8*: What should the indexical relativist say about observation C?*

There is a fourth problem for indexical relativism, which is more momentous than the other three. However, I shall postpone discussion of it to §5.5.

5.4 Non-indexical Relativism about Taste

Let us go back to the motivation for examining the prospects for indexical relativism about taste. There were two motivations. One was that with sen-

67 Cf. Grice 1989, chs. 2 and 17.

tences like "Pickled herring is tasty." or "Depp is better looking than Pitt.", there seemed to be a discrepancy between speaker-centered and assessor-centered assessments of utterances for correctness. We noted that unless we ask explicitly for speaker-centered assessments, i.e. assessments of whether the speaker of an utterance has complied with the conceptual and linguistic norms to which he or she is subject, our data of user assessments will not be sufficiently uniform. In other words, with sentences of this sort, competent users will diverge amongst one another in their assessor-centered assessments. Since speaker-centered assessments seem to vary with context of use, we looked for ways in which the semantics could predict this. Indexical relativism and non-indexical relativism were possible answers.

Resumé

The other motivation came from a hypothetical situation in which two speakers utter sentences one of which is syntactically the negation of the other: (5a) and (5b). It seemed clearly possible that both speakers are faultless in so far as they assert what they believe and believe what they ought to given the conceptual norms they are subject to. But a semantics that treats sentences like "Depp is prettier than Pitt" as neither indexical, nor as circumstance-sensitive in non-standard ways predicts quite clearly that not both speakers can be faultless. So wishing to preserve the intuition that such scenarios are possible, we have to look for a semantics that construes sentences on matters of taste as somehow indexical or as involving a novel form of circumstance sensitivity.

Non-indexical relativism

In the last section we explored the first of these solutions, let's now turn to the second: non-indexical relativism. The non-indexical relativist about matters of taste claims that when we use sentences about matters of taste, we express contents that may vary in truth-value not only with a time and possible world, but also with a standard of taste. Thus, just as according to the temporalist the sentence "MK is hungry." expresses a content that changes its truth-value over time, the non-indexical relativist about taste claims that "hip-hop is better than swing." expresses a content that varies in truth-value with a standard of taste.

Non-indexical relativism provides a fairly direct and natural explanation of the pattern of speaker-centered and assessor-centered assessments of utterances on matters of taste, and why the two diverge so clearly. The contents we are dealing with are functions from world-time-standard of taste triples to truth-values. If we think of these contents as propositions i.e. the objects of propositional attitudes, they triply violate Frege's methodological principle that propositions (thoughts) have absolute truth-values. Let's think about one central propositional attitude in particular: belief. As outlined in chapter 2, believers are subject to norms. If a content of belief varies in its truth-value from world to world, then it is the world of

the believer (the actual world for that believer) which is normatively relevant: a believer ought to believe a content only if that content is true at the believer's world, i.e. the world she shares with all the co-inhabitants of her world. If a content of belief varies in truth-value with a time, then it is the time at which the belief occurs which is relevant: a thinker ought to believe a content only of the belief is true at the time at which she has the belief. If, now, we have contents the truth-value of which varies with standards of taste, then it is the believer's standard of taste that is relevant: a thinker ought to believe a content only if is true relative to that thinker's standard of belief. Thus, one ought to believe a content only if that content is true at one's world, at the time at which one has the belief and relative to one's own standard of taste:

(NB) A thinker *a* ought to believe a content *p* at time *t* only if *p* is true at *a*'s world, at *t* and relative to *a*'s standard of taste.

Non-indexical relativism elaborated

That there should be contents of this relativistic sort fits well with the joint ideas that contents are constituted by concepts and that some concepts, such as the concept of tastiness, are governed by norms of correct application that are sensitive to certain personal features, such as aesthetic preferences, of the thinker who employs the concept, as I was suggesting in chapter 2.

Some contents not to vary with some or all of the circumstantial parameters

This does not mean, of course, that all contents must vary with all three circumstantial parameters. There may be logically necessary or contradictory contents, temporally fixed contents and gustatorily fixed contents. For certain ranges of contents, some of the parameters may be inert.

When we think about the language use and the assessments by competent users that we hope our semantic theory will be able to predict, then the question arises what are the relevant norms of correctness that govern the correct use of sentences that have these multiply variable contents. Under what conditions is it correct to assert such a content? I.e. under what conditions is it correct to utter a declarative sentence in a context, and thereby to assert the content that the sentence has in that context? Now, there are many different accounts of assertion, and I do not here want to enter that debate. However, all parties may agree that there is a prima facie norm of sincerity which requires us to assert a content only if we believe that content. It may be controversial whether this norm is a basic norm of assertion or whether it is derived from some basic norms or properties of assertion together with further, perhaps prudential, moral or social norms. But that there is a norm of sincerity of this sort and has at

least *some* pro tanto[68] force seems to me to be uncontroversial. This gives us the following norm of correct assertion:

(NA) It is correct for a thinker to assert a content only if the thinker believes the content.

We have, since chapter 3 been operating with the assumption that uttering a declarative sentence amounts (under certain normal conditions) to an assertion of the content of the sentence in the context in which the sentence was uttered:

(A) Uttering a declarative sentence *s* in a context *c* amounts (under certain normal conditions) to an assertion of the content of *s* in *c*.

(NB), (NA) and (A) allow us to derive that a speaker violates neither the norm of belief nor the norm of sincerity only if the declarative sentences she utters express a content at the context of utterance that is true at the world and time of utterances and relative to the standard of taste of the speaker. This allows us to articulate at least a necessary condition for the correctness of an utterance in terms of the semantic value function SV:

(U) The utterance of a declarative sentence *s* in a context *c* is correct only if SV(*s*, *c*, *a*, <T(*c*), W(*c*), S(*c*)>) = 1.

Minimal account of normative role of contents

Here, "S(*c*)" refers to the standard of taste *of the context*, which we can take to be the standard of taste of the context's agent at the context's time.

Explanation of diverging assessor-centred assessments

(U) merely provides a necessary condition for the correctness of an utterance. Let us suppose, for a moment, that there are no further requirements on the correctness of an utterance. Then (U) provides a straightforward explanation of the speaker-centered and assessor-centered assessment data that we were having trouble to explain. When a competent user is making a speaker-centered assessment of an utterance, she considers, amongst other things, whether the utterance complied with (U). This explains why speaker-centered assessments of, for example, the sentence "Depp is more handsome than Pitt." will vary from one utterance of the sentence to another, when the respective speaker of the utterance has a standard of taste according to which Depp is or is not more handsome than Pitt. When competent users make assessor-centered assessments, they are assessing directly the content of the utterance, that is,

68 A pro tanto norm is a norm that may compete with other norms when an agent comes to an overall practical judgement as to what she ought to do. The idea is that the agent considers and weighs all the pro tanto norms and thus arrives at an overall normative assessment.

they consider whether they themselves would comply with (NB) if they believed that content, or whether they themselves would comply with (NA) and (U), if they uttered the same sentence.

Explanation of observation A, B and C

The same goes for our scenario involving Anna's and Barbara's contradictory yet faultless utterances in (5a) and (5b). Each is faultless because each complies with (U). However, what they are asserting (following (A)) are contradictory contents, i.e. what Anna asserts is true at a circumstance if and only if what Barbara asserts is not true at that circumstance. This explains why Anna could not coherently respond by saying "That's true." or "I agree." (observation A). For these utterances would be tantamount to her asserting what Barbara has asserted, and given that the two assertoric contents contradict one another, this response could only comply with (U), if her earlier assertion did not comply with (U), against our hypothesis. She would need to change her mind first in order not to conflict with (NA), and her standard of taste would need to change in order for her not to conflict with (NB). It also explains why Barbara could prefix her utterance with "No, …" (observation B): what she is asserting is incompatible with what Anna has asserted in the sense that no one person could coherently assert (or believe) both contents at the same time. Finally, there is an explanation why Anna can, at a later time at which she knows that her standard of taste has relevantly changed, say of her earlier assertion (5a) "I was wrong.". Again, saying "I was wrong." amounts to rejecting the content of (5a), much like an assessor-centered assessment.

More full-blooded accounts of assertion

We were working under the assumption that (U) is not just a necessary but also a sufficient condition for correct assertion. How problematic is this? What further requirements might there be on correct assertion? Let me at least consider two further conditions that have been dicussed. One might require in addition that asserters (as well as believers!) ought to have justification for what they assert (and what they believe). Such a requirement would be violated by any assertion that merely happens to be of a content that is true relative to the utterances world, time and standard of taste, but for which the asserter did not have any justification. With this further requirement on correct assertion, the explanations of both the assessment data and the faultlessness of Anna's and Barbara's utterances still work. Williamson's knowledge account requires furthermore that an asserter should know what she is asserting.[69] This adds a complication: we have to ask ourselves whether gustatorily variable propositions can be known. If we say "yes", then the scenario with Anna and Barbara can be strengthened in such a way that both Anna and Barbara know what they are asserting, and the argument as well as the explanation still go through. If we

69 Williamson 1996 and 2002.

say "no", then presumably we should also say that these contents cannot be genuinely asserted, for contents of this sort are not appropriate objects of knowledge, and only potential objects of knowledge can be properly asserted. But in that case, we should introduce a broader category of speech-acts, quasi-assertion, so that gustatorily invariant contents can at least be quasi-asserted. Quasi-assertion would then be governed by (NA) and possibly by a further justification requirement. Our explanations could then be rephrased in terms of quasi-assertion.

Usefulness of communicating variable contents

In any case, our account does not depend on providing an account of assertion. What people believe and assert is subject to a range of norms, some of them more narrowly tied to assertion than others. The usefulness of communication via assertion may be facilitated by a range of normative requirements on asserters. Certain accounts of assertion say that by asserting, an asserter undergoes certain justificatory obligations, and also issues a license to her audience to rely on what has been asserted. (e.g. Brandom 1983, Kölbel 2002, MacFarlane 2005). (NB) and (NA) are sure to play some role in any account. Surely, the ideal and most straightforward cases of assertoric communication involve asserters who believe correctly what they assert, and their audiences can profit from the assertion if they take the assertion to be sincere and correct by the asserter's lights. When the content asserted is gustatorily invariant, a content correctly believed by the asserter is also one that can be correctly believed by the audience. When contents are gustatorily variable, the audience has to be more careful. If the audience thinks that the asserter's standard of taste relevantly converges with her own, then again, correctness for the asserter coincides with correctness for the audience. But even in the absence of converging standards of taste, assertions of gustatorily variable contents can be of some use, either in acquiring information about the asserter or in acquiring information indirectly about the subject matter of the assertion.

Non-indexical relativism does not involve abandoning objective truth altogether

We all share the same world, so our beliefs, in so far as the truth of their contents depends on the world, are subject to the same standard. We are all subject to the norm that our beliefs should represent the actual world correctly. However, in so far as the truth-value of belief contents depends on the time at which we have the beliefs the correctness of our beliefs only depends on the same standard when we have those beliefs at the same time. And in so far as the truth-value of belief contents depends on a standard of taste, we are subject to the same standard only to the extent to which our standards of taste, i.e. our aesthetic preferences, converge. If we restrict our attention to the range of contents that are invariable as to standard of taste and time, we can still say, with Frege, that these contents have absolute truth-values: what it is correct to believe for one person at one time can be correctly believed by anyone at any time. Thus, the picture painted by the non-indexical relativist about taste and temporalist

does not force us to give up the idea that there are objective truths, truths that describe the world that we all share. It merely adds some further belief contents that allow us to make sense of our thought and talk more generally.

5.5 Operators

Temporal intensional operators

Temporalism is the view that some sentences, such as "MK is hungry.", have temporally variable contents. In §§4.1 and 4.4 I mentioned that one motivation for such a view is that it allows a treatment of certain expressions as intensional temporal operators which shift or fix the time parameter in the circumstances of evaluation. Thus, if the content of "MK is hungry." is temporally variable in truth-value, then we can treat "sometimes" in "Sometimes MK is hungry." as fixing the time-parameter in the circumstance of evaluation. Thus, for all sentences α, 'Sometimes α' is true at a context *c* and circumstance *p* just if α is true at *c* and some circumstance that varies from *p* only in the time component.

Are there any operators that might operate on the standard of taste parameter in the circumstance of evaluation, as postulated by the non-indexical relativist about taste? There certainly could be such operators. Let us imagine a language L5 which results from L4 by adding simply a standard of taste parameter to the circumstances of evaluation. Make all the adjustments needed for that change, in particular the re-definition of the set P of circumstances of evaluation. Now let's imagine there are also predicates whose extension varies with the standard of taste parameter. For example the 1-place predicate **is pretty** with the following semantic clause:

5b. **is pretty**: For all $c \in$ C, for all $a \in$ A, for all $p \in$ P:
SV(**is pretty**, *c, a, p*) = {*x*: *x* $\in$ O and *x* is pretty at T(*p*) at W(*p*) relative to S(*p*)}

"S(*p*)" here refers to the standard of taste of the circumstance of evaluation *p*. We could now add an operator **by all standards** which fixes the standard of taste parameter in exactly the way in which **always** fixes the time parameter and **necessarily** fixes the world parameter. This might be the semantic clause:

15k. **by all st** … if α is a sentence, then SV('**by all standards** α', *c, a, p*) = 1 iff for all $p^* \in$ P such that T(p^*) = T(*p*) and W(p^*) = W(*p*), SV(α, *c*, *a*, p^*) = 1.

Sample clauses for taste predicates and intensional taste operators

Similarly, we could add operators of the form **for Sue**, **for John**, **for** x_1 etc. These would be complex operators formed by **for** and a name. So the expression **for** would be in the category of expressions that form sentence operators from names. We could then add the following semantic clause:

15l. **for** … if α is a name and β is a sentence, then SV('**for** α β', *c, a, p*) = 1 iff for all $p^* \in$ P such that S(p^*) is compatible with the aesthetic preferences of SV(α, *c, a, p*), T(p^*) = T(p) and W(p^*) = W(p), SV(β, *c*, *a*, p^*) = 1.

Thus, a complex sentential operator '**for** α' fixes the standard of taste parameter to the referent of the name α. I believe that **for** and **by all standards** are prima facie good models for the English "for" and "by all standards", as they might occur in sentences like "For Peter, Sue is pretty." or "By all standards, Sue is pretty.".

Reason why these operators should not be construed as sentential

It has been argued by Cappelen and Hawthorne (2009, p. 75) that a sentential operator treatment of "for Peter" is problematic because sometimes this qualifying phrase qualifies merely a part of a sentence, as for example in "She ate something tasty for Peter on a pretty plate.". Here, the suggested sentential operator would not be sufficiently selective: "For Peter, she ate something tasty on a pretty plate.". This seems to me to be a very good point, a point that may equally apply to some other operators that have traditionally been treated as sentential, such as "necessarily". This means that we should adapt the two qualifiers and classify them as being of a different syntactic category, for example the category of a predicate modifier or an adjective modifier.

This, however, leaves the intensional nature of the qualifiers intact: a predicate modifier can also fix the standard of taste parameter in the circumstances of evaluation. It will turn a gustatorily variable predicate into one that is gustatorily fixed. In other words, applying "for Peter" to a predicate will turn any predicate whose extension varies with the standard of taste parameter into one that no longer varies along that parameter. Let us briefly look at how such a treatment of such predicate modifiers of the form '**for** α' might look:

Revised clause for **for** as predicate modifier

15l.* **for** … if α is a name and β is a 1-place predicate, then SV('β **for** α', *c, a, p*) = {*x*: for all $p^* \in$ P such that S(p^*) is compatible with the aesthetic preferences of SV(α, *c, a, p*), T(p^*) = T(p) and W(p^*) = W(p), $x \in$ SV(β, *c*, *a*, p^*)}

The possibility of this kind of treatment of this type of natural language embedding is, I believe, an added bonus of non-indexical relativism about taste. It is worth noting that **for**, on both accounts, combines with a *name*

to form sentence operators/predicate modifiers. This means that it can also combine with *variable* names and thus give rise to binding phenomena.[70]

Exercise 9: *articulate a semantic clause for a predicate modifier* ***by all standards****.*

Binding phenomena

[At this point it will be worth commenting briefly about certain binding phenomena that allegedly require a free-variable approach. Consider this sentence:

(7) Every man kisses some pretty woman.

It has been observed that two readings of this sentence are distinguishable. On one reading an utterance would be correct iff each man kissed a woman who is pretty according to the standard of taste of the circumstance relative to which the whole utterance is being evaluated. On another reading, each man needs to kiss a woman who is pretty according to that man's standard. This second reading seems to involve binding, and where there is binding, there must be variables that are bound, and if an expression involves an implicit variable once, then it does so always (Stanley 2000, 2005, but see also Recanati 2004, Lasersohn 2008). Our non-indexical relativist says that "pretty" is just a one-place predicate, and that it does not involve a second argument place for a standard of taste. So how can the non-indexical relativist account for the second reading? Is binding involved?

In my view, the two readings of (7) corresponds to the following kinds of base language sentences (in a language similar to L4 – I am standardizing the syntax somewhat):

(7a) **Every man** (λx_1 (**some pretty woman** (λx_2 (x_1 **admires** x_2)

(7b) **Every man** (λx_1 (**some pretty for** x_1 **woman** (λx_2 (x_1 **admires** x_2)

Response to Stanley's Binding Argument

I am assuming that we have also added an attributive version of **pretty** and that "**for**" in (7b) is just the adjective modifier described in 151.* above. (7a) will express contents that standard-variable, while (7b)'s contents have truth-values that are constant with respect to the standard parameter in circumstances of evaluation. With an operator like **for**, given

70 The possibility of variables in this position is in effect what Recanati calls a "variadic function" (2007) in the context of a discussion of location operators. The fact that a variable can be attached to qualifier-forming expressions like **for** also provides the resources for dealing with a range of binding phenomena to which Stanley (2005) has drawn attention in the same debate.

that it can combine with variables, we can generate any required bound reading without having to treat "pretty" as containing an implicit free variable.]

What does the *indexical* relativist about taste have to say about qualifiers like "by all standards" or "for Peter"? For the indexical relativist, a sentence like "Sue is pretty." will have a non-constant character, i.e. it will assign different intensions to it in different contexts of use, depending on the standard of taste of the speaker of the context. This will, no doubt, be due to the indexicality of "is pretty", which might receive the following semantic clause in a language like L4:

5b.^{IR} **is pretty**: For all $c \in$ C, for all $a \in$ A, for all $p \in$ P:
SV(**is pretty**, *c, a, p*) = {x: $x \in$ O and at T(p) in W(p)
x is pretty according to the standard of taste that A(c)
has at T(c)}

Sample indexical relativist clause for a taste predicate

Here, "A(c)" and "T(c)" refer respectively to the agent of context c and the time of context c. This means that the intension of "Sue is pretty" in any context of use will at most vary along the time and world parameter. Thus there is no relevant unspecificity or variability on which an intensional operator "by all standards" or "for Peter" could operate.

Nevertheless, we want these qualifiers to shift or fix a standard of taste parameter, so that "For Peter, Sue is pretty" no longer has a non-constant character, i.e. that this sentence no longer depends in its intension on the standard of taste of the speaker of the context in which it is uttered. This means that the qualifiers must be operators that shift a contextual parameter, i.e. they must be construed as "monsters" in Kaplan's sense. We already saw that the indexical relativist is committed to a monstruous treatment of speech reports, so this is nothing completely new and shocking. A semantic clause for a sentential operator version of **for** might then look like this:

Sample indexical relativist clause for **for**

15l.^{IR} **for** … if α is a name and β is a sentence, then
SV('**for** α β', *c, a, p*) = 1 iff for all $c^* \in$ C such that the standard of taste that A(c^*) has at T(c^*) is compatible with the aesthetic preferences of SV(α, *c, a, p*) at T(c), and otherwise c^* is like c, S(c^*), SV(β, c^*, a, p) = 1.

Content not compositional

This means that the language is no longer compositional in the sense that the intension of an expression at a context is determined by the intensions of its parts at the same context. For now the intension of a sentence at a context sometimes depends on the intension of its parts at *other* contexts. This may not be a problem, as long as the *character* of all expressions is

still compositional, i.e. that the character of an expression solely depends on the characters of its parts.

A potential problem

There is one potential problem: contexts of use, in Kaplan's framework, are highly restricted in the sense that the speaker of a context is always at the place of the context at the times of the context in the world of the context. So, contextual parameters cannot be freely shifted individually. This means that there may not always be an appropriate context in C which can provide the intension of the complex expression. But the set of contexts C was restricted for good reason: Kaplan wanted to be able to define logical truth as truth at all contexts, so that "I am here now." and "I exist." come out as logically true sentences (though the intensions they express are contingent). I believe that this is a problem the indexical relativist would need to solve, by either offering a different, non-monstruous treatment of the qualifiers in question, or by providing an alternative way of characterising logical truth.

Alternatives to indexical relativism

This sort of problem is already familiar from our discussion of anti-temporalist approaches in §4.4. There we explored two approaches that did not treat sentences like "MK is hungry" as strictly indexical: the predicate view and the free variable view. Analogous approaches are available in the realm of sentences concerning matters of taste. We could either view "pretty" and similar predicates as containing an extra argument place or as containing a free variable. Both of these approaches are alternatives to the non-indexical relativist approach. Thus we are not forced to adopt sentence contents that vary in truth-value with standards of taste. Let us briefly look at both.

The predicate view

On a predicate view of sentences on matters of taste, an apparent sentence like "Sue is pretty" is in reality a one-place predicate. So we can easily account for qualifiers like "for Peter" or "by all standards". These will just be construed as names or quantifier phrases which naturally combine with 1-place predicates. However, an account has to be given of what it is that we are doing when we utter predicates like "Sue is pretty." on their own. One option is to say that the context "provides" an argument for the predicate intension of the sentence. It is the resulting content (proposition) that is asserted by the utterance of the predicate. This comes pretty close to the indexical relativist's intitial intention. The only difference is that on the predicate view, the object of assertion, i.e. the contextually completed proposition) is not identical to the compositional content of the expression used. Thus we are departing from our idea that when we use an apparent

sentence to make an assertion, it is the sentence's intension at the context of use that we are asserting.[71]

On a free variable view of sentences on matters of taste, a sentence like "Sue is pretty." involves a 2-place predicate "is pretty" which is combined with a name "Sue" and with a an implicit variable. Thus, the sentence, an open sentence, will vary in truth-value with the assignment parameter. The obvious treatment of our qualifiers is then clear: they must be the sort of thing that can bind a free variable, such as a name and a binder or a quantifier phrase and a binder. (Such an approach would suggest that we could also treat temporal variability and world variability in this way, thus doing away altogether with the circumstance of evaluation.) As I already remarked in §4.4, this is an approach that has many similarities with the non-indexical relativist approach. Here, the objects of assertion will coincide with the intensions of the sentence uttered. However, the objects of assertion vary in truth-value with assignments of values to variables. Similarly for the objects of belief. We will have to articulate norms of correct assertion and belief that are similar to the norms discussed above in §5.4.

The free-variable view

I conclude that indexical relativism about taste faces some problems with the compositional account of qualifiers like "for Peter" or "by all standards". An indexical relativist treatment of these qualifiers meant giving up the compositionality of *contents* (by treating the qualifiers in effect as Kaplanian "monsters"). More problematically, it meant having to extend the domain of contexts to include contexts <*a*, *l*, *t*, *w*> where agent *a* is not located at location *l* at time *t* at world *w*. This in turn means that the Kaplanian account of logical necessity would need to be replaced. However, the predicate view and the free-variable view are less problematic alternatives. According to the predicate view, apparent sentences concerning matters of taste are not in fact *sentences* but *predicates*, and have the corresponding semantic contents (functions from world-time pairs to extensions). This view might still maintain that the *assertoric* content (not the semantic content) is a proposition that is the result of applying the intension of the apparent sentence to a contextually provided object. The free-variable view similarly has to explain which objects are the objects of assertion, given that an operator-free sentence like "Peter is pretty." would express a sentence-intension only relative to an assignment of values to variables. Again, some account of how the context determines the relevant assignment would yield a propositional object of assertion. Thus both the predicate view and the free-variable view have in common that they distin-

71 This distinction between content of assertion and semantic content is similar to Dummett's distinction between ingredient sense and assertoric content (Dummett 1991).

guish between the compositional semantic content and the content of assertion (much like Dummett's distinction between assertoric sense and ingredient sense, see Dummett 1991, p. 48 and Stanley 1997, pp. 574–8). Non-indexical relativism can treat the qualifiers unproblematically as intensional operators but has to give an account of what it is to assert contents with additional parameters of truth-value variation. Thus, there were three viable accounts of the qualifiers, but a strictly indexical relativist account involved a major rehaul of the framework.

5.6 Indexical and Non-indexical Relativism in Other Cases

Similar cases

The problem of accounting for the possibility of superficially contradictory utterances that are nevertheless fault free is not unique to discourse about matters of taste, even though perhaps this is the case that is most intuitive. The same goes for the problem of sentences on which the speaker-centered and assessor-centered assessments by competent users might diverge. The same problems and also the various possible types of response, recur with remarkable similarity in a wide range of further cases, some of which I shall now briefly review. I shall only expound the initial motivation and remark on some specific differences. It will be left to the reader to restore the detail provided in the taste case to the other cases.

Epistemic Modals:

Imagine Anna is trying to find out who emptied her bottle of whisky. At the beginning of her investigations, at time t1, she thinks that it might have been Barbara, because at that time she has no evidence that rules out the possibility of Barbara having emptied the bottle. Much later, at t2, she finds out that Barbara could not have emptied the bottle because she was far away at the relevant time. Anna makes two utterances:

(7a) Anna at t1: "It might have been Barbara."

(7b) Anna at t2: "It's not the case that it might have been Barbara."

Faultless disagreement scenario for epistemic modals

Let us assume that the contexts of (7a) and (7b) are such that in both cases "it" picks out the same emptying of the same whisky bottle. Thus, on the face of it, the two utterances look like the assertion and denial, respectively, of the same epistemic possibility. The utterances *seem* to express contradictory contents: the proposition that it might be the case that Barbara emptied the whisky bottle and the proposition that it might not be the case that Barbara emptied the whisky bottle. Nevertheless, there is a clear sense in which Anna is right both times. Given that her knowledge at t1

does not rule out the possibility in question[72], what she says in (7a) is the right thing to say at t1 (and its content is the right thing to believe at t1). Given that her state of knowledge at t2 does rule out the possibility in question, the sentence she uses in (7b) is the right thing to say (and its content the right thing to think) at t2. Both utterances comply perfectly with the norms of assertion and of belief.

As in the taste case, there are several ways in which one can try to accommodate the impression that both utterances are correct. First, one could make a diagnosis of hidden indexicality: the propositions expressed by (7a) and (7b) respectively are not contradictories, for the sentences used exhibit a form of indexicality that causes some change of intension from t1 to t2. For example, the first sentence used might have a similar character to:

(7a*) Anna at t1: "My current state of knowledge does not rule out that Barbara did it."

Indexical relativism

As in the taste case, the indexical hypothesis generates various seemingly problematic predictions. One prediction concerns speech reports. We would expect, given (SR), that Anna cannot, at t2, correctly report her utterance (7a) homophonically. For (7a*) clearly could not be reported homophonically at t2. However, contrary to this prediction (generated from (SR) and the indexical hypothesis), a homophonic report at t2 is perfectly correct: "I said that it might have been Barbara.". Explaining this would require dropping (SR) and saying that the hidden indexicals differ from ordinary indexicals in that they are subject to certain "monstruous" context-shifting operators. A similar complaint is that Anna might add at t2: "But I was wrong. For it could not have been Barbara.".[73] If the indexical hypothesis were correct, one would not expect this. One would expect her to continue to endorse the proposition asserted at t1: "What I said was true: my knowledge at t1 did not rule out Barbara being the perpetrator.". But it seems that this is not how we would retrospectively assess such an utterance.[74]

The problem is that the way we report utterances like (7a) and (7b) and the way we evaluate them sometimes suggests that they have contradictory contents, and the indexical relativist cannot easily accommodate that,

72 And perhaps given there is also no easy way in which Anna could have obtained such knowledge. See DeRose 1991.

73 See MacFarlane forthcoming a, for a version of this argument against an indexical view.

74 However, compare the objections raised against this reasoning in Wright 2008.

except by invoking the generalized implicatures I mentioned in §5.3. It seems that after the new information has come in Anna re-assesses the *same* content from a different perspective, namely that of the knowledge that she now has. A non-indexical relativist construal is a natural alternative: the content in question varies in truth value with an extra parameter in the circumstances of evaluation. The content that it might have been Barbara is true relative to Anna's earlier state of knowledge, and it is false relative to her later state of knowledge.[75]

Non-indexical relativism

Which state of knowledge, then, is relevant for assessing the correctness of a belief with such a content? The most plausible account, it seems, would say that an assertion of the content that it might be that *p* is only correct if *the speaker's* knowledge at the time of assertion is compatible with *p*.[76]

Operators

Possible intensional qualifiers that pick up the variability with respect to state of knowledge might include "given what Peter knows …". My earlier observations about possible treatments of such operators by the indexical and non-indexical relativist, as well as the predicate- and free-variable views will apply *mutatis mutandis* to these cases.[77]

Probability:

On a subjectivist construal of probability, probability ascriptions are plausible candidates for a relativist treatment similar to that of epistemic modals.[78] Whether it is right for some thinker to call some outcome probable will depend on the evidence available to that thinker or even on the probability function with which that thinker started out.

Knowledge Attributions:

One of the areas of most active debate in recent epistemology has been a position called "contextualism about knowledge".[79] According to the most

75 Relativism about epistemic modals is defended by Egan, Hawthorne and Weatherson 2005 and MacFarlane forthcoming a. See also Egan 2011 and Gillies & von Fintel 2006.

76 Some uses, however, would suggest that it is not always the speaker's knowledge that is relevant. For example, consider the game show host who says: "The prize might be behind this door." while knowing it is not. Compare Gillies & von Fintel 2006, Bach 2006 and Egan 2008.

77 However, see Yalcin 2007, who discusses a problem for a semantics of epistemic modals that does not seem to be easily resolved by any of the views here discussed.

78 See Price 1983 and Kölbel 2002 for some discussion.

79 Cohen 1986, DeRose 1992, Lewis 1996 are classic articulations. For a recent overview of the vast literature on this topic see Rysiew 2007.

common form of this view, sentences used for knowledge attributions, such as "Charles knows he has hands.", are context-sensitive in unobvious ways. The sentence is obviously indexical because in different contexts of use, it will concern different times in Charles' life. The contextualist's claim is that the sentence "Charles knows he has hands." is in addition sensitive to the standard of knowledge associated with the context of use. Thus, in a context C1 (perhaps when waking up in hospital after a horrific accident), the standards required for knowledge are relatively low: Charles' feeling his right hand with his left and vice versa may be enough for knowledge that he has hands. But in the context C2 of discussing scepticism, when a sceptic has just raised a sceptical possibility – perhaps that of Charles being a handless brain in a vat tricked by a computer into thinking he has got feelings in his hands – the standard sufficient for knowledge is much higher. Thus the following two of Anna's utterances may both be correct:

(8a) Anna at C1: Charles knows (at noon 1/1/2010) that he has hands.

(8b) Anna at C2: Charles does not know (at noon 1/1/2010) that he has hands.[80]

Faultless disagreement scenario for knowledge ascriptions

Being able to say that knowledge attributions vary in this way with the contextual standards for knowledge enables contextualists to give an answer of sorts to sceptical arguments. It allows them to concede sceptical conclusions such as that expressed by (8b) when faced with a sceptical argument, without having to give up all or even most claims to knowledge. For (8a) and (8b) are both correct.

Indexical relativism

This result is achieved, in ordinary forms of contextualism about knowledge, through the type of indexical hypothesis with which we are by now familiar. The sentence

(8) Charles knows (at noon 1/1/2010) that he has hands.

is claimed to have the same character as the sentence

(8*) By the currently salient standards for knowledge, Charles knows (at noon 1/1/2010) that he has hands.

80 I added the time specification in order to highlight that the contextualists are not talking about ordinary time-sensitivity of knowledge attributions. In other words, they are not claiming that Charles "forgets" that he has hands from C1 to C2 (or that he "learns" that he has hands, in case C2 precedes C1).

Thus, the two utterances (8a) and (8b) do not express contradictory contents. Rather, (8a) expresses the proposition that Charles meets C1's standard for knowing (at noon 1/1/2010) that he has hands, while (8b) expresses the proposition that Charles fails to meet C2's standards for knowing the very same thing. Postulating a hidden form of indexicality enables the contextualist to claim that both (8a) and (8b) are true.

Problems for indexical relativism

The parallelism with the previous cases is not difficult to detect. Similarly, it is not difficult to see that the indexical hypothesis will be subject to some of the same problems that arose for indexical relativism in previous cases. Thus, for example, it would seem to be correct to report (8a) in C2 by saying "Anna said that Charles knows that he has hands.". Indeed, if one pointed out to Anna after her utterance (8b) that she said earlier that Charles does know that he has hands, then it would be quite normal for her to say: "Yes, I did say that. But I was wrong.". These observations seem to be in conflict with epistemic contextualism in the form currently discussed,[81] so the epistemic contextualist will need to adopt some of the more sophisticated strategies I adopted on behalf of the indexical relativist above in section 5.3.

Non-indexical relativism

Alternatively, she can go for a form of non-indexical relativism.[82] Instead of saying that sentences like (8) express different contents in different contexts of use, due to variations in the relevant standards for knowledge, we could say that the content expressed remains constant but that this invariably expressed content is unusual in that it varies in truth-value with an extra circumstantial parameter. Rather than complicating the function from context of use to content (i.e. the character), we instead complicate the function from circumstances of evaluation to extensions by adding a circumstantial parameter.[83]

Again, relativists about knowledge attributions will have to say something about how the changes in the semantics affect the requirements on correct assertion and belief. I believe that the mainstream of contextualism is a form of *attributor* contextualism, thus the most obvious starting point would be to say that correct assertion and belief of knowledge attributions

81 See Yourgrau 1983 and Kompa 2002 for similar observations about contextualism.

82 See MacFarlane 2005b.

83 As pointed out earlier, MacFarlane (2005b) prefers to introduce in addition to context of use and circumstance of evaluation, a "context of assessment". The effect is similar.

requires that the content asserted be true relative to the standard of knowledge determined by the context of the assertion or belief.[84]

Moral Values:

The case of moral value is somewhat similar to that of aesthetic value already discussed at length. However, in the moral case, the driving assumption that there is some extra factor on which value depends is much more controversial. To give a very brief sketch, let us consider two utterances, one made in an Indian moral context by Arvind, one made in a Western European moral context by Barbara

(9a) Arvind in C1: One ought not to marry outside one's own caste.

(9b) Barbara in C2: It's not the case that one ought not to marry outside one's own caste.

Now, to say that both of these utterances may be correct is highly controversial, by far more controversial than in the taste case. Nevertheless, some philosophers and some anthropologists have been driven by a variety of different considerations to the conclusion that moral judgements depend for their correctness on some implicit parameter such as a moral code. Usually positions of this sort are called "moral relativism". The best known position that goes under that name, namely Harman's moral relativism[85], is actually a thesis of indexical relativism. Again, (9a) and (9b) do not express contradictory propositions and this explains the (controversial) starting datum that neither (9a) nor (9b) were mistaken. Much the same problems have been pointed out with this view, and a non-indexical relativist version has been suggested.[86]

Fiction:

Another potential case of application is that of fictional utterances. The very same sentence, for example "Mozart's requiem was commissioned by Salieri." can be used to make a correct remark about Forman's movie *Amadeus*, and an incorrect remark about the real world. The standard treatment involves the postulation of contextually implicit operators.[87] However, implicit operators can be avoided if one takes seriously the idea

---

84 However, MacFarlane (2005b) argues that an adequate pragmatics must allow us to evaluate the same utterance by recourse to varying standards.

85 Harman 1975. See also Dreier 1990 and 1999.

86 See Gibbard 1990, Kölbel 2002, 2004 and 2007.

87 For example, Field 1973 and Lewis 1978.

that the truth of a proposition can vary with an extra parameter.[88] The extra parameter is presumably something like a possible world, so that the world of Forman's *Amadeus* and the actual world can figure as values of such a parameter.[89] If such a parameter is allowed, we can be much more differentiated in the pragmatic evaluation of utterances as true, and thus do justice to the impression that the very same proposition – that Salieri commissioned Mozart's requiem – can be asserted correctly in one context and incorrectly in another.

Painted Leaves:

The debate between traditional semanticists and "contextualists"

There is a seemingly unrelated debate in theoretical linguistics and the philosophy of language about the general role of context in linguistic communication. Contextualists, in this debate, challenge a certain traditional model of linguistic communication associated with formal semantics. The formal semanticists believe that the linguistic meaning of a natural language sentence will determine for each context of use a proposition that it semantically expresses at that context. The route from context to proposition expressed is, according to the traditionalists, a purely mechanical one completely anticipated by the semantics of the sentence used. The paradigm for this role of context is given by typical indexical expressions such as "I" or "here" or "now". However, according to contextualists, the route from context to content is not one that can be anticipated in this way in a semantic theory. Contextualists put much more emphasis on pragmatic processes of determining which proposition is expressed by any particular utterance. They often claim that a sentence by itself together with the context of use only determines an incomplete or truncated proposition, one that needs completion via an interpretive process of "free enrichment". Now, the cases presented by contextualists in support of their view bear an uncanny resemblance to the sort of cases we have been discussing.

The green leaves scenario

One of the best-known examples is from Charles Travis (1997). Suppose Pia utters the sentence "The leaves are green." twice, once talking to a photographer who is looking for a green motif, and once talking to a botanist looking to classify the plant in question. The twist is that the leaves that each of the utterances refers to are the leaves of a Japanese maple. As is normal with Japanese maples, the leaves were originally russet in colour. But subsequently they have been painted green. According to

88 The basic idea for this is contained in Predelli 1997, 2005. However, Predelli seems to have changed his mind – see his 2006.

89 However, the job of the new parameter cannot be done by the standard possible world parameter. For there are possible worlds in which Salieri does commission the requiem and Forman's fiction portrays him as not commissioning it.

Travis (1997), the first utterance is true, while the second utterance is false. The lesson is supposed to be that a semantic rule alone cannot anticipate which propositions might be expressed by a sentence in which contexts, and more examples on the same model can be easily constructed.

Travis' answer to the pragmatic solution

The typical defensive move of a traditional semanticist would be to insist that even though the utterances may *convey* or *communicate* different propositions which then differ in truth-value, the proposition semantically expressed by the two utterances are the same. Thus, the formal semanticist might say that both utterances are true, but the second is misleading in that it also pragmatically conveys a false proposition, namely that the natural colour of the leaves is green. Or she might say that they are both false, but that the first one also pragmatically communicates a true proposition, namely that the leaves are superficially green. The problem Travis points to is that the formal semanticist cannot offer a reason to prefer one of these two options to the other. Now, if this problem were restricted to only one case, it might have been feasible to provide some rationale for preferring one of the options. However, given that many more examples can be constructed, this will not remove the problem.

Now, in order to preserve the continuity with the earlier discussion, let me adjust Travis' example. Consider two utterances of superficially contradictory sentences:

(10a) Pia in C1: The leaves are green. [directed to a photographer who is looking for something green as background for his photo]

(10b) Pia in C2: The leaves are not green. [directed to a botanist who is trying to classify the tree]

New type of problem with indexical solution

There is a strong intuition that both utterances are correct. How can we accommodate that intuition in a traditional semantic framework? As before, we have an indexical and a non-indexical relativist option, as well as the predicate- and free-variable views. According to indexical relativism, the predicate "is green" is indexical in such a way that in C1, it expresses a property that applies to the painted leaves, and in C2 it expresses a property that does not apply to the leaves. Thus the propositions expressed in (10a) and (10b) do not contradict one another. By contrast to the previous cases, it is not easy to generate intuitions about speech reports or retrospective evaluations that would create a problem for indexical relativism. However, there are other reasons why it seems ill-advised to postulate indexicality in this case. If the current case provides evidence for "is green" being indexical in the way suggested, then it is likely that we will

be able to find evidence for the indexicality of any predicate whatsoever. For Travis-style examples abound.[90]

Non-indexical relativism

The conclusion contextualists want to draw is that a formal semantic theory for a language just cannot spell out in advance the semantic contents that will be expressed by the language's sentences in all possible contexts. A form of non-indexical relativism may help avoid this conclusion: instead of concluding that the contents expressed by sentences vary in unpredictable ways from context to context, we can say that the contents predictably expressed by sentences in contexts are evaluated, and used in communication, in unpredictable ways. The content expressed by the sentence "The leaves are green." remains invariant as long as we are talking about the same leaves at the same time. However, whether that content counts as true, whether asserting it counts as correct, depends on the specific purpose against which we are evaluating that content. We cannot in advance predict for what purposes people may venture to assert the content that those leaves are green at that time.[91] Thus again, we need to add a new parameter to the circumstances of evaluation: a purpose for believing or asserting something.

As before, once the circumstances of evaluation are enriched, the relativist also needs to say something about the norms of correct assertion and belief. The impression that (10a) and (10b) are fault free will need to be explained by reference to some such account. Perhaps it is correct to believe or assert a content only if it is true relative to the purposes one is pursuing at the time of having the belief or making the assertion.

Again, we might consider a predicate view and a free-variable view as a solution to the Travis problem: If we regard form of words used and (10a) and (10b) as either a predicate or as a sentence containing an unbound variable, we can explain why both utterances are correct by claiming that the propositions asserted by them do not contradict one another.

Advantage of non-indexical relaivism

Nevertheless, there is something unsettling about all these answers to Travis' problem. They face a problem analogous to the indexical relativist's problem. If, as Travis alleges, cases like the green-leaves scenario can be multiplied ad libitum, then the non-indexical relativist, would seem to be forced to postulate a multitude of new circumstantial parameters. Similarly,

90 This is the slippery slope of which Cappelen & Lepore warn us in their 2004.

91 Stefano Predelli's (2005) answer to Travis' problem seems to follow roughly these lines. MacFarlane has also offered a similar answer to context-shifting arguments in his forthcoming b. In MacFarlane's terminology, an invariant proposition can have a varying intension, rather than the proposition *being* the intension and varying in truth-value with respect to a purpose.

the predicate view would have to contemplate mutiply incomplete sentences, and the free-variable view would have to multiply implicit unbound variables. Perhaps one advantage of the non-indexical relativist is that this semantic indeterminacy complicates merely the semantic clauses for the sentences in question. The other three solutions are also committed to saying that the the syntax of the base language sentences (logical form), on which the semantic theory operates is highly indeterminate. Thus, while the latter cannot even start to articulate semantic clauses (because they can't be sure about the syntactic objects to which they are assigning semantic values in context), the former (nonindexical relativist) is merely handicapped by having to say something highly indeterminate about quite definite syntactic objects.

Travis cases are genuinely difficult

But the problem is urgent for all four. After all, semantics aims centrally to explain how converging linguistic competences can be acquired by the members of a speech community. If the conditions under which a given concept is correctly applied (non-indexical relativism) vary unsystematically, then it becomes a mystery how language users can be trained up to acquire converging competences. If the deep syntax (logical form) of natural language sentences varies unsystematically across different uses, then again, how do language users acquire the ability to recognize these logical forms and to interpret them (predicate view and free-variable view)? Or alternatively, if all the syntactic complexity is claimed to be present on all uses[92], then how are language users going to acquire competence with all these forms of variability, given that many of them will rarely or never manifest themselves in their linguistic environment?

Possible antidote to Travis cases

All four approaches could be rescued from this problem if one of Travis's premisses could be plausibly denied, e.g. the premiss that green-leaf scenarios can be multiplied ad libitum, or the premiss that each new case requires a new type of indexicality/circumstantial parameter/incompleteness/free variable. Indeed, there is reason for hope. One of Travis' other problem cases: "Smith weighs 80 kg" (Travis 1985) may be seen as a symptom of a two recurring problems: vagueness and an ambiguity is the meaning of tense. It is unclear whether Smith needs to weigh exactly 80 kg, and whether he needs to do so on average, or in the morning before breakfast, with his clothes on, etc. If there are a general answers to problems concerning vagueness and the exact interpretation of tense, then several such cases could be dealt with in the same way, thus the systematicity of semantics could be upheld. Similarly with another one of Travis's examples: "The ball is round.": is the squash ball round when it hits the

92 This seems to be the view underlying Stanley's binding argument, see Stanley 2000, 2007.

wall and, at least for a short time, has an ovoid shape? There may simply be an ambiguity between "being round" in the sense of being round under certain normal conditions (a condition that persists throughout most if not all the ball's life-span), and the sense of "being round" as a description of the shape of the ball's boundaries at a particular moment in the ball's life. Moreover, it is not clear that all problem cases cited by contextualists fall victim to Travis's objection against a pragmatic solution.

6. Evans' Problem and Radical Relativism

6.1 Introduction

Recap chapters 2-4

Let us recapitulate. I have provided a sketch of how a semantic theory that compositionally assigns semantic values to sentences in context can play a part in a wider model of human linguistic communication. In Chapter 2, I attempted a justification of a Fregen anti-psychologistic model of the contents of language and thought. In Chapter 3, I introduced the basic idea of a semantic theory, as it is familiar from logic and semantics, outlining some motivations for going beyond a simple extensional model and adopting the idea of an intension, as pioneered by Carnap. In Chapter 4 I then showed how this basic framework can be enriched to model the phenomenon of context dependence in natural language, i.e. the phenomenon of sentences that it is literally correct to utter on some but not on other occasions. We saw that in Kaplan's system of "Demonstratives", there are two semantic mechanisms of context dependence. On the one hand, the semantics no longer assigns intensions directly to expression-types, but rather, it assigns these intensions relative to a context of use. We can describe this as the semantics assigning a *character* to each expression, i.e. a function from a context to an intension. Equivalently, we can describe it as the semantics assigning intensions to expression-context pairs. Within this first model of context dependence, which we might call "indexicality", the intension of a sentence can vary with the context in which it used. The second kind of context dependence in Kaplan's system does not involve such variation in intension. Even if the intension of an expression remains the same across contexts, it may be an intension that varies in extension with some factor in the circumstance of evaluation, as for example in the case of tensed sentences in the temporalist model. Thus, the very same intension may be true as evaluated against one circumstance of evaluation and false as evaluated against another. The temporal intension that MK is hungry, for example, behaves in this way. This second form of context dependence might be called "circumstance-sensitivity". Thus the source of a given variation in the correctness of the same (closed) sentence can be attributed either to non-constant *characters* (different content in different contexts), or to non-constant intensions (contents that have different truth-values with respect to different circumstances of evaluation). I also considered the possibility of treating natural language sentences as either predicates or open sentences (sentences containing free variables) in the semantics. The resulting "Predicate View" and "Free-variable View" represented two further ways in which context-dependence can be modelled semantically. These are proposals to model context dependence in

ways that are not initially envisaged in Kaplan's system, but for which room can be made without touching the core framework.

Chapter 5

In the last chapter, I considered a range of purported phenomena of context dependence and how they are best modeled within an intensional semantics like Kaplan's. I compared the option of treating them as forms of indexicality and as forms of circumstance relativity. I also considered the "predicate view" and the "free-variable view". I offered a tentative assessment, according to which the indexical approach to these phenomena faced problems in the treatment of certain qualifiers or operators ("For Peter …", "Given what Peter knows, …" etc), problems the other proposals didn't face. In the comparison between the relativist model, i.e. the model that postulates circumstance-sensitivity in these cases, with the predicate- and free-variable-view, I conceded that there might be a fundamental case for giving up intensions altogether and operating with predicates or sentences containing free variables instead. This might mean, for example, treating modal as well as tense operators generally as quantifiers that can bind the free variables in question. However, this would mean a more radical departure from the customary semantic framework, so that, within this framework, the model of these phenomena that was based on circumstance sensitivity remained the option of choice.

Evans' objection

The success of the circumstance-sensitivity model of context dependence, however, depends on its ability to meet a fundamental objection that has been made some time ago by Gareth Evans against the temporalist view of tensed sentences, in his "Does Tense Logic Rest upon a Mistake?". Temporalism, as we have seen, is one of the best known semantic proposals involving circumstance-sensitivity and Evans' worries about temporalism equally concern other forms of circumstance sensitivity. Evans draws attention to some important issues that need to be addressed, so the current chapter will do this. I shall begin by explaining Evans' argument in detail. It will emerge that Evans' worries can be answered decisively in a range of cases, those which I shall label "moderate relativism". Evans' raises deeper worries in a more restricted range of proposals involving circumstance sensitivity. I shall examine two such areas: sentences concerning future contingents and those involving vague predicates. In both of these cases, Evans' point needs to be addressed, but can be addressed satisfactorily.

6.2 Empirical Bridge Principles

In the semantic framework we have been discussing, the semantics assigns contents to sentence-context pairs. It does so by defining a semantic value function which assigns to each sentence *s*, context *c*, assignment *a*

and circumstance of evaluation *p* the extension SV(*s*, *c*, *a*, *p*). Now, such an assignment by itself, as we saw, is devoid of empirical content. In order to fill it with empirical content, we need to add some bridge principle that connects the abstract semantic value function with independently accessible data. On our simplified model, such a bridge principle was given by the thought that the literal truth of utterances was to correspond to the truth-value of the content expressed by the sentence used in the context in which it was used, with respect to the circumstance of evaluation given by that context of use. Kaplan explicitly articulates such a principle:

(K) An utterance of a (closed) sentence is *true* just if the content (intension) expressed by the sentence in the context of the utterance assigns the value true to the circumstance of evaluation *of* the context.[93]

Kaplan's empirical bridge principle

Kaplan's idea, presumably, is that "*true*" on its first, italicized occurrence in (K) expresses a notion of the truth of "occurrences" of sentences to which we have some independent empirical access, so that we now have an empirical handle on the theory. So, the semantics tells us which contents sentences express in which contexts of use, and principle (K) allows us to use that information to make predictions about the *truth* of utterances, and the *truth* of utterances is meant to be something which competent speakers of the relevant natural language are able to recognize in a theory-independent way. For example, competent speakers of English with good information about the meteorological situation at various times and places would be able to assess independently whether an utterance of "It's raining." in a certain context is true, and this result can then be compared with the prediction as to the truth of utterances of that sentence that the semantic theory makes in conjunction with (K) and the same meteorological information.

***Exercise 10**: Suppose that L4 is a semantics for a fragment of English, and that the correct logical form sentence corresponding to "Peter dances." is "**dances**(**Peter**)". Spell out what predictions L4 together with (K) would allow us to make about utterances of "Peter dances".*

93 Kaplan 1977, p. 522: "If *c* is a context, then an occurrence of φ in *c* is true iff the content expressed by φ in this context is true when evaluated with respect to the circumstance of the context.". In terms of Kaplan's formal system, the stipulation is that an occurrence *u* of a sentence *s* in context *c* is *true* iff for all/some *f*, T(*s*, *c*, *f*, T(*c*), W(*c*)), where T(*c*) and W(c) are the time and the world of the context *c* respectively. Kaplan is speaking of "occurrences", i.e. sentence-context pairs. This only allows us to test the semantics empirically if we make *further* assumptions about how occurrence-*truth* is related to utterance truth, i.e. of concrete events of utterance (or some other pretheoretical notion). For reasons of presentation, I am presenting his principle as a principle that is directly about utterances.

(K) also involves the idea of an utterance of a sentence occurring *in* a context of use, and the idea of a circumstance of evaluation being that *of* a context of use. Thus, applying (K) requires grasp of the relation *in* which determines for each utterance the context *in* which the utterance takes place, and also grasp of the relationship *of* such that for each context of use there is a circumstance of evaluation that is the circumstance *of* that context. Using (K) as an empirical bridge principle requires that each utterance determine a context of use, that each context of use determine a circumstance of evaluation and that these two determination relations figure in (K) and make it true.

Mechanisms of context-dependence.

Now, given an empirical bridge principle like (K), we can ask ourselves again: what forms of context dependence can a semantics like L4 model? Context dependence, broadly conceived, consists in a variation of *truth*-value of utterances of the same sentence. So let's assume that we have two utterances U1 and U2 of the same sentence S, and U1 is literally *true*, while U2 is not literally *true*. What are the theoretical possibilities of how this could arise in a Kaplanian semantics like L4 combined with a bridge principle like (K)? First, S might be ambiguous, i.e. it might be associated with different logical form sentences S1 and S2, and this would explain the difference in *truth*-value. This variation is due to "pre-semantic" variability: the variation in *truth*-value is explained by the input into the semantics being different in the two utterances. Secondly, S might have non-constant *character*, so that its intension in the context of U1 is different from its intension in the context of U2. The difference in intension then explains the difference of *truth*-value. Here the mechanism is properly semantic: we have the same logical form sentence as input to the semantics, and the difference in *truth*-value of U1 and U2 comes about because the intension expressed by the sentence depends on the context, and the context of U1 was relevantly different from that of U2. Call this mechanism "indexicality". Thirdly, the sentence might be circumstance-sensitive. Thus, the intension S has in the context of U1 is the same as the intension that S has in the context of U2. However, the truth-value of that intension varies from some circumstances to others, and the circumstance determined by the context of U1 differs relevantly from that determined y the context of U2.

If S is not a closed sentence (e.g. if it is a predicate or an open sentence), then the case is not covered by (K). So the distinctive pattern of explanation for contextual variation offered by the predicate-view and the free-variable views would need to be based on some further bridge principle over and above (K).

6.3. Temporalism and Other Approaches to Tense

Actual examples of using circumstance sensitivity

Tensed sentences, like "MK is hungry." or "Peter is dancing." clearly exhibit context dependence: an utterance of the very same sentence may be *true* while a different utterance of the very same sentence is not. Tense logicians or temporalists use the third mechanism mentioned above to account for this variation. That is, they treat such sentences as expressing temporally variable intensions, and they explain the variations in truth-value just mentioned by pointing towards relevant differences in the circumstance of evaluation determined by the context determined by the utterance. This approach is exemplified by, for example, Prior (1967), (1977) and Blackburn (1994), and Recanati (2007). The mechanism is also invoked by a good number of people to account for different phenomena of purported *truth*-value variation (see Chapter 5 above). But for the moment, let us focus on temporalism.

Temporal intensional operators

As we saw in Chapters 4 and 5 above, this sort of account allows one to treat a range of temporally qualifying expressions or inflections as sentential intensional operators, i.e. operators that express functions from contents to contents. For example, "Sometimes" can be seen as expressing a function that takes us from a content p to a content **Sometimes**(p) which is true at a circumstance $<t, w>$ just if for some t^*, p is true at $<t^*, w>$. Similarly, the past tense can be treated as expressing a function that takes us from a content p to the content **Past tense**(p) which is true at a circumstance $<t, w>$ just if for some time t^* prior to t, p is true at $<t^*, w>$.

Indexical view of tense

Such an approach to tense must be contrasted with an indexical approach which employes the second mechanism mentioned above: on that approach, a sentence like "MK is hungry." has non-constant character: the sentence expresses different (eternal) contents in different contexts. If t is uttered in two different contexts $c1$ and $c2$, and the respective times of these contexts, T($c1$) and T($c2$), are different, then the two utterances express different (eternal) contents. Now, if the time T(c) of a context c is always the same as the time t in the circumstance of evaluation $<t, w>$ *of* context c, then bridge principle (K) will ensure that the indexicalist approach predicts the same truth-values for utterances of unembedded tensed sentences as the temporalist approach.

Predicate view and free-variable view of tense

The temporalist and indexicalist approaches must also be distinguished from a predicate approach to tense (exemplified by Quine 1960) and from a free-variable approach (probably exemplified by King 2003). On both of these views, the intension of "MK is hungry." in a context c is not a content, i.e. a sentence-intension. But they can be viewed as delivering truth-values at the world of c relative to individuals, i.e. times. If we take the rel-

evant time to be the time *of c*, then the predictions will again be the same as those made by the temporalist.

However, there are nevertheless significant differences between the four approaches. They will have to give a substantially different treatment of temporal qualifiers such as "Sometimes" and the past tense. A proponent of the predicate view will be able to regard the temporal qualifiers simply as noun phrases. Proponents of the free-variable view will be able to regard the temporal qualifiers as binders of sorts (such as binding quantifiers). The indexical approach has some difficulty in explaining what is going on. Consider

(1) MK is hungry.

(1S) Sometimes MK is hungry.

(1P) MK was hungry.

"Sometimes" and past tense

(1S) no longer exhibits the intuitive variability in utterance truth that (1) exhibits: no matter what the time of the context in which we utter (1), there is no variation in utterance truth, so on the indexical account, while (1) has non-constant character (1S) has constant character – "Sometimes" has removed the context dependence of (1).[94]

(1P), on the other hand still has its non-constant character, though one that differs from the character of (1). Suppose MK has always been hungry until time *t*, but at *t*, his hunger is finally satisfied. Then an utterance of (1) at *t* will be false, while an utterance of (1P) at *t* will be true. So the characters of (1) and (1P) must be different.

Disadvantage of indexical approach

Which semantics of the two temporal constructions would achieve this? One possibility is to treat "Sometimes" and the past tense as constructions that operate on the character of sentences, in a manner akin to Kaplan's "monsters". But the main point of differentiating between context dependence and circumstance dependence seemed to be to differentiate between shiftable and unshiftable parameters of context dependence, with indexicals constituting the unshiftable kind. So treating some indexicals as shiftable would undermine this motivation. So the predicate view, the free-variable view and the temporalist view have an easier time with the temporal qualifiers. In other words: approaches that regard the character of (3) to be constant (i.e. "invariantist" approaches) seem to be at an advantage. It is therefore interesting that Evans (1985) singles out the temporalist (pat-

94 I am pretending that "Sometimes" always quantifies unrestrictedly over all times, past, present and future.

tern 5) approach as facing problems with a compositional treatment of temporal constructions.

6.4 Evans' Objection

The difficulties Evans sees are due to certain requirements he imposes on the semantic values a semantic theory should assign to sentences in context. Evans believes that the semantic values assigned by a semantic theory to sentences in context should *immediately* and *as part of the semantic theory* yield evaluations of utterances as correct or incorrect. This conception of semantic values leads him to reject the temporalist account as just outlined.[95]

Evans is considering bridge principles

In considering various versions of tense logic, Evans explicitly discusses bridge principles similar to (K) and how they interact with a semantic theory that assigns intensions of the sort mentioned, i.e. temporal contents, to tensed sentences. Evans is simplifying by ignoring indexical phenomena and considering semantic theories that simply define a truth-predicate "True$_t$" with a time index, which can be applied to sentences. Thus, in effect, the truth-predicate to be defined in the temporalist semantics is a two-place predicate relating sentences with times. Ultimately the semantic truth-predicate must be significantly linked with some notion of correctness for *utterances*: Evans says that "it seems reasonable to require" that a semantics enable someone who knows it "to determine that certain utterances are correct or incorrect" (1985, p. 346). It seems clear that Evans has in mind precisely the sort of considerations discussed in section 6.2 above: in order for the semantics to be empirically significant, it needs to yield predictions as to the correctness of utterances, i.e. what we were calling utterance *truth*. The three versions of tense logic Evans discusses differ only in the (K)-like bridge principle with which they link utterance correctness and Truth$_t$ as defined by the semantics.

The first version, T1, postulates a bridge principle that leaves even utterance truth variable (Evans 1985, p. 347):

(E6) $(\forall S)(\forall u)(\forall t)$ [Of $(S, u) \supset$ (Correct-at-t$(u) \equiv$True$_t$(S))]
[For all sentences *S*, utterances *u*, times *t*, if *u* is an utterance of *S* then *u* is correct at *t* if and only if *S* is True at *t*.]

95 By parity of reasoning, it should also have led him to reject the two other invariantist accounts, i.e. the predicate view and the free-variable view. This should worry those who want to invoke Evans 1985 against temporalist and other relativist approaches but who themselves pursue a predicate view or a free-variable view.

(E6) defines a notion of "correctness" for utterances – Evans wants to be clear and explicit about the difference between the relative notion of truth$_t$ defined for sentences in the recursive clauses of the semantics and the theory-independent notion of correctness defined for utterances or assertions. However, the idea that the definition of such a notion for utterances in terms of the semantic notion of truth is needed in order to view a semantic theory as a substantial empirical theory is the same as in Kaplan.

Evans: T1 is incoherent

According to Evans, T1, with its bridge principle (E6), does "not provide for the stable evaluation of utterances as correct or incorrect" (p. 356). He regards such an approach as "not coherent" (p. 349). I shall come back to this claim below. For now it suffices to observe that the version of temporalism we were discussing above is not like Evans's T1, for we were operating with Kaplan's bridge principle (K), which does provide for the stable evaluation of utterances as correct or incorrect, namely as *true* or not *true*.

The two other versions of tense logic Evans discusses are quite similar. T2 is similar to Kaplan's version of temporalism as I described it above: the semantics defines a predicate "true$_t$", and then there is a bridge principle much like Kaplan's (K):

(E7) $(\forall S)(\forall u)(\forall t)\ [(\mathrm{Of}(S, u) \wedge \mathrm{At}\ (t, u)) \supset (\mathrm{Correct}(u) \equiv \mathrm{True}_t(S))]$
(Evans 1985, 348)
[For all sentences *S*, utterances *u*, times *t*: if *u* is an utterance of *S* and *u* occurs at *t*, then *u* is correct if and only if *S* is True at *t*.]

Version T2 of tense logic

Even though Evans does not distinguish, like Kaplan, between the time of the context and the time of the circumstance of evaluation, his characterization is nevertheless recognizably that of a position where the semantics defines a notion of truth$_t$ that is relative to times, and a bridge principle then makes an additional stipulation that links utterance *truth* with truth$_t$.

Evans' criticism of T2

Evans puts a lot of emphasis on the idea that in T2, bridge principle (E7) constitutes an additional stipulation, one that did not already follow analytically from the definition of "true$_t$". This is precisely the point at which the third version, T3, differs from T2, and the point at which Evans criticises T2. For Evans says that if principle (E7) is not already contained in our definition of "true$_t$", then (E7) in effect assigns a second semantic value to sentences in context. The definition of "true$_t$" assigns one semantic value (a function from times to truth-values), and (E7) then assigns a second semantic value to sentences in context, namely a truth-value.

Answering Evans' criticism of T2

Evans compares T2 to a supervaluational semantics and claims that it "has the consequence that the utterance of a present-tense sentence (a present-tense sentence in context) has *two* semantic values, and is (per-

haps harmlessly) ambiguous" (p. 353). However, it is not clear why Evans wants to treat (E7) as assigning a second semantic value to sentences in context. Bridge principles like (E7) need not be regarded as part of the semantics. Rather, they might be part of a further theory of language use or of illocutionary force. Evans implies at one point that truth-values are the semantic values appropriate to sentences (p. 353), so Evans may be operating with a special constraint on what can count as a semantic value of a sentence in context, namely that only a truth-value can. However, I see no reason why we should follow Evans in this.

Evans then goes on to outline T3, which has (E7) not as a bridge principle, but as immediately following from the definition of "true$_t$", so that the semantic value of a sentence in context is simply the "stable" or absolute correctness value assigned to it by (E7), and not the intension (function from times to truth-values).[96] This ensures that there is only one semantic value per sentence in context, but there is another consequence which Evans then elaborates on, namely the treatment of the temporal operators is no longer compositional in the sense of making the semantic value of a complex sentence in a context a function of the semantic values of its constituents in the same context. He says:

T3 asserts that the semantic value which the sentence 'P(X)' has in a context is a function of the semantic value which X would have in *another* context. For, on the present interpretation, the recursive clause (2)[97] says *roughly* that the utterance of 'P(X)' is true iff the utterance of X at some earlier time would have been true. (p.357)

Evans' criticism of version T3 of tense logic

This does seem to be the consequence of treating (E7) as defining the semantic values of sentences in context, and the most plausible way of construing temporal constructions. However, as we saw, there is no need to prescind from taking temporal intensions as being the semantic values of sentences in context, and I suggest that this is the best interpretation of the temporalist's intentions.

I conclude that temporalists or tense-logicians like Kaplan can coherently take Evans's approach T2 and regard (K) as an additional bridge principle. Contra Evans, there is no need to say that (K) defines extra semantic values of sentences in contexts.

96 There is an analogy with the definition of "satisfaction" or "truth" of formulae (open or not) in the standard Tarskian semantics for first-order languages: this notion does not by itself provide semantic values for sentences, but only has the status of an auxiliary notion in the definition of truth for closed sentences.

97 (2) is the recursive clause for the past tense: "(2) For any time *t,* and any sentence S, true$_t$('P'^(S)) iff there is a time *t'*, earlier than *t,* such that S is *true$_{t'}$*." (1985, p. 344).

6.5 Moderate Relativism in General

A more general conclusion we may draw at this point is that it is unproblematic to operate with mere intensions (i.e. functions from some points of evaluation to truth-values, and not truth-values) as the semantic values of sentences in context, as long as we have a bridge principle like (K), which ensures that there is an absolute (or "stable") notion of utterance *truth*. The crucial point here is that each utterance is taken to occur *in* a context, and that each context determines *its* circumstance of evaluation, so that we can then define the *truth* of an utterance as the truth-value of the uttered sentence's intension in the context of the utterance for the circumstance of evaluation determined by that context.

Put in these general terms, the approach can be generalized to different ways in which the content of a sentence in context may depend on factors in the circumstances of evaluation. In Kaplan's framework, the possible world factor in the circumstances of evaluation is another example: we can treat a sentence like "MK is hungry." as expressing, in each context, a proposition that varies in truth-value with a possible world, and the corresponding intension (function from circumstances to truth-values) can figure as the semantic value of the sentence in the context. For (K) will ensure that utterances can be stably evaluated as correct or incorrect. It is usually regarded as unproblematic to assume that each utterance determines (via its context) a possible world, because each utterance will be made in some possible world, and this is the world that figures as the world in the utterance's context's circumstance of evaluation.[98]

If we extend Kaplan's framework and add further factors to the circumstances of evaluation, the same holds. If we add, for example, an agent (in order to accommodate *de se* propositions), all we need to ensure is that each context of use determines an agent, so that we may continue to assume that contexts of use determine relevant circumstances of evaluation. If we add a standard of taste or "judge" (as proposed by Kölbel 2002, 2003, Lasersohn 2005, and others), we need to ensure that each context of use determines a relevant standard of taste or judge, so that utterances in these contexts can be evaluated with respect to the circumstances containing the relevant standard or judge. The same holds for factors like states of information (for epistemic modals – Egan, Hawthorne, Weatherson 2005, MacFarlane forthcoming), interests (for knowledge ascriptions – MacFarlane 2005b) and several other recent so-called "relativist" proposals.

98 Though the next section will show that the assumption is, after all, problematic.

Suppose we have reason to add a novel factor *f* to the circumstances of evaluation. Then the general recipe for meeting Evans' requirement is this:

1. Find a suitable function F1 that determines a unique context of use F1(*u*) for each utterance *u*.

2. Find a suitable function F2 that determines a unique circumstance F2(*c*) for each context of use *c*.

3. Use a bridge principle like (K) to evaluate utterances in the following way: An utterance *u* of a sentence *s* is correct just if the content ex pressed by *s* in F1(*u*) has the value "true" at F2(F1(*u*)) (or in other words: just if T(*s*, F1(*u*), *a*, F2(F1(*u*))) = 1)

Functions F1 and F2 are suitable, of course, if both together support the (K)-like bridge principle mentioned in point 3., i.e. if the predictions as to utterance *truth* generated by the semantics and the bridge principle are empirically adequate.

Following Recanati (2007), I shall call proposals that follow this recipe forms of "moderate relativism". Moderate relativisms disagree with Evans's criticism of T2, but they comply with his requirement that the semantics (plus the (K)-like principle) should provide a stable evaluation of utterances as correct or incorrect.

6.6 Radical Relativism: Future Contingents

Suppose we had reason to add a new factor *f* to the circumstances of evaluation but could not find suitable functions F1 and F2 which determine for every utterance *u* a suitable value of *f*, and thereby provide a circumstance of evaluation F2(F1(*u*)) with respect to which *u* can be evaluated as correct in accordance with a (K)-like principle. In that case we cannot follow the above recipe for moderate relativists, and cannot use a bridge principle like Kaplan's (K) to provide a stable evaluation of utterances as correct or incorrect. Thus, it seems, we would fail to comply with the requirement Evans uses to reject T1.

There are two questions: Is there such a factor *f* that we have good reason to introduce? and Is it coherent to fail to meet Evans's requirement of stable evaluation? I shall argue that the existence of sentences regarding the contingent future provides a reason for positive answers to both questions.

Suppose we operate within Kaplan's original framework and consider the following sentence

(2) Obama is the winner of the 2008 US presidential elections.

The semantics should assign to (2) the same intension/proposition in each context of use, namely a non-constant function from circumstances of evaluation <*t*, *w*> to truth-values. (K) tells us that an utterance of (2) is true/correct just if this function has the value "true" at the circumstance <*t*, *w*> determined by the context in which the utterance is made. This may be unproblematic if the utterance occurs after the elections. But suppose an utterance of (2) is made in early 2008 before the elections. Then the moderate relativist's recipe tells us that an utterance of (2) in early 2008 already determined a context of use, which in turn already determined a circumstance of evaluation <*t*, *w*> with respect to which the utterance was to be evaluated. This may be unproblematic on a block-universe view of the future according to which each utterance event determines a unique future course of events, i.e. a complete possible world. However, a genuinely indeterministic branching universe view is incompatible with this idea. For if the utterance in early 2008 already determined a unique world of evaluation, then it was already determined in early 2008 that Obama would win.[99] Worlds other than the world determined by the utterance, in particular worlds in which Obama fails to win, will not, at that point, have been genuinely metaphysically open possibilities.

Thus, the view that the future is genuinely open (as opposed to merely unknown), is prima facie incompatible with the moderate relativist's recipe for stably evaluating utterances as correct or incorrect. An utterance at best determines a set of worlds, namely the ones that are compatible with the course of events up to the time of utterance. However, in the case of an utterance like the utterance of (2) in early 2008, this set contains worlds at which Obama wins and worlds where he doesn't win.

This by itself does not yet provide a conclusive motivation for violating Evans's desideratum of a stable evaluation of utterances. For we could adopt a new type of bridge principle that takes into account the fact that there are no appropriate determination relations, such that each utterance determines a unique context of use, and each context of use determines a unique circumstance <*t*, *w*>, such that the utterance must be evaluated with respect to *that* circumstance. Let us assume for simplicity that each utterance does indeed determine an appropriate context of use, but that the context of use fails to determine a unique world, and therefore a

99 See Belnap et al. 2001, pp. 163-4.

unique circumstance of evaluation.[100] Then we could adopt a new bridge principle:

(K*) An utterance of a sentence is *true* (*correct*) just if the content (in tension) expressed by the sentence in the context of use determi ned by the utterance assigns the value true to every circumstance of evaluation compatible with that context.

(K*) does indeed provide a stable evaluation for utterances. However, according to (K*), our utterance of (2) in early 2008 receives the stable evaluation "not true". This is very hard to accept, for we would want to say today that that utterance was true, or at the very least *is* true. If (K*) is to be our bridge principle, we would have to interpret our view today that the utterance was true or correct in such a way that it does not involve truth or correctness as defined by (K*).

Those prepared to swallow the consequence that the utterance of (2) in early 2008 was not true/correct according to the bridge principle we use to evaluate utterances may have some room for manoeuvre. They might say that our intuitive judgements as to the correctness of the utterance concern some different norm of correctness distinct from true/correct as defined by (K*). It is not so easy to articulate such a supplementary norm. If we want to retain indeterminism as well as our impression that the supplementary norm is met after the election, it would seem that the norm should be such that in early 2008 it is not yet determined whether it is met, while after the election it is met. However, if this is the norm that tracks our intuitive judgements of correctness, then it is not clear whether we haven't again violated Evans's desideratum of stable evaluation. We now have a stable evaluation in terms of truth as defined by (K*) which does not track intuitive correctness in the relevant cases, and a supplementary norm which does track intuitive correctness but is not stable. Thus it would seem that it is better to stick to one norm that tracks intuitive correctness, which is not stable in the relevant cases, but which is stable when indeterminism allows it.

This concludes my case for a positive answer to the first question, i.e. the question whether there may be a factor *f* which prevents us from defining an absolute notion of utterance truth.

[100] Alternatively, we might say that utterances fail to determine a unique context of use. This would take us in the direction of Cappelen's "content relativism", namely the view that the very same *utterance* is taken to have different contents at different points of interpretation (Cappelen 2008), or the Egan's view in Egan 2009. Similar considerations apply.

The second question was whether failing to meet Evans' stable evaluation requirement leads to incoherence, as Evans claims. Is there a suitable bridge principle that violates Evans's desideratum but nevertheless provides a coherent account of the correctness of utterances? I believe there is, and it has a form analogous to the defining principle (E6) of Evans's T1:

(K**) An utterance *u* of a sentence *s* is *true* (*correct*) at a context of eva luation *e* just if the content (intension) expressed by *s* in the con text of use determined by *u* assigns the value true to every cir cumstance of evaluation compatible with *e*.

Let us go with MacFarlane (2008) in taking a context of evaluation simply to be a concrete situation where some thinker evaluates some utterance. How can a principle like (K**) "assist the subject in deciding what to say" (Evans 1985, p. 349)? What is the point of making an utterance of a sentence like (8), be it before or after the election? Indeed, what is the point of believing the content of such an utterance? What are we aiming for?[101]

There are indeed some deep problems to which these questions point. But these are not problems unique to, or caused by, any semantic proposals. Before I got on the plane, I believed that the plane would not crash. Before I went to the bar, I believed that it would be open. What is the point of these beliefs? They are supposed to guide my actions. If they are correct at the current context of evaluation, i.e. now that the plane has landed safely and I am happily sitting in the bar having a drink, they guided my actions well (*ceteris paribus*). But if the future is genuinely open, then there was nothing in the situation in which I originally had these beliefs which determined that I would later get to be in a context of evaluation at which they are true.

If indeterminism is to be compatible with the idea of our beliefs about the contingent future rationally guiding our actions, then we must assume that despite the openness of the future there are nevertheless reasons available to us that justify our beliefs despite their being true in some but not true in other currently possible futures.

One kind of reason would be constituted by knowledge that in each of the possible futures having the belief was preferable to not having it, even in those where it is not true.[102] This is a Pascal's wager-type of justification

[101] NB: these question arise equally for the sort of supplementary norm I mentioned in connection with (K*), so even an alternative supervaluationist solution would have to deal with these issues.

[102] Cf Belnap et al. 2001, p. 171: "We shall argue that it makes sense to wonder about what history has not yet decided so long as history *will* decide the matter. We shall

that may be available in some cases, but surely does not cover all the cases in which we take ourselves to have rational beliefs about the contingent future. In a future where the bar is not open, contrary to my belief that it would be, I might very well regret having had the belief without thinking that it was irrational.

What else could justify my belief that *p* at moment *m*1 at time *t*1, when it is not determined whether at *t*2 I'll be at moment *m*2, where *p* is true or at moment *m*3, where *p* is not true? Presumably the justification will have to consist in evidence that somehow favours some futures as more worthy of preparing for them than others. For example, on some views, we may be rationally justified in assigning subjective probabilities (degrees of belief) to propositions concerning the contingent future. On other views, some futures may be objectively more probable than others, and our evidence is a rational guide to these objective probabilities. Or perhaps some futures stand in non-deterministic causal relations to the current state of the world, and we have reasons for thinking these obtain. I take it that these are the sorts of options an indeterminist has for explaining the rationality of belief concerning the open future. Whatever the answer, the problem of rational belief and action concerning the contingent future is a general problem any indeterminist has to face.

I want to argue that if the indeterminist has a solution to this general problem of rational belief about the future, then a number of different answers to Evans' questions are available. In other words, once the indeterminist has a general solution to the problem of rational belief concerning the contingent future, she can explain how (K**) assists speakers in deciding what to say, and say what the point is of asserting or believing a proposition in a situation compatible both with future courses of events that evaluate the assertion or belief as true/correct according to (K**) and with those that don't. Let's say that our answer to the general problem is that it is rational for an agent to believe what is "likely by her evidence", and that this is rational because it will make her prepare in her actions (if rational) for those futures that most deserve it (leaving open what exactly this amounts to).

We can now say, for example, that asserters are subject to a norm of sincerity, i.e. that speakers aim to assert only propositions they believe (a pro tanto norm which can be outweighed by other conderations). Or we can say that asserters are subject to a norm of truth and should avoid situations of evaluation in which they have asserted something that is not

also argue that it makes sense to assert *A* when *A*'s truth-value is not settled at the moment of use; the idea is that assertion is an act that has consequences for the speaker no matter how things turn out."

true/correct (according to (K**)). Or we might say that asserters are subject to the norm that they have to provide justifications for what they assert if they are asked to, and that they have to withdraw any past assertion in a situation with respect to which that assertion is not true/correct according to (K**). In each of these cases, the indeterminist's solution to the general problem of rational belief will allow us to make sense of the special problem of rational assertion. On the simple sincerity norm, it is belief that is governed by the general norm that requires believing what is likely by one's own evidence. What the speaker is allowed to assert then falls out from what she believes. On the truth-norm, the speaker will, in her assertions, prepare for those futures that are by her evidence more likely. On the justification/withdrawal norms, again, the speaker will prepare for those worlds that are, by her evidence, more likely.

This conclusion seems to show that a relativist semantics for sentences about the contingent future by itself does not yet provide motivation for any particular account of assertion.[103] Once the basic problem of the indeterminist is solved, a range of different solutions are available, and some other considerations will have to decide between alternative accounts of assertion.

I conclude that the view that the future is genuinely open provides good reason to violate Evans's requirement of stable evaluation and to adopt a bridge principle analogous to Evans's (E6), namely (K**). The resulting view is coherent, *if* the idea of rational beliefs about the contingent future is coherent at all under the assumption of an open future.

6.7 Radical Relativism: Vagueness

Vague predicates give rise to the sorites paradox

Many natural language predicates are vague in the sense that they seem subject to tolerance constraints and therefore generate sorites paradoxes. For example, the predicate "is rich" is vague because it seems to be subject to the constraint that if someone is not rich then receiving a small amount of money such as one cent will not make that person rich. Thus a sorites paradox can be formulated as follows:

(A) A person with possessions worth 0 Euros is not rich.

[103] MacFarlane 2003, 2008 and Kölbel 2002 seem to think it does.

(B) If a person with possessions worth *n* Euros is not rich, then a person with possessions worth *n* + 0.01 Euros is not rich either.[104]

(C) A person with possessions worth 100 million Euros is not rich.

(C) seems to follow from (A) and (B), but while (A) and (B) seem clearly true, (C) seems clearly not true.

The normality of vagueness

This phenomenon is widespread. Countless natural language predicates are vague in this sense. The vagueness of these predicates does not seem to be an impediment to their usefulness in communication. Similarly, the concepts expressed by vague predicates do not seem to create any major problems for our everyday thought. The vagueness of natural language predicates and the concepts they express is therefore not some deficiency, shortfall or malfunction. Vagueness is perfectly normal.

If vagueness is normal, then semantic frameworks for natural languages, ought to be able to accommodate it. There are several ways in which room can be made for vagueness within standard semantics.

There is a minority of philosophers, the *epistemicists*, who hold that vagueness is not a semantic problem, but rather reflects our inability to know the exact borderlines of the extensions of the predicates (and concepts) we use. On this view, premiss (B) in the above sorites is simply false. There is a truth of the form

(D) A person with possessions worth *n* Euros is not rich and a person with possessions worth *n* + 0.01 Euros is rich.

Epistemicism about vagueness

But we cannot know that truth because of general principles concerning knowledge.[105] It is this fact that explains why (B), despite its falsity, is so attractive. According to epistemicism, then, vagueness is an epistemic, and not a semantic phenomenon. The meaning of vague as well as non-vague predicates determines for each object whether it is in that predicate's extension or not.

The majority, however, finds the epistemic view incredible, in large part because it remains mysterious how the precise extensions of vague predicates are determined. The majority instead believes that vagueness is a semantic phenomenon, i.e. that the meanings of vague predicates fail to determine exact extensions. I will not provide any reasons to favour se-

[104] There are different ways of formalizing "a person" in (B), see Pagin 2009. These differences will not matter for the current discussion.

[105] See Williamson 1994 and Sorensen 1988.

mantic views over epistemicism. I shall merely assume that the semantic view is correct. Starting from that assumption, I will make a case for a treatment of vagueness in terms of circumstance sensitivity, understood as semantic, within standard semantic frameworks like the ones we have been considering in chapter 4 and 5.

Vagueness as semantic

On the semantic view of vagueness that we are adopting, vague predicates fail to determine for each object whether the predicate applies to it. For short, vague predicates are extensionally indeterminate. In the framework we have been discussing, there are three ways in which a semantically non-deficient predicate (thought of as a phonetic type) can fail to determine an extension: the predicate may be ambiguous, indexical or its extension may vary with the circumstances of evaluation.

Vagueness as ambiguity

It is not easy to say why exactly we shouldn't say that the source of the indeterminacy is ambiguity. This is in part because there are no clear criteria that allow us to distinguish ambiguity from indexicality. Presumably, the reason why we treat the word "coach" as ambiguous in English has something to do with the fact that it just seems a historical accident that the same phonetic type bears the two distinct meanings, and that we need to decide with the help of the context, what the intended meaning is. We do not treat the word "she" as ambiguous, presumably in part because it seems to be a systematic, non-accidental feature of the word that it pcks out different people in different contexts. But not all cases are clear. For example, it is hard to understand why, in general, personal names like "John" or "Mary" are treated as ambiguous rather than indexical. Presumably in the case of vague predicates, the kind of ambiguity we would have to postulate would be quite systematic, and not at all a historical accident, which, in general, is a reason not to attribute extensional indeterminacy to ambiguity.[106]

Vagueness as indexicality

The second possible source is indexicality. A predicate may fail by itself to determine an extension because its character is a non-constant function. For example the predicate "is my uncle" expresses different properties when used by different speakers. When used by you it expresses a (relational) property instantiated by your uncles (if any), and when expressed by me it expresses a property instantiated my uncles. Accounting for the vagueness of predicates such as "rich" in terms of a special kind of indexicality is in fact a widely discussed option. Many so-called "contextualists" about vagueness seem to have in mind such a view.[107] The central idea is

[106] See Kölbel 2010a for more detailed consideration of the possibility that vague predicates are systematically ambiguous.

[107] Kamp 1981, Raffman 1994, 1996, Soames 1999, 2002, Fara 2000 and Shapiro 2003 have all defended the view that vague predicates are context-relative in some way. It

that the tolerance premiss (B) above only seems to be true because of a special type of indexicality of "rich". The intension of "rich": varies with the context of use in such a way that the borderline between the rich and the non-rich never divides two people who very similar in terms of their wealth when this similarity is salient in the context. Our attention to two similar cases will therefore move the boundary out of sight, which creates the illusion that the boundary does not divide any similar cases. In fact, however, similar cases that are out of sight, and whose similarity is not salient, will be divided by the boundary between the rich and the non-rich.

Problems with the indexical view

Arguably, the indexical view of vagueness is simply a contextualist version of epistemicism: vague predicates, according to it, do have sharp boundaries. But these boundaries move in ways that make them difficult, if not impossible to find. However, what these views cannot explain is why the boundary remains undetectable even retrospectively, i.e. long after the utterance has taken place and when our attention to similar cases can no longer affect the location of the boundary.[108]

Vague contents

Let us, therefore, explore the proposal that vagueness is a certain form of circumstance sensitivity. The proposal is to mimic the temporalist by adding another parameter to the circumstances of evaluation, and saying that the propositions or contents expressed by sentences containing vague predicates vary in truth-value with this parameter. They vary in this way because the vague predicates used to express these contents express vague concepts which themselves vary their extensions with this parameter.

A sharpening parameter in circumstances of evaluation

What are the values of the circumstantial parameter with which the truth-values (extensions) of vague contents (predicates) vary? They are ways of making vague predicates precise consistently with clear cases and with certain a priori principles, i.e. functions that assign to the vague concepts expressed by vague predicates precise extensions. We could call these functions "reasonable standards of precisification". But I will here rely on the terminology familiar from supervaluationism and call them "(admissible) sharpenings". On this view, then, vague predicates express intensions that vary in extension with a sharpening component in the circumstances of evaluation. Consequently sentences containing vague predicates will sometimes express vague contents, i.e. contents that are

is not always clear whether these authors are accounting for the vagueness of vague predicates by claiming that they are indexical in our current sense, or whether they are making proposals of circumstance sensitivity. Raffman (2005) denies explicitly that her proposal is one of indexicality.

108 See Kölbel 2010a for a more detailed account of, and objection to, the view that vagueness is a form of indexicality.

sensitive in their truth-value to a sharpening parameter in the circumstances of evaluation.[109]

Just as in the case of temporal contents, in order now to make sense of vague contents as the objects of assertions or beliefs, we need to answer Evans' question: what constitutes asserting or believing such a vague content *correctly*? To put it in a different way: if our semantics assigns to some sentences in context a vague content as semantic value, what kind of bridge-principle allows us to fill this assignment with empirical content? It seems clear that we cannot just adopt Kaplan's principle (K):

(K) An utterance of a (closed) sentence is correct just if the content ex pressed by the sentence in the context *of* the utterance assigns the value true to the circumstance of evaluation *of* the context.

The moderate relativist's recipe does not apply

The moderate relativist's recipe does not applFor if the circumstance of evaluation involves a specific value for the sharpening parameter for vague predicates, then this would mean that for each utterance involving a vague predicate there is only one sharpening that is relevant for the literal correctness of the utterance. However, this does not seem to be a viable proposal given that the point of introducing vague contents was precisely that there was no single correct sharpening. This means that the recipe for moderate relativism cannot be employed in this case. But perhaps we can articulate an alternative bridge principle.

What we need is an appropriate completion of the following schematic principle:

(VP) An assertion (belief) that *p* occurring in context *c* is correct only if the content that *p* is true at …

Evans' challenge

Once we have such a specification of the conditions under which it is correct to assert or believe vague contents, we can easily articulate a suitable bridge principle.

I would like to broach this task by first distinguishing in the abstract two dimensions in which completions of (VP) can vary, and then argue for each of these dimensions what our completion should look like. In principle, the completion of (VP) could either

109 For simplicity, I am ignoring the indexicality of many vague predicates, such as their sensitivity to a contextually salient comparison class. Thus I am strictly speaking considering only a subclass of vague predicates, namely those whose character is constant – such as, perhaps, "is tall for a British male born between 1975 and 1980". Such predicates are no doubt still vague.

(a) privilege a unique sharpening

or

(b) privilege a range of several sharpenings

and it could either

(i) privilege the same sharpening(s) in each situation of assertion/belief

or

(ii) privilege a different sharpening(or different sharpenings) in different situations of assertion/belief.

Two decisions

I shall advocate a (b)-(ii) completion of (VP). Contextualists about vagueness teach us how option (ii) helps us avoid sorites paradoxes. Let's consider a non-inductive version of the sorites discussed above. Consider a sorites series of people, P_0, P_1, P_2, … $P_{1,000,000}$, such that P_0 has €0, P_1 has €1, P_2, has €2 and so on, each P_i having exactly *i* Euros. Consider the following series of one million applications of modus ponens:

(P0) P_0 is not rich.

(P1) If P_0 is not rich then P_1 is not rich.

(C1) P_1 is not rich.

(P2) If P_1 is not rich then P_2 is not rich.

(C2) P_2 is not rich.

.

.

.

(P1,000,000) If $P_{999,999}$ is not rich then $P_{1,000,000}$, is not rich.

(C1,000,000) $P_{1,000,000}$, is not rich.

Now, on an indexical contextualist approach, each of the constituent modus ponens arguments is valid in the sense that if its premisses are true in a context *c*, then the conclusion is also true in *c*. However, there is no context *c* such that all the non-conditional premisses (C1)–(C1,000,000) are true in it, and there is no context such that all conditional premisses (P1)–(P1,000,000) are true in it. This is because considering more and more people in the series will accumulatively change the context until at some point the context undergoes a sudden reversal (cf. Raffman's 1994 "gestalt switch" and Kamp's 1981 "incoherent" context). If someone were to begin pronouncing the entire argument, then in each of the premiss pairs (P*n*)/(C*n*), "rich" would express a *slightly* different property, until suddenly it would express a significantly different property.

How contextualists treat the sorites

This means that the corresponding generalized conditional premiss

(GP) For all *x*, *y*: if *x* is not rich and *y* has only €1 more than *x*, then *y* is not rich either.

> [There are no *x*, *y*, such that *x* is not rich, y has only €1 more than *x* and *y* is rich.]

is false in every context. Nevertheless, there is no context in which a counterexample of the form

(B) *a* is not rich, *b* has only €0.01 more than *a*, and *b* is rich.

could be uttered and be true at that context. This explains (GP)'s appearance of truth.

If sentences containing vague predicates express vague propositions in the sense outlined above, and if the normative significance of propositional truth is given by a type (ii) completion of (VP), then a structurally analogous response to the sorites is available. On the non-indexical approach, the premisses of the sorites argument are not indexical, but express the same propositions in all contexts of use. However, the propositions expressed vary in truth-value with the sharpening parameter in the circumstances of evaluation. The question we are now considering concerns the normative significance of these relative truth-values, e.g. under what conditions it is correct to assert or believe such a proposition (i.e. how to complete (VP)). According to response (ii), the sharpening or sharpenings relevant for evaluating an assertion will vary as a subject is marched along the sorites series. At the beginning of the series, when we are asking ourselves whether some P_n is rich, we'll have to say that she is not, because we have just previously ruled that P_{n-1} is not rich. The sharpening(s) relevant for the correctness of an utterance in a context *c* obey the following constraint: if two individuals *x* and *y* are relevantly similar (e.g. they differ only by €1), and their similarity is salient[110] in *c*, then the sharpening(s) relevant for judging correct assertability (believability) in *c* will not classify differently the proposition that *x* is not rich and the proposition that *y* is not rich. As the subject moves further and further along the series, however, there will come a point at which the context undergoes a sudden leap (perhaps just because proximity to clear cases of rich is becoming all too obvious).

Using the contextualist's idea

This explains why (GP) is not correctly assertable (believable) in any of the contexts, yet each of its instances is. This in turn explains the deceptive pull exerted by (GP) despite its unacceptability.

Difference beteen indexical contextualism and an approach based on circumstance-sensitivity

Despite emulating some aspects of indexical contextualism about vagueness, the non-indexical approach here proposed clearly differs in other

[110] I here go with Fara's (2000) "saliently similar" rather than with Soames' (1999) "similar and salient".

respects. According to the indexical approach, "rich" expresses a different intension at each stage of the march through the sorites series, whereas on the relativist approach, the intension expressed by "rich" typically remains constant as a subject is moving along a sorites series. It is merely the correctness of calling an individual "rich" and the correctness of believing an individual to be rich that varies as we move along the series.

So, when completing (VP), we should make assertability and believability depend on a *variable* (range of) sharpening(s):

(VP) An assertion (belief) that *p* occurring in context *c* is correct only if the proposition that *p* is true at *S*(*c*).

where "*S*" is some contextual function that I will describe further in a moment. It remains to argue that in completing (VP) we should privilege a *range* of sharpenings rather than an individual one, and then to supervaluate. The motivation for this comes from our intuitions about borderline cases. There are three obvious options for construing *S*:

Three versions of (VP)

(VPa) An assertion (belief) that *p* occurring in context *c* is correct only if the proposition that *p* is true at **the** sharpening determined by *c*.

(VPb) An assertion (belief) that *p* occurring in context *c* is correct only if the proposition that *p* is true at **all** sharpenings in the range determined by *c*.

(VPc) An assertion (belief) that *p* occurring in context *c* is correct only if the proposition that *p* is true at **some** sharpenings in the range determined by *c*.[111]

Not (VPa)

The consequence of (VPa) would be that vague propositions are, in any context, either correctly assertable or correctly deniable, and never both (where correct deniability of *p* is equivalent to correct assertability of not-*p*). This goes against all intuitions: against the intuition that in borderline cases of a predicate one may neither assert nor deny and also against the intuition that in borderline cases one may both assert and deny. Thus, I believe, (VPa) can be discarded.

(VPb) provides an account of borderline cases

As for the remaining two options: it seems that there are two ways of thinking about borderline cases. According to one view (I believe the majority's), borderline cases of richness are cases of where it is neither cor-

111 These three options are clearly not exhaustive. For example we might replace "all" in (VPb) with "most", "many", "a few", or even with "twenty". However, I do not see any reason to think that any of these options is promising.

rect to affirm nor to deny richness. And thus, for some *n*, it may neither be correct to call P_n rich, nor to call her not rich, at least in certain contexts (not, for example, when one has just judged P_{n-1} to be not rich). Option (VPb) is the way to make room for this intuition.[112]

The point of (VPc)

Some have argued that borderline cases are cases where both verdicts are permissible (e.g. Wright 2003). According to them, in a borderline case it is both correct to assert and deny the property in question. A theorist supporting this view would naturally opt for (VPc). However, I am persuaded by the more common conception of borderline case.

Difference from supervaluationism

There is a close structural similarity, then, between the characterization of borderline cases adopted here and the supervaluationist position. What exactly is the difference? – Supervaluationists typically claim that truth is super-truth and that falsity is super-falsity. Thus, supervaluationist semantics involves the claim that some utterances are neither true nor false. The relativist here described, however, does not superevaluate in the semantics: the semantics does not specify super-truth conditions. Rather, the relativist superevaluates at the pragmatic level, when it comes to spelling out the normative significance of the semantic properties of expressions.

One of the difficulties of supervaluationism is that it is committed to the truth of the negation of the general premiss in the Sorites:

(¬GP) For some *x*, *y*: *x* is not rich, *y* has only €1 more than *x*, and *y* is rich.

For (¬GP) is supertrue, i.e. true on each admissible sharpening. This seems to be a problem because there does not seem to be a true instance of (¬GP). It might seem that the current proposal is similarly committed to (¬GP) being correctly assertable. However, this is not so because, as pointed out above, the range of sharpenings which are relevant for assertability vary with the context. A similarity constraint will ensure that (¬GP) is not assertable in any context, while (GP) is assertable in every context.

Advantage over supervaluationism

The upshot, then, is that a principle along the lines of (VPb) states the normative significance of propositional truth. This explains the seductiveness of the sorites and makes good sense of borderline cases without in any way departing from standard semantics. What is new is simply the

112 It s worth noting that even an epistemicist like Williamson can accept this characterization of borderline cases as cases where it is neither correct to assert (believe) nor to deny (disbelieve). For according to Williamson, correct assertability requires knowledge, and belief that is not knowledge is "botched". See Williamson 2000.

role semantic truth plays in assessing assertions and beliefs for correctness. This allows us to articulate a replacement for (K) that is in line with the stable evaluation requirement:

(VK) An utterance *u* of a sentence *s* is *correct* just if the content *s* expresses in the context *c* of *u* has the value "true" relative to all circumstances of evaluation <*t*, *w*, *f*> where *t* = the time determined by *c*, *w* = the world determined by *c*, and *f* = is one of the sharpenings admissible in *c*.

In conclusion, let me return to the two questions I raised at the beginning of §6.6, namely: 1. is there any factor *f* that we have reason to add to the circumstances of evaluation, such that we cannot find a function F1 that determines the context of every utterance and a function F2 that determines the circumstance of evaluation of every context. And 2. Is it coherent to violate Evans' stable evaluation requirement? In the last section I argued that utterances about the contingent future provide a reason to answer both questions in the positive. In this section, I argued that vague predicates provide a motivation for a positive answer to the first question, but that nevertheless a principle of stable evaluation can be found, namely (VK), so that Evans' requirement was met in a way that diverges from moderate relativism's recipe.

6.8 Conclusion

In this chapter I discussed in detail how to respond to Evans' objection to tense logic or temporalism, which at the same time constituted an objection to any approach that postulates, as relativists do, circumstantial factors that are not possible worlds. I showed how temporalists could answer the objection conclusively as long as they used a bridge principle similar to that used by Kaplan. In fact any relativist can use this as a general recipe for answering Evans. However, in the last two sections of the chapter, I showed that there are at least two cases where the general recipe does not apply. First, the case of utterances about the contingent future provided a deep problem for anyone with a branching universe type of indeterminist view. However, I argued that a relativist semantic position here does not introduce any problems that weren't already present in the metaphysical position of this indeterminist. Secondly, the case of vague predicates showed how there is reason to postulate contents of speech that vary in their semantic truth-value with a new circumstantial parameter "sharpening", but that this parameter was such that its relevant value is not determined (as in other cases) by the context of use. Rather, there is a range of relevant values, none of which is privileged. By contrast with the

future contingents case however, this case permitted a principle of stable evaluation to the taste of those impressed with Evans' requirement.

7. Conclusion: the Metaphysics of Perspectival Representation

I began this text by assuming what I called “the basic picture”, namely the idea that we all share the same reality, and we can have more or less successful beliefs about that reality, beliefs that we sometimes share by using the medium of language. I argued that we can legitimately theorize about our thought and speech by describing them in terms of certain abstract entities called “contents” or “propositions”, which are individuated in terms of the conditions under which they can be believed or asserted correctly. These contents are shared in the sense that several different members of a speech community can employ the same concepts and therefore have thoughts with the same content. People make these contents of thought available to others by using language that expresses precisely these contents. Thus, one simple idea of how a language helps this interchange of information is the idea that the meanings of the sentences of a language consist centrally in these contents or propositions. In the main body of the text, I filled this idea with content by expounding some of the ways in which semantic theories describe the properties of natural language expressions in order to explain aspects of linguistic communication.

The focus was in particular on the type of contents that the semantics should assign to, for example, sentences concerning questions of value or questions of what might or might not be. In chapters 4 and 5 we saw that there are several different ways in which a semantics can model the contents expressed by such sentences. The main protagonists were what I called “indexical relativism” and “non-indexical relativism”.

Indexical relativists are impressed by the idea that a proposition or content must be an absolute bearer of truth or falsity. This idea comes to us as a methodological principle from Frege. Thus, if we want to say that the very same expression was uttered truly on one occasion and not truly on another, then we are forced to say also that the propositions expressed on the two occasions were different. Thus, when I say “I am hungry.” at one time, and you say “I am hungry.” at a later time, then we express different propositions. I express a proposition about myself and the earlier time, while you express a proposition about yourself and the later time. The principle also applies if (if!) we want to say about two utterances of a sentence like “Buster Keaton is funny.” or “Marzipan is tasty.” that one is true and the other isn’t. On the Fregean approach, we are forced to say that the same sentence expresses different propositions in the different utterances. In the case of “I am hungry.”, it is obvious how the propositions dif-

fer: one is about me, and one time, the other is about you and another time. But in the case of two utterances of “Buster Keaton is funny.” that have different truth-values, it is not clear how these propositions will differ. One option is to say that the utterance are, in some sense, about the respective speaker. This means that we suffer from a kind of illusion of talking about the same thing, expressing the same thought, when using a sentence like this.

But one advantage of this kind of approach is the straightforward way in which the basic picture is validated: we share a reality, a reality that we represent in our thought and talk. There does seem to be something at least superficially puzzling about the idea of two people representing the same reality in the same way (i.e. with thought or speech that has the same content), but nevertheless one of them doing so correctly, the other not. If they really are representing the same thing in the same way, then either both are right or both aren’t.

Those impressed by the non-indexical relativist paradigm resist this thought. They think that two representations may have the same content, and be representations of the same object, yet one of them may be correct, the other not. Thus, my judgement that Buster Keaton is funny may be correct, while your judgement that Buster Keaton is funny may be incorrect, simply because he makes me laugh, but not you. This does not mean that we represent different realities. The representations are representations that answer to the same reality for their correctness. However, it will also depend on the representer’s individual features whether it is correct to represent the world in this way. One might ay that on this approach, representation is perspectival, for the correctness of a representation depends not only on the object represented, but also on the specific perspective of the representer.

Once we see things in this way, we will find it easy to diagnose perspectival representation in many unsuspected places. Thus we might choose to say that Napoleon’s utterance of “I am Napoleon.” in fact has the same content as the contemporary schizophrenic’s utterance of the same sentence. Both think the same thought, but only Napoleon does so truly. Or my utterance of “I am hungry.” at one time and your later utterance of the same sentence. There does seem to be a sense in which we are both saying and thinking the same thing. However, whether it is correct to represent the world in that way will depend on who represents it that way and at what time.

It may in many cases be simpler and more natural to make use of the idea of perspectival representation. For example, we do not need to postulate ad hoc departures of logical form from surface form: we can treat the sen-

tence “Buster Keaton is funny.” as having the same logical form as “Buster Keaton is American.”. But there is also a worry. If we say that different thinkers’ representations depend for their correctness on certain factors which are not shared between the speakers, then there is a good sense in which these thinkers’ representations do not answer to the same object of representation. If the correctness of thought and talk depends on a circumstance of evaluation against which we evaluate the content of the thought and talk, but different thinkers face different circumstances of evaluation, then it seems that our basic idea of a shared reality that we all face is not accurate. Different thinkers seem to face different circumstances.

In order to answer this worry, we should say that reality, the shared environment to which the thought and talk of all of us answers, is only part of the circumstances against which we evaluate contents. More precisely, it is those factors that do not vary from the circumstance of one context of utterance to the circumstance of the next. It is the invariant aspects of the circumstance of evaluation. In Kaplan’s framework, this is the world parameter. In a framework like that explored in § 6.6 when talking about future contingents, the shared reality is just the actual past, i.e. the set of worlds that overlap until the moment of the utterance context.

If we can speak of the reality that everyone shares, then, presumably, we can also speak of the reality that is shared by a smaller group of thinkers. Thus, everyone with a certain state of information will answer to the same standard of correctness when it comes to judgements of what is epistemically possible. So reality or objectivity may be relative itself: it makes sense to say that one group with converging states of information face the same reality of epistemic possibility, and that they face a different such reality from those who are in different states of information. But this does not mean that anything goes. On the contrary: once a group do share a certain group-related reality in the sense mentioned, they can approach their thought in this area in a more collaborative spirit, by assuming, for example, that if one of them is believes correctly a content of the subject matter in question, then everyone of them would be correct in believing it also. This means that these people can profitable debate the matter or benefit from the testimony of reliable sources.

Annex: Notes on Exercises

The following notes are meant to provide some help with the exercises that occur throughout the text.

Exercise 1, p. 22: Read chapter 3 of Boghossian's *Fear of Knowledge* (2006). Essay question: "What does Boghossian mean by 'fact constructivism', and do his 'three problems' show that fact constructivism is wrong?"

Note: Boghossian's three problems can be found on pp. 38–41. I suggest a word limit of 1000 words for this essay.

Exercise 2, p. 29: Read Hardwig 1985 and Blais 1987. Essay question: What is Hardwig's conclusion, and has he offered convincing arguments for it? Is the objection by Blais successful?

Note: A careful reading of Hardwig 1985 (bibliographical details in the references) will allow you to find out what conclusion Hardwig is arguing for, and what his arguments are. In assessing whether these arguments are convincing, your own judgement is called for. Blais (1987) thinks the arguments are not convincing. You are asked to assess his view as well. I suggest a word limit of 1500 words for this essay.

Exercise 3, p. 48: Read Perry 1979. Essay question: What kind of solution to the problem just mentioned here would Perry propose?

Note: The problem just mentioned was a problem for the view that "here" expresses different concepts when used by different people, and that, more specifically, "here" expresses in the mouth of Grover the concept that "where Grover is" expresses in the mouth of others. The problem was that "here" and "where Grover is" seem to express different concepts in Grover's mouth, for he might forget that he is Grover, in which case he might think concerning some place that it is here without believing that it is where Grover is (or vice versa). But we want to say that "where Grover is" expresses the same concept in everyone's mouth.

In Perry 1979, Perry considers a number of solutions to a related problem: the "Problem of the Essential Indexical". He articulates this as a problem for a certain position: "the doctrine of propositions". He considers several

ways of modifying this position, only to reject them. In the end, he makes his own proposal. In this essay question, you are asked to explain Perry's take on beliefs such as the belief that Grover would express by saying "There is immediate danger here.", and the belief he would express by saying "There is immediate danger where Grover is.". Is the propositional content of these beliefs the same on Perry's account?

I suggest a word limit of 1500 words for this essay.

Exercise 4, p. 70: if the SVs of expressions in N are objects and the SVs of expressions in S are truth-values, then what kind of function would be the SV of "kissed Mary" and of "kissed"?

Note: This is a simple exercise to test your understanding of the simple compositional rule and categorial grammar. The idea of this rule is that whenever two expressions are concatenated (i.e. whenever a new expression is created by prefixing one expression by another), the semantic value of the new complex will be the value of the semantic value of the first when we take the semantic value of the second as argument. The proposal was now that syntactically, "Mary" and "John" are in category N (names), "kissed Mary" is in S/N (the category of expressions that yield an S, i.e. sentence, when concatenated with an N) and "kissed" is in (S/N)/N (the category of expressions that yield an S/N when concatenated with an N). The exercise asks you to specify what kinds of functions the semantic values of "kissed Mary" and "kissed" would need to be if the semantic values of names are objects, and the semantic values of sentences are truth-values. It is sufficient to indicate the domain and range of the relevant functions. Thus, your answer could have the form: "The semantic value of 'kissed Mary' needs to be a function from … as arguments to … as values; and the semantic value of 'kissed' needs to be a function from … as arguments to … as values."

Exercise 5, p. 75: Actually, this suggests that we should treat these operators as complex, so that $\lambda \boldsymbol{x_1}$ results from applying λ to $\boldsymbol{x_1}$, that $\lambda \boldsymbol{x_2}$ results from applying λ to $\boldsymbol{x_2}$. Question: in which category is λ?

Note: The text says that for each variable $\boldsymbol{x_i}$, there is a lambda operator $\lambda \boldsymbol{x_i}$ in category (S/N)/S, i.e. in the category of expressions that form predicates (S/N) from sentences. The fact that there is such a lambda operator for each variable suggests that these operators are syntactically complex, i.e. that $\lambda \boldsymbol{x_1}$ is a compound from the simple expression λ and the variable $\boldsymbol{x_1}$;

that λx_2 is a compound from the λ and the variable x_2 etc. The question is simply in which category λ would be.

Exercise 6, p. 86: Articulate an appropriate semantic clause for **it is contingent whether** on the model of 15a. and 15b.

Note: On an intuitive understanding, it is contingent whether *p* just if there are possible circumstances in which *p* is the case and there are possible circumstances in which *p* fails to be the case. You are asked, simply, to provide a clause for an operator **it is contingent whether** which follows the model of 15a. and 15b.:

15a. If α is a sentence, then for all a $\in$ A and for all w $\in$ W:
SV('**possibly**(α)', a, w) = 1 iff there is a w*$\in$ W such that SV(α, a, w*) = 1.

15b. If α is a sentence, then for all a $\in$ A and for all w $\in$ W:
SV('**necessarily**(α)', a, w) = 1 iff for all w*$\in$ W, SV(α, a, w*) = 1.

In other words, you are asked to construct a clause 15* which begins as follows:

15*. If α is a sentence, then for all a $\in$ A and for all w $\in$ W:
SV('**it is contingent whether**(α)', a, w) = 1 iff …

Exercise 7, p. 111: In order to account for the fact that "Peter always frowns now." can be used to convey that Peter frowns at all times in a range somehow indicated by "now", we might construe "always" as an indexical shifting operator of some sort. Attempt this.

Note: The clause for **always** in L4 runs as follows (p. 105):

15h. **always** … if α is a sentence, then
SV('**always** (α)', c, a, p) = 1 iff for all p* $\in$ P such that W(p*) = W(p), SV(α, c, a, p*) = 1.

As explained on p. 109, according to 15h., **always** is a non-indexical fixing operator: when applied to a sentence with temporally variable content, the result is a sentrence expressing a temporally fixed content. Moreover, the way always does this is not dependent on the context. Similarly, now, as introduced on p. 105:

15d. **now** … if α is a sentence, then
SV('**now**(α)', c, a, p) = 1 iff SV(α, c, a, <T(c), W(p)>) = 1.

turns a sentence expressing a temporally variable content in a given context into a sentence that expresses a temporally fixed content at that context. However, now does this in an indexical way: how the content is fixed depends on the context of utterance.

As also explained on p. 109, this means that iterating **always** or **now** is redundant: when α expresses a temporally fixed content at some context, then '**now(**α)' and '**always**(α)'express the very same content. This does not seem to be true for the English expressions "always" and "now": the sentence "Peter always frowns now." seems to be a non-redundant application of "now" to a sentence containing "always" (or perhaps an application of "always" to a sentence containing "now"?). The question asks you to attempt to construct a semantic clause for **always** that is capable of capturing the intuitive content suggested in the question. Hint: presumably the alternative clause for **always** will need to construe it as a shifting rather than a fixing temporal operator. In other words, the content of sentences formed with **always** will still be temporally variable, so that **now** has a temporally variable content to work on non-redundantly.

Exercise 8, p. 137: What should the indexical relativist say about observation C?

Note: Observation C was the observation, on p. 94, that if Anna in a later context C48 were to have changed her aesthetic preferences in such a way that in C48 it is no longer correct for her to judge that Depp is more handsome than Pitt, then Anna could in C48 utter correctly "I was wrong." while referring to her original utterance of (5a), but it would not be correct for her to utter this with respect to an utterance of (5a*), i.e. of the indexical relativist paraphrase of (5a).

One way to approach this question would be to explore whether the proposed indexical relativist explanation of observations A and B in terms of a Gricean generalized implicature (see p. 95–6) can be extended to account for observation C.

Exercise 9, p. 145: articulate a semantic clause for a predicate modifier **by all standards**.

Note: On the previous page, an intensional sentence operators **by all standards** has been introduced:

15k. **by all st** … if α is a sentence, then SV('**by all standards** α', c, a, p) = 1 iff for all p* $\in$ P such that T(p*) = T(p) and W(p*) = W(p), SV(α, c, a, p*) = 1.

Subsequently, **for** was introduced as an expression that allows forming an intensional sentence-operator by concatenating it with a personal name:

15l. **for** … if α is a name and β is a sentence, then SV('**for** α β', c, a, p) = 1 iff for all p* $\in$ P such that S(p*) is compatible with the aesthetic preferences of SV(α, c, a, p), T(p*) = T(p) and W(p*) = W(p), SV(β, c, a, p*) = 1.

After discussion of an objection (by Hawthorne and Cappelen), according to which "for Peter", as in "The dish is tasty for Peter." should be construed as a predicate modifier rather than a sentential operator, the following replacement for 15l. was proposed:

15l.* **for** … if α is a name and β is a 1-place predicate, then SV('β **for** α', c, a, p) = {x: for all p* $\in$ P such that S(p*) is compatible with the aesthetic preferences of SV(α, c, a, p), T(p*) = T(p) and W(p*) = W(p), x $\in$ SV(β, c, a, p*)}

15l.* treats the operator formed by prefixing **for** to a name as predicate modifiers. The exercise now asks you, simply, to formulate a similar treatment of **by all standards**, which also treats it as a predicate modifier rather than a sentence operator.

Exercise 10, p. 163: Suppose that L4 is a semantics for a fragment of English, and that the correct logical form sentence corresponding to "Peter dances." is "**dances**(**Peter**)". Spell out what predictions L4 together with (K) would allow us to make about utterances of "Peter dances.".

Note: The point of this exercise is for you to apply to a concrete case what has just been said about the empirical significance of semantic theories via a bridge principle like (K) on p. 163. So you ought to look up which intension L4 assigns to "**dances**(**Peter**)", and then explain what predictions one can now make with the help of (K). Suggested word-limit: 500 words.