

Forthcoming in *Semantics: An International Handbook of Natural Language Meaning*, edited by Claudia Maienborn, Klaus von Heusinger, and Paul Portner.

66. Propositional Attitudes

Eric Swanson

University of Michigan

Department of Philosophy

2215 Angell Hall

435 South State Street

Ann Arbor, MI 48109-1003

USA

Phone: 1-734-764-6285

Fax: 1-734-763-8071

ericsw@umich.edu

September 30, 2009

Propositional Attitudes

1. Attitude ascriptions
2. The variety of attitudes
3. Selected coreference puzzles, and the analysis of ‘believes’
- 5 4. Some treatments of coreference puzzles
5. Modeling doxastic states
6. ‘Local’ presupposition satisfaction: Another puzzle of attitude ascription
7. Conclusion
8. References

10 *Verbs like ‘believes,’ ‘knows,’ ‘suspects,’ ‘hopes,’ and ‘worries’—verbs that, at the level of logical form, can take clauses as their complements—are generally taken to denote intentional attitudes borne to a proposition. For this reason they are known as propositional attitude verbs. It is difficult to construct a semantics and pragmatics adequate to the features of these verbs. Any successful theory must explain why, within the scope of an attitude ascription, substitution of coreferring terms sometimes seems to change the truth value of the ascription.*

15 *This feature of attitude ascriptions seems to entail that coreferring terms can have different semantic values; other compelling arguments seem to show that coreferring terms must have the same semantic value. After surveying other important features of propositional attitude verbs, and presenting several coreference puzzles, this article discusses conceptions of mental content intended to help resolve such puzzles. It then explores the importance of subjective uncertainty to attitude ascriptions and to formal semantics in general.*

20 *It concludes by sketching an approach to the semantics of attitude ascriptions that coheres with the standard ways of representing subjective uncertainty. This approach also unifies the treatment of coreference puzzles and the treatment of presupposition carrying expressions in attitude ascriptions.*

1. Attitude ascriptions

Attitude ascriptions include sentences like

- 25 (1) Liem hopes that Santa is coming soon.

(2) He suspects that Santa will give him a rocket.

(3) Some children worry that if Santa gets lost, the reindeer won't know how to find their house.

These sentences have the form 'NP VPs that φ ', where 'VP' denotes some intentional attitude, and (bracketing the question of its denotation) 'that φ ' is a finite clause headed by a complementizer. Attitude ascriptions come in many other forms, of course. But everyone agrees that (1)–(3) represent the sort of claim we are aiming to give a semantics for when we give a semantics for attitude ascriptions in general.

It is very common to think that any semantics for attitude ascriptions presupposes a theory of the relations that believers bear to propositions—abstract objects that represent or embody truth-evaluable intentional content. These relations are commonly called *propositional attitudes*. The nature of propositions and the nature of the relations we putatively bear to them is hotly debated. So it is helpful for certain purposes to explore the features of attitude ascriptions without making substantive presuppositions about propositions and propositional attitudes. Temporarily bracketing questions about the nature of content makes it easier to appreciate considerations from all the different fields that bear on attitude ascription: (at least) epistemology, philosophy of mind, philosophy of language, semantics, and pragmatics.

2. The variety of attitudes

The 'propositional' attitudes are a motley bunch: we can know, learn, regret, believe, imagine, fear, wish, want, pretend, suppose, surmise, suspect, predict, speculate, doubt, prove, disprove, infer, ..., that φ . Following Huddleston & Pullum (2002), I will say that a verb is *factive* just in case it carries the presupposition that its complement clause is true, and that a verb is *entailing* just in case in the positive declarative it entails the truth of its complement (1008–1009). (Throughout I have in mind pragmatic presupposition, in Stalnaker's sense (1974); cf. article 102 *Presupposition*.) These categories cut across each other, although it's sometimes overlooked that verbs can be factive without being entailing, and vice versa.

	ENTAILING	NON-ENTAILING
FACTIVE	find out, know, remember	confess, regret, resent
NON-FACTIVE	discover, establish, prove	believe, infer, suspect

TABLE 1: CATEGORIES OF ATTITUDE VERBS

Verbs like ‘believe’ are neither factive nor entailing; I leave the production of examples to the reader.

For entailing verbs that differ primarily in their factivity—that is, in whether they carry the presuppo-

55 sition that their complements are true—contrast

(4) NASA just found out that there’s life on the moon!

(5) NASA just discovered that there’s life on the moon!

Utterances of (4) generally presuppose that there’s life on the moon and utterances of (5) generally

don’t. But both sentences entail that there is life on the moon: you can neither find out nor discover

60 that φ unless it’s actually the case that φ . So ‘find out’ is factive and entailing, whereas ‘discover’ is not

factive, but is entailing. And here are some examples that show that ‘confess’ and ‘resent’ generally

presuppose the truth of their complements without entailing it:

(6) She confessed to taking the money, but later recanted. It turned out that she had been trying
to cover up a friend’s mistake.

65 (7) I resented him for leaving all the work to me, until I learned how much he had already done.

Finally, some verbs with an attitudinal component to their meaning, like ‘disprove,’ ‘refute,’ ‘lie,’ and ‘fib,’

entail the falsity of their clausal complements. These verbs more naturally take noun phrases, but in

negative or contrastive environments they take clausal complements fairly easily: “It wasn’t disproved

that the earth was flat until ...”; “His investigators claim UnitedHealth manipulated data and even lied

70 that its reimbursement rates were based on national research” (Associated Press, “Cuomo to sue major

health insurers,” February 13, 2008).

A successful theory of attitude ascriptions must explain why, within the scope of an attitude ascription, substituting for one term another term that corefers with it sometimes seems to change the truth value of the ascription. I discuss this phenomenon at length in sections 3, 4, and 6. A successful theory must also explain how the facts that some attitude ascriptions seem to express entail that believers stand in intentional relations—relations of ‘aboutness’—to nonexistent objects. As Brentano (1874, 50) seminally put it: “Every mental phenomenon is characterized by what the scholastics of the Middle Ages called the intentional ...inexistence (*Inexistenz*) of an object (*Gegenstand*), and what we could call ...the reference to a content, a direction upon an object (by which we are not to understand a reality in this case), or an immanent objectivity.” These features of attitude ascriptions are more than a little mysterious, and motivate a vast amount of work (cf. article 3 *Sinn and Bedeutung* and article 4 *Reference*).

And they are distinctive enough that other expressions that exhibit these features are sometimes taken to be covert attitude ascriptions. For example, Quine (1956) analyzes ‘NP1 is hunting NP2’ as, roughly, ‘NP1 strives to make it the case that NP1 finds NP2’. This lets him apply explanations of substitution failure and ‘intentional inexistence’ in the complements of clausal intentional verbs to intentional verbs with complements that are not overtly clausal. It also lets him distinguish between two readings of sentences like ‘Ernest is hunting lions’—one that relates Ernest to a particular lion and another that leaves open the question whether there is a particular lion that he is hunting:

(8) $\exists x(x \text{ is a lion and Ernest strives that Ernest finds } x)$

(9) Ernest strives that $\exists x(x \text{ is a lion and Ernest finds } x)$

Considerable ingenuity has been applied to the project of analyzing all intentional attitude ascriptions as being, fundamentally, propositional attitude ascriptions. (See especially den Dikken, Larson & Ludlow 1996 and 1997. McCawley 1974, Karttunen 1976, and Ross 1976 marshal syntactic considerations in favor of the view.) Such efforts have hefty burdens to discharge, since they are not successful unless every apparently non-clausal intentional verb can be plausibly paraphrased without using any such verbs. Whether or not ‘I want x ’ is plausibly paraphrased as ‘I desire that I have x ’, it’s hard to

find clausal paraphrases of ‘ignore,’ ‘love,’ ‘insult,’ and many other non-clausal intentional verbs. (See Montague 2007 for discussion of potential clausal semantics for ‘love.’)

100 Verbs used to ascribe desires, like ‘want,’ ‘hope,’ and ‘wish,’ raise problems that it will be instructive to linger on. On an extremely simple approach to attitude ascriptions, which I will call the *naïve possible worlds approach*, the complement of an attitude ascription denotes a set of possible worlds, or *possible worlds proposition*. (According to textbook intensional semantics, clauses denote possible worlds propositions, so this approach looks quite natural from that theoretical point of view.) On one
105 version of this approach, ‘Liem hopes that φ ’ is true just in case Liem’s most preferred worlds are all worlds in which it is true that φ . Similarly, ‘Liem believes that φ ’ is true just in case all of the possible worlds compatible with Liem’s beliefs—that is, all of the possible worlds in which all Liem’s beliefs in the world of evaluation are true—are worlds in which it is true that φ . This semantics for ‘hopes’ founders on the fact that our preferences are generally not closed under entailment. From the fact that

110 (10) John hopes that you leave later.

it does not follow, intuitively, that

(11) John hopes that you leave.

But all the worlds in which you leave later are worlds in which you leave. So on the naïve possible worlds semantics we are considering, (10) wrongly entails (11) (cf. Prior 1958, Forrester 1984).

115 It would be overhasty to conclude that we should not model the contents of our preferences using possible worlds propositions. Following Stalnaker (1984, 89–90), Heim (1992) details a theory according to which (11) means (roughly) “that John thinks that if you leave he will be in a more desirable world than if you don’t leave” (193), and (10) means that Jones thinks that if you leave later he will be in a more desirable world than if you don’t leave later. (Cf. Hansson 1969, van Fraassen 1972 and
120 1973, Lewis 1973, article 64 *Modality*, and article 65 *Conditionals*.) More precisely, on Heim’s semantics, (10) means that for any world w compatible with what John believes, John prefers every world in which you leave later that is maximally similar to w (among worlds in which you leave later) to any

world in which you do not leave later that is maximally similar to w (among worlds in which you do not leave later). On this semantics, then, the truth of (10) does not entail that John has the preferences that would make (11) true. But as far as this particular problem is concerned, we may still model the contents of our preferences using possible worlds propositions, with the help of an ordering source (that itself may be modeled using possible worlds propositions: see Lewis 1981a). As Heim argues, this kind of approach also helps explain a host of other features of desire reports (195–206). The prima facie failure of a ‘possible worlds approach’ here spurred refinements resulting in a more explanatory overall theory than we might have come to otherwise. It is more complicated to refine our models of belief states enough to avoid the result that what is truly ascribed by ‘believes,’ in a context, is closed under entailment; I take up that project in section 5.

3. Selected coreference puzzles, and the analysis of ‘believes’

‘Believes’ is one among many intentional verbs, as we have seen. But it is the uncontested central case in the literature on attitude ascriptions. By restricting our attention to it we can focus on some of the most challenging features of attitude ascriptions. It’s quite plausible that a successful treatment of those features as they are exemplified by ‘believes’ could be applied to verbs that express relevantly similar attitudes.

In his ground-breaking “On Sense and Reference,” Frege introduces a puzzle about belief—and so, to some extent, about ‘believes’—that has to do with the cognitive significance of identity statements (cf. article 3 *Sinn and Bedeutung*).

If the sign ‘ a ’ is distinguished from the sign ‘ b ’ only as an object (here, by means of its shape), not as a sign (i.e. not by the manner in which it designates something), the cognitive value of $a = a$ becomes essentially equal to that of $a = b$, provided $a = b$ is true. (Frege 1892, 152)

But, as Frege observes, the two sentences differ in their cognitive significance, in the sense that

$a = a$ holds *a priori* and, according to Kant, is to be labelled analytic, while statements

of the form $a = b$ often contain very valuable extensions of our knowledge and cannot always be established *a priori*. (151)

150 The substitution of coreferring terms in attitude ascriptions elicits similar phenomena: it's easy to realize that $a = a$, but to discover that $a = b$ can be a hard-won achievement.

Frege puts the problem in terms of phonologically and orthographically distinct signs, but this isn't essential to coreference puzzles. Kripke (1979, 902) asks us to consider Pierre, who has very different 'London' and 'Londres' beliefs, and Peter, who believes that there are two people named
155 'Paderewski'—one a famous pianist with considerable musical talent and the other a Polish nationalist leader with no musical talent. We can imagine contexts in which it's appropriate to say

(12) Peter believes that Paderewski had musical talent.

and we can also imagine contexts in which it's appropriate to say

(13) Peter believes that Paderewski had no musical talent.

160 even if we believe that Peter is not guilty of any logical errors. I think that we can even imagine contexts in which *either* (12) or (13) can be used appropriately to target beliefs of Peter's that correlate with his 'famous pianist' and 'Polish politician' beliefs. (This isn't to say, of course, that we could appropriately use one right after the other; using one changes the context to make uses of the other inappropriate.) This leads Kripke to suggest that in this case, and in others like it, "our normal practices interpreta-
165 tion and attribution of belief are subjected to the greatest possible strain, perhaps even to the point of breakdown. So is the notion of the *content* of someone's assertion, the *proposition* it expresses" (906).

It's not even essential to coreference puzzles that the coreferring terms be *names*. Adapting a case due to Mark Richard: Imagine that Al, talking to Betty on the telephone, sees a woman in a distant phone booth about to be hit by a runaway steamroller, and wishes he could warn her of the danger.
170 But Al does not realize that Betty is the woman he sees in the phone booth. Then it seems Al could say truly to Betty

(14) I believe I can inform you of her danger over the telephone.

But Al could not say truly

(15) I believe I can inform her of her danger over the telephone. (Richard 1983, 439–440)

175 Examples like these suggest that the substitution of coreferring proper names, pronouns, and demonstratives can make a difference to the truth conditions of an attitude ascription. Given other plausible assumptions, these examples seem to show that coreferring proper names, pronouns, and demonstratives can have different semantic values (cf. article 4 *Reference*, article 41 *Pronouns*, and article 101 *Deixis and demonstratives*).

180 4. Some treatments of coreference puzzles

Indeed, semantic theories with this commitment were once extremely common. Many philosophers (including Frege, Russell, Carnap, and Searle) held that there could be differences between coreferential terms that made for semantic differences between attitude ascriptions of which they were a part. Some such semantic treatments, like Frege's, posit a systematic shift in meaning whereby proper names have a different semantic value in the context of an attitude ascription than they otherwise would. For Frege, the "customary" semantic value of a finite clause is a truth value, but when a clause is the complement of an attitude ascription its semantic value is instead a "thought" (160). (I translate '*Bedeutung*' as 'semantic value'; this is a bit anachronistic, since it's now common to presuppose that semantic values are intensional, although *Bedeutungen* are not.) Frege attributes a wide range of features to sense, listed below:

1. "By employing a sign we express its sense and designate its *Bedeutung* [referent]" (156).
2. Senses are compositional: the sense of a sentence is determined by the senses of its parts, and their arrangement. (156).
3. "Every grammatically well-formed expression figuring as a proper name always has a sense,"
195 whether or not it has a referent (153).

4. The thought expressed by “Odysseus was set ashore at Ithaca while sound asleep” “remains the same whether ‘Odysseus’ has a *Bedeutung* [referent] or not” (157).
5. The sense of a sign “contains” a “mode of presentation” of the object designated by the sign (152).
6. A sign’s sense “may be the common property of many people, and so is not a part or a mode of the individual mind.” This supposed to help explain how it is that “a common store of thoughts ...is transmitted from one generation to another” (154).
7. The sense of a proper name “is grasped by everybody who is sufficiently familiar with the language or totality of designations to which it belongs” (153).

It’s not uncharitable, I think, to think of Frege as saying that the phenomena associated with attitude ascriptions make it plausible that *something* has these features, and that by ‘sense’ he means whatever it is that in fact *does* have these features (cf. article 3 *Sinn and Bedeutung*).

Some of the features on their own are mysterious—if senses are not “a part or a mode of the individual mind” then what kind of thing *are* they, and how do we have access to them?—and there’s considerable tension between items on the list, as well. For example, it’s hard to see why linguistic competence ensures “grasp” of an object’s modes of presentation. And it’s hard to see how whether Odysseus existed could fail to make a difference to the modes of presentation associated with ‘Odysseus.’ Contemporary neo-Fregeans thus usually abandon or modify one or more of these desiderata. But Frege’s enumeration of the features of sense helps make it clearer what an *ideal* theory of attitude ascriptions might aim for. And Frege’s decision not to complicate the semantics of names that occur outside attitude ascriptions is also instructive. The puzzles of coreferring terms in attitude ascriptions, like so many in philosophy of language, are rooted in ignorance: what body or bodies of information make the difference between being ignorant of the fact that $a = b$ and knowing that $a = b$? So it’s entirely unsurprising that the semantics of attitude ascriptions can be sensitive to what the ascriber does and doesn’t know. But it’s another matter altogether, as Frege must have seen, to think that facts about what is and isn’t known have any bearing on the semantics of ‘simple’ sentences that do not make attitude ascriptions.

Russell (1919, 76), by contrast, treats ordinary proper names in a semantically uniform way: wherever they occur, they are disguised definite descriptions. On any plausible analysis, definite descriptions enter into scope relations with quantifiers and other operators, including the operators whereby
225 we make attitude ascriptions. For example, uttered today

(16) The president of the United States will always be male.

can be interpreted as making the uncontroversial claim that Barack Obama will always be male, or as making the very controversial claim that for all future times t , the president of the United States at t is male at t (Heim 1991, 7). Similarly, Russell (1905, 44-45) observes that when we say

230 (17) George IV wondered whether Scott was the author of *Waverley*.

...we normally mean “George IV wished to know whether one and only one man wrote *Waverley* and Scott was that man”; but we *may* also mean: “One and only one man wrote *Waverly*, and George IV wished to know whether Scott was that man.”...[which might also] be expressed by “George IV wished to know, concerning the man who in fact wrote
235 *Waverly*, whether he was Scott.”

Because coreferential but distinct proper names might well ‘disguise’ different definite descriptions, Russell has a straightforward answer to the question how “Al knows that $a = a$ ” and “Al knows that $a = b$ ” can come apart in truth value. These attitude ascriptions are, on Russell’s view, not essentially different from the attitude ascriptions

240 (18) Al knows that the F is the F .

and

(19) Al knows that the F is the G .

Even if the *F* is the *G*, (18) and (19) clearly can have different truth values. And Russell's analysis gives a lucid account how believers can apparently stand in intentional relations to nonexistent objects denoted by proper names: they simply falsely believe that there is something that satisfies the definite description associated with the proper name. As I mentioned earlier, on any plausible analysis of definite descriptions they can enter into scope relations. On Russell's own analysis, the scope facts are a result of the quantificational structure that definite descriptions contribute to logical form. But a sufficiently rich intensional semantics can also capture the scope facts, without treating definite descriptions as quantificational (Heim 1991, 19–22; cf. article 4 *Reference*; article 42 *Definiteness and indefiniteness*; and article 44 *Quantifiers*).

The development of modal logic in the middle of the twentieth century brought in its wake an assault on 'descriptivism'—broadly construed to include theories of proper names on which they were associated with either Fregean senses or Russellian descriptions—that found powerful and synoptic expression in Kripke's *Naming and Necessity* (1980). There Kripke forcefully argued that proper names were rigid designators, where a rigid designator is an expression that designates the same object in all possible worlds. (More precisely: an expression is rigid just in case its intension is constant.) One crucial datum was that "it's a contingent fact that Aristotle ever did *any* of the things commonly attributed to him today" (75). This was inconsistent with many descriptivist views of the time. Searle (1958, 172), for example, claimed that "it is a necessary fact that Aristotle has the ...inclusive disjunction of properties commonly attributed to him: any individual not having at least some of these properties could not be Aristotle" (cf. article 64, *Modality*).

Until attention was focused on proper names' modal profile, it was common to presuppose that they had variable intensions in the sense that the world of evaluation—a semantic parameter shiftable by modal expressions—could make a difference to what a name contributed to the determination of the truth value of a sentence in which it occurred. This presupposition is arguably latent, for example, in the view that (1) the identity of Hesperus and Phosphorus is knowable only a posteriori and (2) all 'a posteriori truths' are contingent (Kripke 1980, 101–105). And broadly descriptivist semantics in effect exploited the putative variability of the intensions of proper names to explain coreference puzzles. But Kripke and others elicited intuitions about the modal profile of proper names that strongly suggested

that their intensions are actually *constant* over (at least) the metaphysically possible worlds. These intuitions thus make obvious trouble for views that explain coreference puzzles by appeal to variable intensions. Many also take intuitions about rigidity to make trouble for views that model the content of attitude ascriptions purely intensionally, with a set of possible worlds: together with the fact that Hesperus is Phosphorus, the rigidity of proper names like ‘Hesperus’ and ‘Phosphorus’ entails that the possible worlds in which Hesperus is *F* are exactly the possible worlds in which Phosphorus is *F*. But then the possible worlds model of the content of the belief that Hesperus is *F* does not differ from the possible worlds model of the content of the belief that Phosphorus is *F*. As Soames influentially put it, “one can always find psychologically non-equivalent sentences which are true in the same circumstances, and, hence, [would be] assigned the same content” if belief were modeled using sets of possible worlds (1987, 231; cf. Scheffler 1955, 41–42).

Together these considerations encouraged the refinement of old treatments of attitude ascription, and the development of new ones. Some philosophers try to accommodate intuitions about rigidity with carefully tailored descriptivist semantics (see, e.g., the discussion in Stanley 1997). Salmon (1986), Soames (1987, 2002), and Thau (2002), among many others, explain coreference puzzles by appealing to implicatures or other pragmatic effects. In both these camps the apparent problems with characterizing mental content in terms of sets of possible worlds encourage many philosophers to defend conceptions of content on which it has more structure than a set of possible worlds. Some argue that content must be individuated so finely that its putative structure quite closely parallels the syntactic structure of sentences (Larson & Ludlow 1993, King 1996). King (2007, 57) goes so far as to hold that “the structure of a proposition will be identical to the syntactic structure of the sentence expressing it.” One point adduced in favor of such views is that they allow us to “see how” it is possible to “believe a proposition while failing to believe another necessarily equivalent to it” (57): “‘1 = 2’ and ‘2 = 1’ express different propositions in virtue of having their constituents differently combined” (97). Some take structured propositions to help solve this important aspect of the problem of logical omniscience; others think solving the problem requires a fuller characterization of the putative difference in the contents expressed by ‘1 = 2’ and ‘2 = 1’ (for more, see Stalnaker 1991 and 1998).

For the most part approaches like these were motivated by features that noun phrases exhibit in attitude ascriptions. Another important line of thought, initially motivated by the behavior of noun phrases in discourse (Karttunen 1976, Kamp 1981, and Heim 1982), posited a different kind of structure in mental content. Discourse phenomena motivate, in Kamp's formulation, "a theory of meaning and context dependent interpretation of English ...that goes beyond what sets of possible worlds are able to reveal" as representers of context (1988, 158). "Contextual structure" (165) cannot be represented with a set of possible worlds, according to Kamp; a (highly structured) Discourse Representation Structure or DRS can represent it. Kamp uses a famous minimal pair due to Barbara Partee to press his point:

- (20) a. Exactly one of the ten balls is not in the bag.
b. It is under the sofa.

- (21) a. Exactly nine of the ten balls are in the bag.
b. It is under the sofa.

The missing ball is readily available as the referent of 'it' in (20-b); it is relatively hard to read 'it' in (21-b) as referring to that ball. Kamp writes: "if propositions are sets of possible worlds, the two assertions [(20-a) and (21-a)] express the same proposition. ...So the resulting contexts [after those assertions] ...will be equal to each other. We must conclude that no difference can be predicted if contexts are identified with sets of possible worlds" (158). By contrast the DRSs associated with (20-a) and (21-a) differ: (20-a) introduces a discourse referent, where (21-a) does not, which "links *it* to the subject of the first sentence" (162, cf. article 38 *Discourse Representation Theory*). Asher, among others, put such rich representations of context and context change potential to work in attitude ascriptions: "DRSs ...take on a new role as characterizations of the objects of mental states" (1986, 134). This approach aspires to unify the representation of context and the representation of mental content. To the extent to which it successfully discharges this ambition it has a prima facie advantage over approaches that say nothing about discourse relations and the behavior of noun phrases in discourse.

5. Modeling doxastic states

There is much to be said in defense of modeling the contents of attitude ascriptions (and contexts) using sets of possible worlds (see especially Lewis 1981b and Stalnaker 1984, 1991, 1996, 1998, and 1999b).
325 Rather than rehearsing that dialectic I want to explore a source of pressure against structured conceptions of content: the representation of credence (cf. Chalmers 2009).

Credence is not an ‘on/off’ attitude: it comes in degrees. For example, if I believe that rain is likelier than snow tonight, I lend more credence to the proposition that it will rain tonight than I do to the
330 proposition that it will snow tonight. The language of attitude ascription reflects the degreed nature of credence. Consider (22):

(22) Al believes it will probably rain.

This sentence *prima facie* ascribes to Al moderate but not full credence in the proposition that it will rain, not full credence in the proposition that it will probably rain: we would not put “probable rain” on
335 the list of things Al is sure about. And some attitude verbs pretty clearly ascribe less than full credence even with ‘unhedged’ complements:

(23) As the clouds grew darker, he slowly became more confident that it would rain.

(24) Betty surmised that it would rain.

(25) Clara suspects and Doug doubts that it will rain.

340 Modeling doxastic states with probability spaces makes it possible to give elegant semantic entries for such verbs. For example, it’s trivial to give entries for ‘suspects’ and ‘doubts’ with the result that (25), for example, is true just in case Clara gives credence above some threshold to the proposition that it will rain, and Doug gives credence below some other threshold to *the very same* proposition—the proposition denoted in the context of utterance by ‘it will rain’. To be sure, one could try to develop a
345 view on which (25) ascribes full beliefs with distinct contents—a ‘suspecting’ full belief to Clara, and a ‘doubting’ full belief to Doug. But such a view would have the unenviable burden of explaining how

it can be that, despite the difference in the contents of their putative full beliefs, Clara suspects what Doug doubts.

Probability spaces are especially attractive and popular tools for representing credence because of their expressive power and accessibility. By a *probability space* I mean a triple $\langle W, \mathcal{F}, \mu \rangle$ such that:

1. \mathcal{F} is a Boolean algebra over W (where a *Boolean algebra* over a set W is a set of subsets of W that includes W itself and is closed under complementation and union);
2. $\mu(\cdot)$ is a function from $\mathcal{F} \rightarrow [0, 1]$;
3. $\mu(W) = 1$;
4. If M and N are disjoint elements of \mathcal{F} , then $\mu(M \cup N) = \mu(M) + \mu(N)$.

For simplicity I assume that W is finite. W is generally a set of possible worlds, making \mathcal{F} a set of sets of worlds—a set of possible worlds propositions. $\mu(\cdot)$ assigns values to the elements of \mathcal{F} . \mathcal{F} may be as sparse as $\{\emptyset, W\}$, it may be as rich as the power set of W , and it also may be any Boolean algebra over W in between. I will occasionally refer to μ , for a given probability space, as its *probability measure* (Halpern 2003, 15–16).

Ordinary talk about subjective uncertainty is often overtly probabilistic: weather forecasters, bookies, and the passengers on the Clapham omnibus all hedge their predictions in ways that conform to the probability axioms. (If they disregarded the axioms, then it would be easy to imagine a forecast that said that frost and no frost were both 70% likely, and a bookie who gave two to one odds for all four horses in a race.) It's also routine to compare credences in ways that are well modeled by probability spaces. I may rightly and justifiably tell you that it's twice as likely to rain as it is to snow, for example, although I do not have an opinion about the likelihood of precipitation. This can be modeled with a constraint on probability spaces that favors lending twice as much credence to rain as to snow. Moreover, belief in indicative conditionals can and arguably should be analyzed as a kind of comparison of probabilities which, thanks to Lewis's celebrated 'triviality results,' provably cannot be reduced to the probability of a proposition. (See especially Lewis 1976 and Gibbard 1981; cf. article 65 *Conditionals*.) The conditional probability of C given A , relative to a probability measure $\mu(\cdot)$, is $\frac{\mu(A \cap C)}{\mu(A)}$. So for the conditional probability of C given A to be high, it must be the case that $\mu(A \cap C) \approx \mu(A)$.

Epistemically hedged claims have distinctive effects on context and on conversational participants' belief states, and are subject to distinctive norms governing their use. It is difficult (if not impossible) to find propositions that give the meaning of epistemically hedged claims. One way to see this is to notice that it is difficult (if not impossible) to find propositions the use of which is governed by the same norms as those that govern epistemically hedged claims. (See MacFarlane 2003, Egan, Hawthorne & Weatherson 2005, Swanson 2006, 2008, 2009, and Yalcin 2007 for detailed work in this vein.) Another way to see this is to notice that if epistemically hedged claims *did* express propositions, then there would have to be a function from every way each proposition can be hedged and each proposition to propositions. Intuitively, this function would take a *partial* belief—a particular credence in $[0, 1]$ borne to a particular proposition—and yield a proposition that a believer fully believes just in case she has that *partial* belief. We shouldn't take the existence of a function with the appropriate features on faith, and to my knowledge no one has tried to argue that one exists. (Kratzer 1991 develops a treatment of a few 'graded modals,' but I do not see how to generalize her approach very far; cf. article 64 *Modality*.) Showing that there *is* such a function, incidentally, would be a discovery of first importance, since it would show that in representing doxastic uncertainty, at least, the probability space is a mere *façon de parler* that is reducible without loss of expressive power to a set of full beliefs.

The attractive features of probability spaces come with significant prima facie drawbacks, however. In particular, in modeling believers with probability spaces we impute to them a kind of logical omniscience. First, the necessary proposition—thought of in this setting as the proposition true in all the elements of W —is ruled a certainty for any believer. Second, the probability axioms constrain degrees of uncertainty in several substantive ways. For example:

- A believer's credence in the proposition that φ is a function of her credence in the proposition that $\neg\varphi$ and vice versa: $\mu(A) = 1 - \mu(\bar{A})$;
- A believer's credence in the proposition that φ or ψ is a function of her credence in the propositions that φ , that ψ , and that φ and ψ ;
- A believer's credence in an entailment of some proposition must be at least as great as her credence in the entailer.

All this appears to have the problematic result that, if doxastic states are accurately modeled by probability spaces, believers believe all the entailments of the things they believe, are certain of the necessary proposition, cannot lend different credences to propositions true in exactly the same possible worlds, and so on. These are analogues of already mentioned problems with non-probabilistic representations of belief using possible worlds propositions (cf. Soames 1987). That is little succor—they are serious problems nevertheless. But I think we should try to solve (or at least mitigate) these problems rather than simply abandoning the probabilistic framework. By way of illustration I now sketch one example of a solution to a prima facie problem for probabilism. The problem can be seen as analogous to closure under entailment—discussed in Section 2—but it is for belief instead of desire.

Because probability measures simply need to be defined over the elements of *some* Boolean algebra over W , we can use probability spaces to model doxastic states without assuming that believers are opinionated about every proposition. This feature of probability spaces allows us to use distinct but closely related probability spaces to model ‘overlooked’ and ‘seen’ possibilities, thereby affording a treatment of overlooked inferences. One probability space is defined over both those possibilities the believer overlooks and those she sees, measuring her credences with respect to all those possibilities. I call this her *fine credal space*. The other probability space is defined only over those possibilities she sees, representing (for any normal person) a proper subset of the credences represented by her fine credal space. This space characterizes all her credences *except* those borne to any possibilities she overlooks; I call this her *coarse credal space*. The domain of the coarse credal function is a subset of the domain of the fine credal function. Given a probability space P and any subalgebra \mathcal{S} of the propositions measured by P , it’s easy to show that a function that is defined on exactly the propositions in \mathcal{S} and that agrees with P on their values must be a probability measure. So the probability measure of the coarse credal space will agree with the probability measure of the fine credal space with respect to the values assigned to any proposition that is measured by both spaces—the propositions in the algebra of the coarse credal space. A coarse credal space defined accommodates a fine credal space just as a map accommodates an overlay: a fine credal space might add information about latent dispositions, for example, without conflicting with the ‘seen credences’ represented by the coarse space.

The assumption that the set of seen possibilities is an algebra makes that set closed under Boolean

operations. This imposes constraints on the work that can be done by the formalism: it does not
430 help represent a believer who sees the possibility that φ and sees the possibility that ψ but overlooks a
possibility yielded by any Boolean operation on the proposition that φ and the proposition that ψ . For
example, the formalism doesn't help us represent such a believer if she overlooks the possibility that
 $\neg\varphi$, or overlooks the possibility that $\varphi \vee \psi$, or overlooks the possibility that $\varphi \wedge \psi$, or overlooks the
possibility that $\varphi \vee \neg\varphi$, or But for many cases I do not think that this limitation of the framework is
435 implausible or unwelcome. For example, although I crack eggs with one hand, it wasn't until I reflected
on how I do it that I came to see the possibility that the right way to crack an egg is with the large end
in your palm. (And this despite the fact that I always picked up eggs that way before I realized that I
had such a disposition.) But once I saw *that* possibility, I ipso facto saw the possibility that the right
way to crack an egg is *not* with the large end in your palm. It's helpful to think of the closure properties
440 of the formalism in this way: each seen possibility imposes a boundary on W , in such a way that any
proposition whose boundaries can be defined purely in terms of the boundaries laid down by seen
possibilities is itself a seen possibility. To see a possibility, in this sense, is to see a way of distinguishing
between possible worlds.

This framework lets us represent believers without imputing full logical omniscience to them, at
445 least in that we can model believers who overlook (certain) entailments and entailers of possibilities
they see. For example, suppose that T entails U , and that U entails V . Suppose also that our believer
sees the possibilities represented by S , T , and V but overlooks the possibility represented by U . Then
the coarse credal space will measure any subset of W that carves W solely along black lines in FIGURE
1, but will omit those subsets that carve along any gray line. The fine credal space may measure every
450 subset of W . The coarse space thus can leave unseen certain entailers and entailments of seen possi-
bilities. This is important because I may see the possibility that my partner castles, for example, while
overlooking the possibility that she castles or moves *en passant*. This same formal device reconciles
the folk conception of belief, and natural semantics for 'believes,' with closure under entailment: we
can say that our fine-grained commitments are closed under entailment, although often we do not
455 *see* all those commitments, and we can say that we *believe* that φ only if we see the possibility that φ .
Beyond the closure properties already discussed, the framework puts no unusual constraints on the

	T	$\sim T$	U	$\sim U$	V	$\sim V$
S						
$\sim S$						

FIGURE 1: COARSE AND FINE CREDAL SPACES

norms that govern the relationships between overall doxastic states and the possibilities a believer sees and overlooks. Moreover, it allows us to precisely characterize one doxastic change induced by ‘might’ statements, as I have argued elsewhere (2006, 2009; cf. Yalcin 2009): they themselves are often used
 460 to ‘raise’ possibilities, making overlooked possibilities seen without committing the speaker to much else.

To be sure, this formal apparatus does not let us represent believers who ‘overlook the necessary proposition,’ if such there be. But the formal tools for the representation of credence will doubtless be refined over time. The prima facie problems with our current tools don’t give us good reason, on their
 465 own, to think that ‘probably φ ’ expresses a proposition, or that ‘surmise’ and ‘doubt’ express sui generis intentional attitudes. Of course, there may be other reasons to try to defend such hypotheses. But this section has provided some reasons not to depart too quickly from the ‘coarse’ individuation conditions for the contents of attitudes that make it possible to represent credence using the tools of probability theory. What of the considerations discussed in Sections 3 and 4, which suggested that content must be
 470 individuated finely? The next section argues that other phenomena exemplified by attitude ascriptions motivate an approach that solves coreference puzzles compatibly with coarse individuation of content.

6. ‘Local’ presupposition satisfaction: Another puzzle of attitude ascription

Suppose that Ken is blindfolded, and that he is trying to guess who is speaking. We can tell from Ken’s guesses that he believes that Louise has spoken once. But we also know that Louise has not spo-
 475 ken—Ken mistakenly thought that someone who sounds like Louise was Louise. That person speaks again, and I say to you

(26) Ken believes that Louise has spoken again.

(26) plainly does not exhibit presupposition failure in this context. But it is not common ground between us that Louise has spoken, and it does not become common ground between us that Louise has
480 spoken. Moreover, (26) *would* exhibit presupposition failure if it weren't common ground between us that Ken thinks Louise has already spoken once. This suggests that, even when embedded in the 'that' clause of a belief ascription, 'again' carries presuppositions—presuppositions that in the conversation as described are satisfied by what we take to be Ken's belief state. The example shows that these presuppositions need not be satisfied by the conversational participants' belief states or the conversational
485 common ground (cf. article 102 *Presupposition*).

I want to give a couple more examples to show that this phenomenon is not overly exotic. (See also Stalnaker 1988, 157–158.) Suppose we believe and presuppose that there are no spies at the party. But it's also common ground between us that Hob believes there are several. The people that Hob thinks are spies leave, and I say to you

490 (27) Hob believes that every spy has left.

(27) does not exhibit presupposition failure in this context. But we might expect that it would, because in simple sentences 'every spy' carries the presupposition that it has a non-empty domain, and it's common ground that it has an empty domain. Fortunately, in the conversation described, the presuppositions carried by 'every spy' are satisfied by what we take to be Hob's belief state.

495 Or suppose we believe and presuppose that Sue has never smoked, but it's also common ground between us that Tom is convinced that Sue does smoke. Then (28) will not exhibit presupposition failure, even though (29) would.

(28) Tom believes that Sue has quit smoking.

(29) Sue has quit smoking.

500 In (28), the presuppositions carried by ‘quit’ can be satisfied by what we take to be Tom’s belief state; in (29) they would have to be satisfied by the conversational common ground.

It’s easy to create more examples like these:

1. Take an expression ‘ α ’ that in simple sentences generally carries the presupposition that ψ .
2. Give an example of a conversation in which it is common ground that $\neg\psi$.
- 505 3. Consider a non-negated belief ascription that includes ‘ α ’ in its ‘that’ clause, as used in that conversation.
4. Notice that the belief ascription carries the presupposition that the ascriber believes that ψ .

(Karttunen (1973a, 1973b, 1974) influentially claims that ‘A believes that φ ’ *always* presupposes that A believes that ψ , for any presupposition ‘ ψ ’ normally carried by ‘ φ ’; see Heim 1992 for a recent devel-
510 opment of the view. Geurts 1998 offers a battery of arguments against Karttunen’s generalization.)

To reiterate, in such examples we have a presupposition that cannot be satisfied by the conversational common ground. It *would* be satisfied by what the conversational participants take to be the ascriber’s belief state, for purposes of conversation. And in fact, and very broadly speaking, there is a sense in which it *is* so satisfied. I follow Geurts (1998, 584–585) in classifying this phenomenon as a
515 kind of local accommodation. To say that an expression is *locally accommodated* in this sense is just to say that some or all of its presuppositions are satisfied by something other than the ‘basic’ or ‘global’ conversational context (cf. Heim 1983, 254–255).

It is important to note that locally accommodated expressions need not be noun phrases, as we have seen with ‘again’ and ‘quit’. (Karttunen 1968a, 1968b, and 1976 are arguably the first serious dis-
520 cussions of local accommodation of noun phrases; the bulk of the literature since then has followed Karttunen in this focus.) My aim is to offer a theory that unifies the behavior of noun phrases and the behavior of other kinds of expressions in attitude ascriptions. Failure to do this, I fear, might tempt us to a theory that looks explanatorily adequate for noun phrases but that turns out, in the end, to be redundant. The guiding ideas of the approach I sketch here are (1) that proper names, demonstratives,
525 and the like are presupposition-carrying expressions, (2) that they exhibit distinctive behavior as a re-

sult, and (3) that they thus warrant the treatment we would give to any other presupposition-carrying expression. On this way of approaching attitude ascriptions, the behavior of presupposition-carrying expressions in attitude ascriptions demands a treatment that brings an explanation of coreference puzzles in its wake. By focusing one or two levels higher in the taxonomy of linguistic phenomena than is common in work on attitude ascription—at the level of presupposition carriers rather than the level
530 of proper names or of definite noun phrases—this kind of approach aspires to unify some otherwise seemingly disparate features of attitude ascriptions. (For a view that is in some respects similar to this one, see Stalnaker 1988. Our approaches diverge in several places, and my approach fills in many details where Stalnaker’s is neutral, but I won’t catalog the differences here.)

535 It’s not immediately obvious how to treat local accommodation in attitude ascriptions. We could try saying that the complement of the attitude ascription is interpreted relative to a single context that is distinct from the conversational context. Or we could say that the complement is interpreted relative to multiple contexts, at least one of which is distinct from the conversational context. On the one-context approach, although the whole sentence (28) is interpreted relative to two contexts, the
540 complement clause “that Sue has quit smoking” is interpreted relative to a *single* context:

(30) [Tom believes]^{c1} [that Sue has quit smoking]^{c2}

Heim (1992) takes this kind of approach. On the multiple-context approach, by contrast, the complement clauses of attitude ascriptions can be interpreted relative to, say, the “basic [global] context” and the “derived [local] context,” which is the “set of all possible situations that might, for all the speaker
545 presupposes, be compatible with [the addressee’s] beliefs” (Stalnaker 1988, 157; see Geurts 1998 for another example of this approach). In principle *both* of these contexts are “available to be exploited” in interpreting the complement clause (158).

Multiple-context approaches provide a straightforward treatment of sentences like (31), uttered when Tom is not present and it’s common ground that the woman demonstrated has never smoked.

550 (31) Tom believes that that woman has quit smoking.

In particular, we can say that the global context satisfies the presuppositions of the demonstrative ‘that woman,’ while the local context satisfies the presuppositions of ‘quit.’ But one-context approaches can handle this sort of example, too, as long as they give an appropriate story about the content of c_2 . Clearly such approaches cannot simply identify c_2 with Stalnaker’s local context because c_2 does not
555 satisfy the presuppositions of ‘that woman.’ But c_2 could be the actual conversational context tweaked just enough so that needed presupposition-satisfying content can come from what we presuppose to be Tom’s beliefs: c_2 could be, as it were, a ‘mix’ of Stalnaker’s basic (global) and derived (local) contexts. This is in effect just taking what Heim (1983) says about local accommodation in general (254–255), and applying it to belief ascriptions.

560 How, then, can we decide between one-context and multiple-context approaches? If our aim were simply to explain how local accommodation in belief ascriptions affects *whether* presupposition failure occurs, then I suspect that there wouldn’t be much basis for this decision. But there is good reason to think that local accommodation also can affect *how* an expression in the complement of an attitude ascription is interpreted, and I think that multiple-context approaches do better at accounting for
565 certain cases of this kind.

The hypothesis that local accommodation can affect how expressions are interpreted—henceforth, the *local interpretation hypothesis*—may sound radical. It contravenes Kaplan, for example, who influentially insists that “no operator can control the character of the indexicals within its scope,” that English does not contain “operators like ‘In some contexts it is true that,’ which attempt to meddle with
570 character” and even that “such operators *could not be added* to English” (1977, 510–511). (He seems to think these claims follow from the hypothesis that “Indexicals, pure and demonstrative alike, are directly referential” (492).) But Kaplan himself cites Rich Thomason’s “Never put off until tomorrow what you can do today” (506), and Partee (1989, 270) offers many fascinating examples like “In all my travels, whenever I have called for a doctor, one has arrived within an hour.” Whether or not these
575 examples are best analyzed as cases of local accommodation, they lend plausibility to the idea that the context relevant to the interpretation of a given expression needn’t be the global context. (Cf. article 4 *Reference*, article 68 *Indexicality and Logophoricity*, and article 100 *Context dependency*.)

Moreover, it is theoretically costly to deny the local interpretation hypothesis. We would need strong reasons to resist it, once we countenance local accommodation as part of the explanation of the readings of (26)–(28) already discussed. It is hard to see what those reasons could be.

(26) Ken believes that Louise has spoken again.

(27) Hob believes that every spy has left.

(28) Tom believes that Sue has quit smoking.

On the treatment of these sentences that I sketched earlier, local contexts are sometimes available to satisfy presuppositions carried by expressions in ‘that’-clauses. And it is widely accepted that context can affect the interpretation of many expressions. It would be invidious to insist that although a local context is available to a given context sensitive expression, the interpretation of that expression is blind to the local context, influenced only by the global one.

It is also fruitful to endorse the local interpretation hypothesis. Consider the following example. Conversational context plausibly makes a difference to the intension associated with ‘best’ with respect to at least the two dimensions of the class of contestants and the scale of comparison. In non-embedded environments these two dimensions are obviously both determined by the global conversational context. But embedded environments are more complicated. Suppose Steve evaluates cakes 1 through 5, ranking 1 best, 2 next best, and so on to 5, which he says is worst. Unbeknownst to him, a cake contest is going on, and we know that exactly cakes 3, 4, and 5 are the competitors. But of the cakes in the contest, Steve thinks that cake 3 is the best. Keeping all this in mind, I think there is a reading on which (32) is true:

(32) Steve thinks cake 3 is the best.

This suggests that the global conversational context here determines the class of competitors relevant to the intension of ‘best.’ Now suppose that it is common ground that Steve is evaluating the cakes on the basis of how light they are—1 is like gossamer, 5 is far too dense—and it’s also common ground that flavor is the only relevant scale of comparison for purposes of the contest. If the global conversational

context also determined the scale of comparison for ‘best,’ then (32) would have to attribute to Steve the belief that cake 3 is the best of cakes 3, 4, and 5 *in flavor*. But it has a reading, I think, on which it
605 attributes the belief that cake 3 is the best of the relevant cakes by whatever *Steve’s* scale of comparison is.

Phenomena like these threaten to crop up for any expression that is sensitive to context in multiple dimensions. We might analyze (32) in a purely intensional way—as, roughly, “In every world *w* compatible with Steve’s beliefs, cake 3 is the best among the actual class of competitors according
610 to the scale of comparison operative in *w*,” but this treatment is committed to an surprising amount of syntactic complexity in an expression like ‘best.’ Moreover, I think it is quite odd to insist that although conversational context determines the scale of comparison when ‘best’ occurs in unembedded environments, it is determined by the binding of intensional variables in cases like (32). Finally, if it’s common ground that Steve is evaluating on the basis of texture (although he is in fact evaluating on
615 the basis of density) then I think there’s a reading of (32) on which it attributes to him the belief that cake 3 is the best with respect to texture. But the intensional treatment cannot explain this reading, because the scale of comparison, if not determined by the conversational context, can be determined only by Steve’s actual beliefs, and not by what we presuppose his beliefs to be.

We can avoid these problems by appealing to local accommodation. For example, on a one-context
620 view we could say that ‘cake 3 is the best’ is interpreted relative to a single context according to which the class of competitors is cakes 3, 4, and 5, and the scale of comparison is lightness. Or, on one multiple-context view we might say that the class of competitors is determined by the global context, whereas the scale of comparison is determined by the local context. Or, because the local context is given not by Steve’s beliefs simpliciter but rather by what the conversational participants *presuppose* to be Steve’s
625 beliefs, we might say that the local context determines both the class of competitors and the scale of comparison. On this line, we treat it as true for purposes of conversation that Steve knows something about the contest (by knowing which cakes are the competitors) without pretending that he knows everything we know about it (since ‘best’ can still be evaluated relative to his scale of comparison). The content of that pretense is the local context. Note that there’s nothing remarkable about this local
630 context: it’s easy to imagine global conversational contexts in which it’s presupposed that cakes 3, 4,

and 5 are the competitors and the scale of competition is either left an open question or resolved to features that are not in play in the actual competition. With the flexibility it affords in our explanations of these phenomena, the local interpretation hypothesis looks quite fruitful.

We have not yet considered any cases that will help us decide between one-context and multiple-
635 context approaches to local accommodation. The most compelling such cases are those in which it seems plausible that occurrences of the same expression in an embedded environment get different interpretations. As Stalnaker notes in passing, his two-context approach can “account for Russell’s notorious yacht” (1988, 159):

(33) Speaking of Russell’s yacht—Moore believes that it is longer than it is.

640 Very roughly, the idea is that the two occurrences of ‘it’ in (33) are interpreted relative to different contexts, and the differences between those contexts—the local and global contexts—are such that the two occurrences are interpreted differently. As a result the complement clause as a whole does not express a or the necessarily false proposition. But if, by contrast, the complement clause of (33) is always interpreted relative to a single context, then *whatever* recipe we give for ‘mixing’ the global and
645 local contexts we will not be able to explain the belief ascription in (33), because both occurrences of ‘it’ will be interpreted relative to the same context. Similarly for

(34) Pierre doesn’t realize that London is London.

By holding that the two occurrences of ‘London’ have different denotations because they are interpreted relative to relevantly different contexts, we have the beginnings of a story of how (34) can mean that
650 Pierre doesn’t realize that some contingent proposition is true. It is very hard to see how to tell such a story on a one-context approach.

In characterizing the effects that context can have on semantic interpretation, it’s helpful to treat any given expression—‘again,’ ‘every,’ ‘quit,’ ‘best,’ ‘it,’ ‘Paderewski,’ ‘London,’ or what have you—as associated not only with an intension or intensions, but also with a *hyperintension*. As I will use the term,
655 the hyperintension of an expression is a relation between contexts and intensions or semantic values.

To a first approximation, a context bears the hyperintension of an expression to some intension or intensions just in case, in that context, the expression is best interpreted by those intensions. (Strictly speaking we should be careful not to assume that there are such ‘best interpretations,’ but the necessary workaround would lead us far afield.) Positing hyperintensions makes it easier to characterize the
660 following hypotheses about the relationship between context, semantics, and linguistic competence:

1. Language users can get by perfectly well without knowing everything there is to know about the hyperintensions of their language’s expressions.
2. Much of the knowledge that is relevant to knowing about the features of an expression’s hyperintension is both a posteriori and, intuitively, non-linguistic. While one can’t be linguistically
665 competent without *some* knowledge of hyperintensions, linguistic competence itself needn’t bring much knowledge of hyperintensions. (Cf. Chalmers 2006 and article 33 *Word knowledge vs. world knowledge*.)
3. Some contexts may not yield an intension for a given expression. So an expression’s hyperintension need not be defined for every possible context.
- 670 4. A context may determine more than one intension for a given expression. For example, in a context that does not resolve whether we are talking about color or weight, I claim that ‘light suit’ has (at least) two semantic values. One is the semantic value that it would have in a context that did resolve that we were talking about color, and the other is the semantic value that it would have in a context that resolved that we were talking about weight. (Intensional functional application can still apply to combinations of *particular* images of contexts under hyperintensions,
675 yielding a proposition for each combination. If context doesn’t determine exactly one intension for an expression in a sentence, then in general it won’t determine exactly one intension for the sentence, either.)

Kaplan, like many others, allows that directly referential expressions are associated with “semantical
680 rules which determine the referent in each context of use” (495). According to the view I am urging here, hyperintensions simply codify the rules whereby particular intensions come to be the semantic

values of expressions in particular contexts. Despite the complexity of typical hyperintensions, the intensions associated with an expression that Kaplan would classify as “directly referential” will, in normal contexts, be simple constant functions. I am bracketing some hard questions about (among
685 other things) which expressions are indexical, the nature of the hyperintensions associated with indexical expressions, and our knowledge of those hyperintensions (cf. article 33 *Word knowledge vs. world knowledge*). Those issues to the side, however, I agree with Kaplan that proper names are not associated with a “cognitive content” that fixes their reference in all contexts (506, footnote 31). More specifically, I agree that competent language users may be ignorant of many features of an expres-
690 sion’s hyperintension (points 1 and 2 above) and that hyperintensions needn’t fix terms referents in all contexts (points 3 and 4 above).

In dramatic cases of identity confusion, like that described in the background for (34), it’s plausible that the contexts provided by local accommodation will determine more than one intension for the relevant locally accommodated expression. (I think that the hypothesis that even in global contexts
695 they sometimes determine more than one intension helps explain anaphoric reference to nonexistent objects, but for present purposes I am not taking a stand on the question.) We also get a better characterization of the content of Pierre’s beliefs if we drop the assumption that ‘London’ in (34) refers, in either of its occurrences, to actual objects: in effect we instead appeal to objects that *would* have been the referent of ‘London,’ if the world had been different in relevant ways, to characterize Pierre’s belief
700 state in an accurate and efficient way. As in the simpler cases considered earlier, we use counterfactual reasoning to determine the relevant features of these hyperintensions: we ask how we would interpret ‘London’ if the presuppositions we brought to bear on its interpretation were more like some of Pierre’s ‘London’ beliefs. Thus a pair of hyperintensions both of which take the actual conversational context to a constant intension—thereby counting as coreferential and rigid—may take a context introduced
705 by local accommodation to quite different intensions.

It is important to see that a proper name need not *always* be locally accommodated when we presuppose that the ascriber believes that the name has a different referent than we do. Even if it’s common ground that Glenda knows Bob Dylan only as her childhood friend Robert Zimmerman, if she thinks he has a beautiful voice then in some contexts (35) seems true.

710 (35) Glenda believes that Bob Dylan has a beautiful voice. (Saul 1998, 366)

I suggest that in such contexts we see the globally accommodated reading of ‘Bob Dylan,’ because it is manifest that what the speaker is trying to convey with (35) is that Glenda believes that *a voice with the qualities of Dylan’s* is beautiful.

This kind of treatment lets us give a simple, clean treatment of coreferential proper names, demonstratives, pronouns and the like. Names that are coreferential in a context of use have the same constant
715 intension (and so the same referent) in *that* context, but in other contexts may have constant intensions picking out different referents. So while we learn much about the semantics of attitude ascription from coreference puzzles, such puzzles do not force us to complicate the semantics of simple sentences. Moreover, because this treatment of coreference puzzles does not involve Fregean senses or any other
720 ‘fine-graining’ of content, it is compatible with the probabilistic representations of doxastic states that claims like (25) make attractive.

(25) Clara suspects and Doug doubts that it will rain.

Finally, this treatment unifies two important features of attitude ascription:

1. Names that are coreferential in the *global* context often seem to make different semantic
725 contributions to attitude ascriptions.
2. Presupposition carrying expressions in attitude ascriptions are often interpreted from a point of view more like that of the ascriber than that of the conversational participants.

I think this account is quite plausible once we think of proper names as just another kind of presupposition-carrying expression, thus warranting the treatment we would give to any other presupposition-carrying expression. The prospects for extending this kind of treatment to discourse anaphora are
730 promising, I think, but extending the approach in that direction must be left for another time.

7. Conclusion

Work on attitude ascription within the philosophical literature has been dominated by the consideration of coreference puzzles. These puzzles are inarguably important: no serious theory of attitude ascription can afford to ignore them. Some aspects of logical omniscience are similarly serious and vexing. Nonetheless, probabilistic language finds a natural home in attitude ascription, as I argued at the beginning of Section 5. We must either work within the constraints imposed by the probabilistic representation of doxastic states, or work to loosen those constraints by developing different ways to represent subjective uncertainty.

There is much, much more to say about attitude ascription. I hope to have made the case that sustained interaction between a broad range of researchers—those who most naturally self-identify as working in semantics, pragmatics, philosophy of language, philosophy of mind, and epistemology—will help make future research fruitful.

Keywords:

propositional attitudes, coreference, belief

Eric Swanson, Ann Arbor (USA)

For helpful discussion thanks to Sarah Moss, Paul Portner, Mark Richard, Bob Stalnaker, and Steve Yablo.

8. References

- Asher, Nicholas 1986. Belief in Discourse Representation Theory. *Journal of Philosophical Logic* 15, 127–189.
- Brentano, Franz 1874. The Distinction between Mental and Physical Phenomena. In: Roderick M. Chisholm (ed.) *Realism and the Background of Phenomenology*, Atascadero, CA: Ridgeview Publishing Company. 39–70. Translated by D. B. Terrell.

- Chalmers, David J. 2006. Two-Dimensional Semantics. In: Ernest Lepore & Barry C. Smith (eds.) *The Oxford Handbook of Philosophy of Language*, Oxford: Oxford University Press. 574–606.
- Chalmers, David J. 2009. Probability and Propositions. Ms., Australian National University.
 760 <http://consc.net/papers/probability.pdf>, September 20, 2009.
- den Dikken, Marcel, Richard Larson & Peter Ludlow 1996. Intensional ‘Transitive’ Verbs and Concealed Complement Clauses. *Rivista di Linguistica* 8, 331–348.
- den Dikken, Marcel, Richard Larson & Peter Ludlow 1997. Intensional Transitive Verbs and Abstract Clausal Complementation. Ms., SUNY Stony Brook. <http://semlab5.sbs.sunysb.edu/%7Erlarson/itv.pdf>,
 765 September 7, 2009.
- Egan, Andy, John Hawthorne & Brian Weatherson 2005. Epistemic Modals in Context. In: Gerhard Preyer & Georg Peter (eds.) *Contextualism in Philosophy*, Oxford: Oxford University Press. 131–168.
- Egan, Andy & Brian Weatherson (eds.) 2009. *Epistemic Modality*. Oxford: Oxford University Press.
- 770 Forrester, James William 1984. Gentle Murder, or the Adverbial Samaritan. *Journal of Philosophy* 81(4), 193–197.
- van Fraassen, Bas C. 1972. The Logic of Conditional Obligation. *Journal of Philosophical Logic* 1, 417–438.
- van Fraassen, Bas C. 1973. Values and the Heart’s Command. *Journal of Philosophy* 70(1), 5–19.
- 775 Frege, Gottlob 1892. On *Sinn* and *Bedeutung*. In: Michael Beaney (ed.) *The Frege Reader*, Oxford: Blackwell Publishers, Ltd. 151–171. Translated by Max Black.
- Geurts, Bart 1998. Presuppositions and Anaphors in Attitude Contexts. *Linguistics and Philosophy* 21, 545–601.
- Gibbard, Allan 1981. Two Recent Theories of Conditionals. In: William L. Harper, Robert Stalnaker & Glenn Pearce (eds.) *Ifs: Conditionals, Belief, Decision, Chance, and Time*, Dordrecht: D. Reidel
 780 Publishing Company. 211–247.
- Halpern, Joseph Y. 2003. *Reasoning About Uncertainty*. Cambridge: MIT Press.

- Hansson, Bengt 1969. An Analysis of Some Deontic Logics. *Noûs* 3, 373–398.
- Heim, Irene 1982. *The Semantics for Definite and Indefinite Noun Phrases*. Ph.D. thesis, University of
 785 Massachusetts, Amherst.
- Heim, Irene 1983. On the Projection Problem for Presuppositions. In: (Portner & Partee 2002),
 249–260.
- Heim, Irene 1991. Articles and Definiteness. Ms., Department of Linguistics and Philosophy, MIT.
 Originally published in German as "Artikel und Definitheit," in von Stechow and Wunderlich
 790 (eds.) *Semantics: An International Handbook of Contemporary Research*, Berlin: de Gruyter, 1991.
- Heim, Irene 1992. Presupposition Projection and the Semantics of Attitude Verbs. *Journal of Semantics*
 9, 183–221.
- Huddleston, Rodney & Geoffrey K. Pullum (eds.) 2002. *The Cambridge Grammar of the English Lan-
 guage*. Cambridge: Cambridge University Press.
- 795 Kamp, Hans 1981. A Theory of Truth and Semantic Representation. In: (Portner & Partee 2002),
 189–222.
- Kamp, Hans 1988. Comments on Stalnaker. In: Robert H. Grimm & Daniel D. Merrill (eds.) *Contents
 of Thought*, Tucson: The University of Arizona Press. 156–181.
- Kaplan, David 1977. Demonstratives: An Essay on the Semantics, Logic, Metaphysics, and Epistemol-
 800 ogy of Demonstratives and Other Indexicals. In: Joseph Almog, John Perry & Howard Wettstein
 (eds.) *Themes from Kaplan*, Oxford: Oxford University Press. 481–563.
- Karttunen, Lauri 1968a. What Do Referential Indices Refer To? Rand Corporation report P3854.
<http://www.rand.org/pubs/papers/P3854>, September 20, 2009.
- Karttunen, Lauri 1968b. What Makes Definite Noun Phrases Definite? Rand Corporation report
 805 P3871. <http://www.rand.org/pubs/papers/P3871>, September 20, 2009.
- Karttunen, Lauri 1973a. The Last Word. Ms., University of Texas, Austin.
- Karttunen, Lauri 1973b. Presuppositions of Compound Sentences. *Linguistic Inquiry* 4, 167–193.
- Karttunen, Lauri 1974. Presupposition and Linguistic Context. *Theoretical Linguistics* 1, 181–194.

- Karttunen, Lauri 1976. Discourse Referents. In: James D. McCawley (ed.) *Syntax and Semantics 7: Notes from the Linguistic Underground*, New York: Academic Press. 363–385. First presented in 1969.
810
- King, Jeffrey C. 1996. Structured Propositions and Sentence Structure. *Journal of Philosophical Logic* 25, 496–521.
- King, Jeffrey C. 2007. *The Nature and Structure of Content*. Oxford: Oxford University Press.
- Kratzer, Angelika 1991. Modality. In: Arnim von Stechow & Dieter Wunderlich (eds.) *Semantics: An International Handbook of Contemporary Research*, Berlin: W. de Gruyter. 639–650.
815
- Kripke, Saul 1979. A Puzzle About Belief. In: (Ludlow 1997), 875–920.
- Kripke, Saul 1980. *Naming and Necessity*. Cambridge: Harvard University Press.
- Larson, Richard K. & Peter Ludlow 1993. Interpreted Logical Forms. In: (Ludlow 1997), 993–1038.
- Lewis, David K. 1973. *Counterfactuals*. Malden, MA: Basil Blackwell Ltd.
- 820 Lewis, David K. 1976. Probabilities of Conditionals and Conditional Probabilities. In: *Philosophical Papers*, Oxford: Oxford University Press, volume 2. 133–56. With postscript.
- Lewis, David K. 1981a. Ordering Semantics and Premise Semantics for Counterfactuals. *Journal of Philosophical Logic* 10, 217–234.
- Lewis, David K. 1981b. What Puzzling Pierre Does Not Believe. In: *Papers in Metaphysics and Epistemology*, Cambridge: Cambridge University Press. 408–417.
825
- Ludlow, Peter (ed.) 1997. *Readings in the Philosophy of Language*. Cambridge: MIT Press.
- MacFarlane, John 2003. Epistemic Modalities and Relative Truth. Ms., University of California, Berkeley. <http://johnmacfarlane.net/epistmod-2003.pdf>, September 20, 2009.
- McCawley, James D. 1974. On identifying the remains of deceased clauses. *Language Research* 9(2),
830 73–85.
- Montague, Michelle 2007. Against Propositionalism. *Noûs* 41, 503–518.
- Partee, Barbara H. 1989. Binding Implicit Variables in Quantified Contexts. In: *Compositionality in Formal Semantics*, Oxford: Blackwell Publishers, Ltd. 259–281.

- Portner, Paul & Barbara H. Partee (eds.) 2002. *Formal Semantics: The Essential Readings*. Oxford: Blackwell Publishers Ltd.
- 835
- Prior, A. N. 1958. Escapism: The Logical Basis of Ethics. In: *Essays in Moral Philosophy*, Seattle: University of Washington Press. 135–146.
- Quine, W. V. O. 1956. Quantifiers and Propositional Attitudes. *Journal of Philosophy* 53(5), 177–187.
- Richard, Mark 1983. Direct Reference and Ascriptions of Belief. *Journal of Philosophical Logic* 12,
- 840 425–452.
- Ross, J. R. 1976. To Have and to Have Not. In: Polome et al. (ed.) *Linguistic and Literary Studies in Honor of Archibald Hill*, Lisse: Peter de Ridder.
- Russell, Bertrand 1905. On Denoting. In: Gary Ostertag (ed.) *Definite Descriptions: A Reader*, Cambridge: MIT Press. 35–49.
- 845 Russell, Bertrand 1919. Descriptions. In: (Ludlow 1997), 323–334.
- Salmon, Nathan 1986. *Frege's Puzzle*. Cambridge: MIT Press.
- Saul, Jennifer 1998. The Pragmatics of Attitude Ascription. *Philosophical Studies* 92(3), 363–389.
- Scheffler, Israel 1955. On Synonymy and Indirect Discourse. *Philosophy of Science* 22(1), 39–44.
- Searle, John 1958. Proper Names. *Mind* 67(266), 166–173.
- 850 Soames, Scott 1987. Direct Reference, Propositional Attitudes, and Semantic Content. In: Nathan Salmon & Scott Soames (eds.) *Propositions and Attitudes*, Oxford: Oxford University Press. 197–239.
- Soames, Scott 2002. *Beyond Rigidity: The Unfinished Semantic Agenda of Naming and Necessity*. Oxford: Oxford University Press.
- 855 Stalnaker, Robert C. 1974. Pragmatic Presuppositions. In: (Stalnaker 1999a), 47–62.
- Stalnaker, Robert C. 1984. *Inquiry*. Cambridge: MIT Press.
- Stalnaker, Robert C. 1988. Belief Attribution and Context. In: (Stalnaker 1999a), 150–166.
- Stalnaker, Robert C. 1991. The Problem of Logical Omniscience, I. In: (Stalnaker 1999a), 241–254.

- Stalnaker, Robert C. 1996. Impossibilities. *Philosophical Topics* 24(1), 193–204.
- ⁸⁶⁰ Stalnaker, Robert C. 1998. On the Representation of Context. In: (Stalnaker 1999a), 96–113.
- Stalnaker, Robert C. 1999a. *Context and Content*. Oxford: Oxford University Press.
- Stalnaker, Robert C. 1999b. The Problem of Logical Omniscience, II. In: (Stalnaker 1999a), 255–273.
- Stanley, Jason 1997. Names and Rigid Designation. In: Bob Hale & Crispin Wright (eds.) *A Companion to the Philosophy of Language*, Oxford: Blackwell Publishers, Ltd. 555–585.
- ⁸⁶⁵ Swanson, Eric 2006. *Interactions with Context*. Ph.D. thesis, Massachusetts Institute of Technology.
- Swanson, Eric 2008. Modality in Language. *Philosophy Compass* DOI: 10.1111/j.1747-9991.2008.00177.x.
- Swanson, Eric 2009. How Not to Theorize about the Language of Subjective Uncertainty. In: (Egan & Weatherson 2009).
- ⁸⁷⁰ Thau, Michael 2002. *Consciousness and Cognition*. Oxford: Oxford University Press.
- Yalcin, Seth 2007. Epistemic Modals. *Mind* 116, 983–1026.
- Yalcin, Seth 2009. Nonfactualism about Epistemic Modality. In: (Egan & Weatherson 2009).