

Perspective in Interpretation

Abstract: A variety of linguistic devices, semantic and pragmatic, guide the addressee during interpretation to consider only certain (types of) situations and to adopt a particular perspective on the information presented about them. Among other things, the perspective suggested leads one to attend to only certain facets of the situations under consideration and only certain properties of the individuals in them; this is sometimes called the intended mode of presentation. I consider one way of modeling perspective in a situation semantics and explore some of its ramifications for interpretation, with attention to domain selection, *de dicto* and *de se* interpretation, incomplete descriptions and specific interpretations of indefinites.

I. Attention, intention, and attitude

Normally, at any given time we attend to what is evident from the perspective we have adopted at that point. Actually, we may even ignore some of the scene before us, if facets of it do not bear on our present goals. Our intentions effectively screen out irrelevant information. Hence, the attitudes¹ we adopt toward the world are central to what we take into consideration.

All normal human beings know this of each other, and in endeavoring to grasp what another is attempting to communicate, we attune ourselves to that person's perspective on the matter at hand by empathetically adopting their perspective. (Autistic children are not generally able to recognize another's stance and intentions—empathy eludes them—and it is a great handicap in acquiring language.) The normal speaker assists in this process by giving clues about the perspective from which the information to be conveyed can most readily, efficiently be grasped. Consider an example:

- (1) Tim and Margaret are sitting at a conference table in her psycholinguistics lab, working on a grant proposal. Tim is making notes on his Compaq-brand laptop.

Margaret: How do you like your Compaq?

¹ American Heritage Dictionary, 4th edition, 2000: **Attitude**: **1.** A position of the body or manner of carrying oneself: *stood in a graceful attitude*. See synonyms at **posture**. **2a.** A state of mind or a feeling; disposition: *had a positive attitude about work...* **3.** The orientation of an aircraft's axes relative to a reference line or plane, such as the horizon. **4.** The orientation of a spacecraft relative to its direction of motion. All these senses are relevant to the current work.

- Tim: It's not bad, but it's getting kind of old. I wish I had a Mac. Macs are far better for graphics, and it turns out that I'm doing a lot more graphics than I'd expected.
- Margaret: [gesturing with her thumb over her shoulder, and slightly turning her head in the direction of her desk in the middle of the room] I just got [that]_F last year.

In the direction Margaret is pointing there's a lot of stuff: a desk with a big pile of papers and a flat screen Dell computer monitor on it, past the desk an eye-tracker, past that a wall calendar.

It is quite clear in (1) that by her use of *that* in the last line Margaret intends to refer to her computer. How does she convey this? It is natural for her to assume that Tim will attend to the vector extending from her through the lab in the direction that her thumb is pointing and her head is inclining. Further, since the question under discussion is about computers, it is natural that she will not expect Tim to take into consideration various other elements located along that vector. So deixis works partly by adopting the perspective of another individual—here the view of the lab from Margaret's position in the direction she is pointing—and partly by recognizing what matters to her at that moment in the conversation.

Perspective isn't static in discourse. A skillful speaker directs the attention, the metaphorical gaze, of her interlocutors first from one situation, seen in the light of certain questions or other considerations, to another—or was it the same situation from a different angle? The resulting perspective corresponds to a vista, a limited scene commanding our attention. And as we shift our perspective in this way, first one set of entities becomes more prominent and then another. Irrelevant entities are masked off from consideration—not non-existent, but as good as not-there because they do not fall under our gaze:

- (2) [Giving directions to a stranger:] After passing through the light at Indianola, go four more blocks, then turn left at the fire station.

Usually, the explicit descriptive content of a given definite NP by itself is just a hint about the speaker's referential intentions: a description which suggests a way of picking out a particular individual in the model or the world, or a particular discourse referent among all those in the interlocutors' common ground. Both the speaker of (2) and the addressee know that there is more than one fire station in the world. But from a practical point of view, the role of descriptive content in definite NPs is not to single-handedly pick out some entity in the world which is unique in virtue of bearing the property corresponding to the descriptive content. In order to assure that the addressee can retrieve the intended (discourse) referent, in the face of the failure of Russellian uniqueness, the speaker need only direct the addressee's attention in such a way that only a limited set of entities come under scrutiny.² (2) gives directions, intended to be understood in a context of use, and at the time of use the intended interpretation of *the*

² This is a more general case of the notion of RETRIEVABILITY that I explore in Roberts (2006), where I was concerned with retrieving the intended referent of an anaphoric or elliptical element.

fire station will be uniquely evident and salient among those entities which are evident from the perspective of a person traveling in the direction suggested by the speaker, at the point where they have gone four blocks past the light at Indianola. There will be nothing incomplete about the description in this context.³

We have concrete mechanisms in discourse that help us at any give time to convey the perspective we have in mind. One class of such mechanisms is what someone (I can't remember who) called *frame adverbials*: sentential adverbials, often clausal, that serve to set the scene about which the speaker is to proffer information in the matrix sentence to which they are adjoined. The order in which such adverbials occur can have an effect on truth conditions. To give a flavor of this complex phenomenon, consider the temporal reflexes of surface order in some sentences with multiple temporal adverbial clauses:

- (3) When Joan was in Kansas City, she took a walk after she gave her talk.
- (4) After Joan gave her talk, she took a walk when she was in Kansas City.

- (3') When Joan was in Kansas City, after she gave her talk she took a walk.
- (4') After Joan gave her talk, when she was in Kansas City she took a walk.

- (5) When John ran the marathon, he sprained his ankle while he was crossing the finish line.
- (6) #While John was crossing the finish line, he sprained his ankle when he ran the marathon.

(3) and (3') are synonymous. Whether the *after*-clause is S-final or is the second preposed temporal adverbial clause makes no difference to the interpretation—we take it that Joan gave her talk, and took her walk, while in Kansas City. But with non-parenthetical intonation for the second temporal adverbial clause, (3)/(3') are not generally taken to be synonymous with (4) or (4'). In both of the latter, we assume that Joan went to Kansas City (where she took a walk), only after she gave her talk. Similarly, in (5) we assume that John crossed the finish line while he was running the marathon; but in (6) the inclusion is reversed, leading to pragmatic infelicity due to the relative lengths of time it generally takes to do such things. Of course, in (4)/(4') and (6), marked prosody can yield the readings in (3)/(3') or (5): For example, one can introduce a break after the main clause in (6) and before the *when*-clause, as well as changing the pitch range of the final adverbial, making the latter seem like an afterthought. The point here is that without such marked intonation, the order of the clauses has pragmatic consequences which can lead to truth-conditional distinctions.

Also note that it isn't that *while/when* sequences are always bad. Compare the acceptability of (7):

³ There are a host of previous authors who have written about the idea that limiting consideration to a narrow situation can limit the domain. Had I more time, I'd list them here. Barwise & Perry (1983) obviously come to mind. The focus here is not on the fact that this occurs, but why and how.

- (7) While John was crossing the finish line, he sprained his ankle when a spectator tripped him.

It's just that you can't run a marathon during the time it takes you to cross a finish line. When both temporal adverbial clauses are preposed, the version in (6') is only pragmatically acceptable when the *when*-clause receives parenthetical intonation, as if it were something the speaker had originally presupposed but decided late in the utterance to remind or inform the hearer about:

- (5') When John ran the marathon, while he was crossing the finish line he sprained his ankle.
(6') While John was crossing the finish line, when he ran the marathon he sprained his ankle.

These effects are most likely pragmatic in origin; one might think of sentential adverbials like the above as syntactic satellites of the sentence to which they adjoin, since their integration into compositional interpretation is almost certainly less well-integrated than is sometimes assumed. But whatever their source, the effects are quite robust.⁴

There are many types of frame adverbials: temporal adverbials of various types, including adverbs like *recently*, PPs like *upon receiving payment*, and a variety of temporal adverbial clauses; locatives like *immediately to her right* or *from where Rudolph sat*; certain free adjuncts and absolutes (Stump 1985) like *ascending the pyramid* or *lashed to the stern*; and hypotheticals like *if*-clauses, *assuming you've arrived at the station* or *in the event that...* All these can co-occur, and they interact in ways that tend to respect the generalizations discussed above. In fact, when iterated adverbials are of mixed type, the correlation between surface order and framing order tends to be even more fixed. To give just the flavor, consider the interpretive effects of order between mixed temporal adverbial clauses and *if*-clauses in the following:

- (8) When John came to visit, if he found some money he hid it.
(9) If John found some money, when he came to visit he hid it.

(8) suggests that the possibility of John's finding money occurred during the visit, whereas in (9) we would presume that the entertained possibility would have taken place prior to the visit.

⁴ On the basis of these examples and many others, in unpublished work from 1994 I hazarded the following generalizations about temporal adverbials in English: (a) The preferred order of framing is identical to surface order, from the first to the last adverbial in preposed position through an S-final adverbial. But (b) the tendency in (a) can be overridden S-initially, especially if the first temporal adverbial clause presupposes the second (e.g., through cataphora). So, (6'') is a bit better than (6'):

(6'') While he was crossing the finish line, when John ran the marathon he sprained his ankle. And (c) there seems to be a single slot for frame-adverbials S-final, and adverbials after that are interpreted as something like afterthoughts. The order of framing of an afterthought S-final temporal adverbial clause is apparently not so fixed as that of an S-final temporal adverbial clause in the syntactic adverbial slot.

Rooth (1985, Chapter 5) noticed that there is a systematic difference between preposed and postposed temporal adverbials:

- (10) John always shaves when he is in the shower.
- (11) When he is in the shower, John always shaves.
 - (a) **always** (John is in the shower) (John shaves)
 - (b) **always** (John shaves) (John is in the shower)

(10) has both the reading in (a), where the *when* clause serves as restrictor on the adverb of quantification *always*, and that in (b), where it is the nuclear scope. These two readings are readily disambiguated by the prosodic patterns associated with them; the focused clause serves as nuclear scope, the other as restrictor. But (11) only has reading (a); the preposed *when*-clause cannot serve as nuclear scope of the operator. To account for this, Rooth speculates that the preposed clause is in a topic position, hence outside the potential scope of the operator in the main clause, while the S-final adverbial is not, so can serve as its scope. Iatridou (1991) observed that preposed adverbials cannot take narrow scope with respect to negation, additional evidence that they are outside the potential scope of main clause operators.

It has been noted in the literature that a wide range of S-initial adverbials can serve to restrict the domains of various types of operators. Consider Lewis 1975; Kratzer 1979, 1981, 1991; Kamp 1979; Stump 1981, 1985; Hinrichs 1981, 1985, 1986; Heim 1982; Farkas & Sugioka 1983; Partee 1984; Rooth 1985; Roberts 1989, 1995; de Swart 1991, 1993; and von Stechow 1994, among others. The operators discussed in that literature include not only adverbs of quantification, but also modals, the temporal operators implicit in tense (the restriction serving as a Reference Time), and ordinary quantificational determiners. But it seems that all types of frame adverbials, including the non-sentential types noted above can serve to restrict operators' domains. Consider this example:

- (11) When John was hiking the Berkshires, from the top of Mount Greylock he saw Worthington to the southeast.

If one takes the tense in *saw* in the main clause of (11), like the pronoun *he*, to be anaphoric, as several of the references cited have argued, then the frame adverbials that precede satisfy those anaphoric presuppositions; but they do much more. The temporal adverbial not only suggests a restricted interval during which the sighting proffered in the main clause occurred—serving as the Reference Time or Reference Eventuality for the past tense in *saw*, but it locates it in space, as well, since the hiking took place in the Berkshires. In the course of doing this, John is made salient and relevant. Then the locative adverbial leads us to zoom in on that spatiotemporal portion of the hike that took John to the top of Mount Greylock. So, it suggests that we consider the view from the top of the mountain. And it constrains the relevant temporal extent of the hike to just that

portion of the hike that took place while John was on the peak.⁵ We understand that John actually saw Worthington (as opposed to hallucinating it, which is perfectly reasonable in *While on peyote, John saw ancient Atlantis*), that he saw it at a distance, from above and from its northwestern aspect. If we continue:

(12) ...Then, with the summit behind him, he continued toward Mount Williams.

we not only know what *the summit* refers to, but (if we know our Berkshire landmarks) we know the direction his walk then took him in. Note that unique intended speaker's referent of *the summit* is clear despite the fact that in the main clause of the same sentence we mention another mountain, Mount Williams, which presumably also has a summit. The reference is clear because we know that typically on a hike after reaching one location we move on, leaving it behind, and hence that when we leave a summit it is behind us, with the one we are going towards lying ahead. The orientation thus plays a central role in licensing felicitous use of the definite and retrieving the intended speaker's referent. (12) is hence just a more complex example of the phenomenon exhibited in (2)—more complex, that is, from the point of view of modeling the practical reasoning involved in grasping the intended meaning, but in practice quite straightforward to competent speakers.

In Roberts (1995), I argued that the fact that all these adverbials may serve to restrict scopes of main-clause operators is not a coincidence, but that it reflects a general, pragmatically-driven phenomenon of domain restriction. The main thesis of the current work is that the mediating factor in this relationship is perspective. Actually, I believe that perspective is a central factor in a variety of pragmatic influences on interpretation, lexical as well as compositional. But domain restriction is one of the ways in which it bears most clearly on truth conditional interpretation, hence the one in which it first began to come to the attention of semanticists. If domain restriction turns out to be driven by perspective in a wide variety of cases, and if frame adverbials function to give clues to the intended perspective one is to take on the eventuality reported in the main clause, then it would not be surprising that this wide range of pre-posed adverbials all happen to also serve the function of domain restriction.

Note that preposed adverbials do not have to restrict the domain of an operator. It is difficult to demonstrate this with temporal adverbs, since these can always be taken to indicate the Reference Time (Kamp, Hinrichs, Partee), which Roberts (1995) argued to be a kind of domain restriction. But though preposed locative frame adverbials can (indirectly) restrict Reference Time, as we saw in (11), they do not always, as we see here:

(13) Mary worked hard all day to clean the yard. She found lots of treasures left behind by previous tenants. Inside a galvanized tub behind the garage, she found some rubber boots and a broken fishing pole. Behind the bushes, she found a small stuffed rabbit and an old soccer ball.

⁵ See <http://www.naturepanoramas.com/68001115.html>, though unfortunately this was at dawn on a cloudy day.

None of the sentences in (13) seems to suggest anything about the passage of time or the way that the sub-events reported in the second through fourth sentences were ordered or otherwise temporally related to each other. The adverbials *inside a galvanized tub behind the garage* and *behind the bushes*, simply serve to locate their respective discoveries spatially in the yard.⁶

And locative adverbials can seem serve to restrict the domains of operators other than tense, as we see in this example from Cresswell (1978:33):

(14) Behind the bushes everybody loves somebody.

But Cresswell points out that while in (14) *everybody* means ‘everybody behind the bushes’, (14) is not synonymous with (15):

(15) Everybody behind the bushes loves somebody.

“[15] does not require that anybody’s beloved be behind the bushes. This suggests that *behind the bushes* modifies the whole sentence [to which it is adjoined] and is true at a given world, time and place iff everybody at that place at that time in that world loves somebody at that time and place in that world.” Hence, the restriction in (14) appears to be global. Notice that it actually has a temporal flavor, as well. It does not suggest that the people in question *always* love somebody else in that group, but implicates that they only do so *while* they are behind the bushes. This also underlines the apparently global character of the restrictive effect of the locative. It leads the addressee to consider only the situation behind the bushes, in which the main clause is asserted to obtain.

As this suggests, I suspect that in general the contribution of frame adverbials to domain restriction is indirect. Their primary semantic contribution, of course, is to attribute additional properties to the eventuality described in the main clause (whether under the scope of an operator or not), properties which help to locate that eventuality in the relevant dimensions. See Stumpe (1985). This they do whether preposed or sentence-final, or in between. But their behavior when preposed is pragmatically distinguished: To say that they are topics does not do it justice. I think they serve to suggest facets of the perspective from which the truth of the matrix clause of the utterance is to be considered. Since (by hypothesis) perspective drives domain restriction, they also appear to serve that function. But examples like (15) are but one indication that we should be cautious about claims that such domain restriction is a straightforward function of syntactically given logical form, especially in view of the fact that presupposed adverbials seem to be outside the scope of operators in the main clause.

⁶ Notice also that here, despite the fact that on the specific occasion that’s the topic of discussion Mary presumably found no more than one broken fishing pole, one small stuffed rabbit and one old soccer ball, the uniqueness of these objects under those descriptions does not license the use of definite articles, which would be infelicitous in (13). We will return to consider the contrast between (13) and (2) in this regard in §4 below.

I have offered here a brief sketch of some of the properties of frame adverbials to make it plausible that when preposed they play a distinguished pragmatic role in discourse—that of offering a perspective from which the interpretation of the main clause is to proceed—and hence to try to motivate intuitively the claim that perspective plays a pervasive role in interpretation. However, this paper is not about frame adverbials *per se*, but about the phenomenon of perspective that they subserve. In the remainder of the paper, I will try to explain what I take perspective to be and suggest how it can be used to develop an account of the semantics and pragmatics of the broad class of phenomena that Mitchell (1986) called *perspectival*. The expressions of interest to Mitchell included those like *local*, *recent* and *above*, *come* and *go*, where the perspectival factor is fairly obvious. But he also considered perspectival relational notions like *foreign*, cathectic predicates like *boring* and *obvious*, and verbs of perception and attitude verbs, and the related phenomenon of interpretation *de se*. I agree that perspective can help shed light on all these phenomena, though I will model it quite differently from Mitchell. I also then want to argue that one measure of the extent to which Mitchell's thesis is correct is that it permits us to offer an improved account of domain restriction. In particular, I will focus on English definite and indefinite descriptions, where at present there is not even a consensus among researchers that the central interpretive problems involve domain restriction. Barwise & Perry (198?), among others, attempted an account of so-called incomplete descriptions which involved a resource situation; I want to argue that we can avoid certain now well-known pitfalls of that approach if we use the richer notion of a perspective, understood as a limitation on access to information about a situation. Similarly, I want to suggest that there is a conventionally perspectival component of the meaning of expressions like *certain* and *particular*, and, pragmatically, of the specific interpretation of indefinite NPs more generally. I agree with Portner (), Schwarzschild () and Breheny () that specificity in indefinites is a reflection of domain restriction. But exactly how does this domain restriction arise? In conjunction with insights due to Kratzer (199?), the exploration of the role of perspective in interpretation suggests a way of accounting for the complex data in this domain, but without assuming that indefinites themselves are ambiguous.

In the following section, I suggest a way of formally modeling the notion of perspective using a kind of epistemic modal accessibility relation. In §3 I turn to consider some of Mitchell's perspectival expressions, and sketch how the notion modeled in §2 helps to explain their contributions to the meaning of utterances in which they occur. In section §4 I apply the theory to descriptions, both definite and indefinite, and to the analysis of the meaning of *certain* and the related notion of specificity. In §5 I summarize, speculate about what perspective tells us about pragmatics and linguistic competence, and mention some unresolved or related issues, including recent debates about contextualism vs. relativism in semantics and the analysis of predicates of personal taste.

§2. Modeling perspective

In this section, I will propose one way of modeling the notion of perspective, which I will call on in the remainder of the paper. My aim is to give some more content to the view of

perspective that underlies the discussion, hoping that this will prove intuitive to the reader. Much more work would be required to provide a fully adequate fragment.

The basic conception is quite simple. The initial intuition is to think of a perspective as a way of viewing some concrete *situation*, which we might call the **ground** of the perspective. The situation is overlaid with an n -dimensional **coordinate system** which defines the relevant notion of *space* within the situation, establishing locations throughout the extent of that space. The canonical conception of perspective is spatial, and it is captured with the standard three-dimensional coordinate system. But that can be extended to characterize spatio-temporal perspective by adding another dimension, or to more abstract notions of point of view by using different coordinate systems, with more or fewer dimensions. A single location, identified by its coordinates, is selected as the **origin**; intuitively, this is the *point of view*. Extending from the origin, we draw a **vector** into the space; the vector represents the *line of sight* from the origin. But of course, typically from a given point of view in a situation we do not have access to information about the entire situation. Humans don't have eyes behind their heads, so the line of sight in connection with the limits of our binocular vision combine to limit visual access to information on either side of the vector. Moreover, there are probably rocks and trees or mountains somewhere along the line of sight or adjacent angle of view; or maybe desks and walls. Maybe it's foggy on Mount Greylock, and we can't see very far. Collectively, we might think of these conditions which limit our view as the **frame** of the perspective, drawing on the visual metaphor. More generally, the conditions in the frame will characterize any kind of limitation on access to information. Frame adverbials, hence, not only suggest a view but delimit it to a certain time span or location.

The situations that intuitively ground perspective in the simple cases are crucially partial—in them, we only have information about a limited circumstance, and this is what makes them useful for domain restriction. The notion of a grounding situation for a perspective is reminiscent of the *resource situations* of Barwise & Perry (1983) and Cooper (), and the *topic situations* in recent work by Kratzer (2005, 2006).

But instead of using situations as primitives, I'm going to model the partiality inherent in the notion of perspective using only whole possible worlds. There are two reasons for this. First, although I have worked with situations semantics of various sorts over the years, in the end I find non-maximal situations rather hard to work with in modeling; for example, in situation semantics it's more difficult to grasp what a proposition is and how it interacts semantically. Entailments and other semantic relations are clearer with possible worlds. Perhaps more importantly, I feel that the partiality which is crucial to understanding phenomena like perspective is not a partiality in the world, but in our information about the world.⁷ And *that* we can grasp and model with the easier-to-work-with possible worlds. The foundation of this approach is that the less information one has about the way things are, the more possible worlds are compatible with that information.

⁷ I started to take this view and its ramifications for the way I build models more seriously after a stimulating conversation with Rich Thomason last year. But of course he is in no way to blame for the direction in which it has taken me.

So I start here by assuming a standard possible worlds model, where **propositions** are sets of worlds. Then, I take a **situation** to be a set of propositions:

Assume a set W of possible worlds.

A **proposition** is a subset of W .

A **situation** is a set of propositions, the **propositions true of that situation**.⁸

If intuitively a situation is about a particular spatio-temporal locale, then the propositions true of that situation will only contain information about entities and eventualities in the relevant spatio-temporal zone.⁹ A **possible situation** s is one such that the intersection of the set of propositions true of it is non-null, $\cap s \neq \emptyset$; i.e. there is at least one possible world w in which all those propositions are true in w . Hence, unlike Kratzer's (1989) primitive situations, situations-as-sets-of-propositions are not parts of particular possible worlds. We might say that a given situation is "part of" every possible world in which all of the propositions constituting that situation are true. For any given situation s , I'll call the worlds in $\cap s$ its **realization worlds**; abbreviated **Real(s)**. For a given s and w , if $w \in \text{Real}(s)$, let's abbreviate this as $s < w$. Of course, the amount of information we have about a given situation s will be inverse to the number of possible worlds in $\text{Real}(s)$ —the more propositions true of a situation, the fewer the number of worlds compatible with that information.

One of the interesting and very useful features of Kratzer's conception of situations as I understand it is that two distinct situations might be co-extensive in terms of certain dimensions within the same world—for example, spatio-temporally co-extensive—and yet one might be "thinner" than the other in the sense that it has less information about that same locale. The thicker situation—say, one reflecting the knowledge of an omniscient being about that locale—might include information about what's behind every mountain and in the top of every tree, as well as of everything that occurs over the course of the temporal extent of the situation. While the other, thinner sub-situation might only include certain information about what was visible in the same realm from a given perspective, with the line of sight blocked by some of the mountains and too low to see what's in the top of every tree. Both situations are sets of propositions, none of them reflecting information about anything outside of the relevant spatio-temporal locale. But since the thicker, more information-rich situation s_{thick} reflects omniscience, the propositions in $\text{Real}(s_{\text{thick}})$ will be a maximal set—everything there is to know about this

⁸ As it stands, this would permit a situation to be a very disjointed entity, with one proposition about the salinity of the biggest lake in Nebraska, another about the atomic weight of helium, and the third about the age of Gary's mother when she married in Ware in 1948. But come to think of it, nothing prevents the situations of Barwise & Perry or Kratzer from being any less abstract. The reader is invited to imagine what might constitute adequate constraints on realistic, intuitive situations. But at least some of them would be easier to capture if we could talk about the situation in relation to the worlds of which it is intuitively part (those in which all the propositions that constitute the situation are true). For example, if there is a notion of space defined over all those worlds, then we could require a situation "in them" to be spatially contiguous in terms of that common notion of space, etc.

⁹ A number of important questions arise at this juncture that I cannot explore here for lack of time. For example, what about generic statements—is it (ever or sometimes) true in a particular situation that carbon has a certain molecular weight? See Kratzer (1989) for extensive discussion of important questions of this sort. Ultimately, any adequate proposal of the sort sketched here would have to address such questions.

situation will be reflected in a proposition in this set. The thinner situation s_{thin} will be modeled as a proper subset of those propositions, $s_{thin} \subset s_{thick}$; the fewer the propositions, the thinner the information in s_{thin} . Then in any reasonably rich model, there will be more worlds compatible with what is known about s_{thin} worlds compatible with s_{thick} , so that $\text{Real}(s_{thick}) \subset \text{Real}(s_{thin})$. If we then relax the requirement that the thicker situation include only information about the limited locale of the thinner situation, we can see that any situation s which contains merely partial information about a given locale will be a thinner version of any world w in $\text{Real}(s)$: In that case, $w \in \cap s$, but there will be propositions p which are true in w such that not only will p not be true of s , $\neg p \in s$, but p will not be entailed by the information we have about s , $\neg[p \subseteq \cap s]$. Thus, even though worlds are not situations in this theory, they are comparable to situations in terms of the information they contain.

Another useful feature of Kratzer's theory is that in it individuals are situations; hence, to model what it is for an individual to be *in* a situation, we take the individual to be a (proper) sub-situation of the locating situation. Kratzer calls individuals without any of their properties *thin situations*. I think we do have to take individuals to be primitives in the present framework, so we cannot say that they are situations. These primitive individuals might exist across different possible worlds or they might be world-particular, with counterpart relations between individuals in different worlds. But in those worlds and the situations in them where an individual (or one of its counterparts) plays a role, the individual might have more or few properties; intuitively, then, the individual would be more or less thick with properties. And, of course, the properties that make a given individual thick might vary (among its counterparts) from world to world. An individual that is singular in one world might even have twin counterparts in another; cf. examples like *If I had been born twins, my mother would have had double trouble* after Lewis (Possible Worlds). When a property is predicated of a primitive individual this yields a proposition; hence an individual with any thickness at all corresponds to a situation. Let us say that **an individual i is in a situation s** , $i < s$, iff there is information about i in s , i.e. iff there is some property P such that $\cap s \subseteq P(i)$ ('the proposition that (a counterpart of) i has property P is true of s '). We can talk about the **true information about an individual i in a situation s** , i_s , as those propositions true of s that are about i , i.e. entail that i has certain properties. Of course i_s , being a set of propositions, is itself a situation, and since all the propositions in i_s are true in s , $i_s \subseteq s$. Now we have a way of talking about having partial information about an individual in a given situation, since for individual i , situation s and world $w \in \text{Real}(s)$, it might well be the case that $i_s \subset i_w$. And if the information in i_s is properly partial relative to i_w , then $\text{Real}(i_s) \supset \text{Real}(i_w)$. This shows that the partial information we have about an individual in a given situation can be resolved in various, incompatible ways as we add additional information about that individual.

Further, in two situations realized in the same world w , there might be partial individuals who are not necessarily the same individual from the perspective of those situations themselves, but who are the same from the perspective of the larger world w . Some epistemic agent Jorge might be in a situation in which someone appears to be a spy—wears a trenchcoat and a slouchy hat, acts suspiciously, speaks with an Eastern European

accent, etc; call this situation s_{spy} . Unbeknownst to Jorge, in the actual world w the guy who has these properties in s_{spy} is his friend and neighbor Ortcutt, dressed up for Halloween, to whom Jorge attaches a completely different set of properties—being open, friendly, and trustworthy, wearing jeans and t-shirts, speaking with a Midwestern U.S. accent, etc.; call a situation in which this information is realized in the same world w s_{Ortcut} . Jorge has only partial information about both the spy and Ortcutt. As far as Jorge’s information goes, i.e. as far as he knows, both these situations are consistent with worlds in which it turns out that the spy-like guy is from Romania and the friendly guy has never been outside the U.S. Let i_{spy} stand for the partial information (a set of propositions) that Jorge has about i in s_{spy} and let i_{Ortcut} stand for the information that Jorge has about i in s_{Ortcut} . Then, although $i < s_{spy} < w$, and $i < s_{Ortcut} < w$ and (since by assumption the spy and Ortcut are the same individual in w) there is an i_w (a thicker individual reflecting all the information about i in w) s.t. $i_w < w$ and $i_{spy} \subseteq i_w$ and $i_{Ortcut} \subseteq i_w$, still it is not the case that $i_{spy} = i_{Ortcut}$. I.e. the partial information packets we have about this individual in the two situations are not the same thick individuals in the informational sense. As a corollary, in that case there will be some world $w' \neq w$ which is epistemically accessible to Jorge and is s.t. $w' \in \text{Real}(i_{spy})$ and $w' \in \text{Real}(i_{Ortcut})$ but there is **no** i' s.t. $i_{spy} \subseteq i'$ and $i_{Ortcut} \subseteq i'$ and $i' < w'$. In the terms of Landman (198?), there are worlds in which the partial information you have about the spy and the partial information you have about Ortcutt grow into information about one and the same individual, and other worlds in which they do not. In the latter cases, i correlates with or stands in a counterpart relation with two (or more) distinct individuals.

Using this way of characterizing information about a situation and about individuals in a situation, we can model the way that perspective in the intuitive characterization above limits access to information in a grounding situation: Think of the thicker situation as the ground, the thinner as reflecting access to information about the ground from the given perspective. One can imagine that one epistemic agent might simultaneously have multiple perspectives on the same situation, that is, be able to access information about that situation in multiple ways, some perhaps richer than others. Ranger Gary, having hiked every square inch of the Berkshires and climbed every tree, and having informants throughout the region, has at least indirect epistemic access to what’s behind every mountain and what’s in the top of all those trees, and knows everything that goes on. But being a human, his visual perspective at any given point and hence his direct perceptual access to evidence for what is taking place in the Berkshires are limited. He can’t see Kathy and her goat, because they’re on the other side of the barn from his angle of sight. Both his general epistemic perspective and his visual perspective might come into play in interpreting a single given utterance, as we’ll see.

Now we can characterize what a perspective is in terms of this notion of situation.

- (16) For situations s, s' , s is a **perspective** on s' if, for some s , *coord*, *origin*, *path*, *frame*, the following holds:
- *coord* an n -dimensional coordinate system over s , defines a comparable sub-space in s' , as well.
 - *origin* an n -dimensional coordinate within *coord*, lies within s .

- *path* a line on *coord*, extends from the *origin* through the space defined over *s*.
- *frame* a set of propositions, characterizes the limitations on access to information evident in any propositions *p* s.t. $p \in s'$ and $p \notin s$.

In this case, we also say that s' is the *ground* of *s*.

In a perspective as defined in (16), think of s' as a relatively thick situation, the ground for the perspective, and the perspective *s* as relatively thin. A given perspective may be more or less informative, so in the limit case of an omniscient being, at least when we limit the ground to the same locale as that of the perspective, roughly $s = s'$. In modeling an epistemic perspective, so that all the information in the perspective *s* would be true, *s* would be a subset of s' . But we might have an unrealistic or distorted perspective on a particular grounding situation, and in that case the modal accessibility relation which makes *s* accessible from s' would not be totally realistic in something like Kratzer's (notional category of modality) sense. Here, that would mean that there would be propositions in *s* that were not true in s' . Partly for that reason, I'm not sure how to define in the general case the relation that must hold between the coordinate system on s' and that on *s*, but intuitively, of course, it's important that this be the same system. Effectively, this system will have to be defined over at least the relevant portions of all of the realization worlds of *s*. In a spatial perspective, the origin is the location of the point of view from which *s* is evident in the defined space within the ground s' .

I borrow the notion of a path from Cresswell (1978), who develops a semantics for locative prepositions.¹⁰ His theory is not intended to capture the general notion of perspective, but only models the notion of point of view as that plays into the semantics of prepositions like *across*, *behind*, and *around*. For these prepositions, a hypothetical *path* through the relevant space from the point of view is used to characterize the relations they denote between participants in a described situation. For example, consider his:

- (150) Two days across the desert we ran out of water.
 (160) Arabella is standing across the meadow from Bill.¹¹

Cresswell's paths are functions from (an) individual(s), an interval, and a world to the set of spaces occupied by the individual across the interval in that world. For a point-of-view sensitive preposition like *across*, the path is defined as commencing at a given spatio-temporal point, the point of view; in (150), this might be the time and place of departure on a journey. In (160), we see the hypotheticality of the path—it isn't that either Arabella or Bill is *following* the path that would connect them across the meadow. Like the present theory, Cresswell's notion of point of view is not anchored in a particular index of interpretation (parallel to *speaker*, *time of utterance*, *location of utterance*, etc.). Rather, the notions of point of view, path, etc. are part of the lexical

¹⁰ Thanks to Sally McConnell-Ginet, who alerted me to Cresswell's work and its relevance for the current project.

¹¹ This is example is slightly modified from Cresswell's, in ways that are inessential to the point here.

semantics of the relevant prepositions, and their values are generally contextually given, although they may sometimes be indicated explicitly, e.g. by adding a modifier like *out of Khartoum* to (150) to indicate the origin. In (160) the origin would seem to be entailed to be Arabella, as a function of the locative implication of *stand*. In earlier versions of the model of perspective in (16), as in the intuitive sketch at the beginning of this section, instead of a path I talked about a vector, which intuitively corresponds to the line of sight in a visual perspective. But the notion of a path is more flexible than that of a vector, as I understand it; it might be curved or zig-zagged, for example, and is more useful for taking into account the temporal dimension. In visual perspective, one can take the path to be a vector, and I will sometimes use that term.

A perspective is defined here in terms of a sort of modal accessibility relation—albeit one between situations instead of directly between possible worlds. The frame makes explicit constraints on that accessibility relation—characterizes limitations on access to information from the given perspective. To get accessibility intuitively right, we have to define what information in the ground s' is accessible (perceptually or otherwise graspable) in the perspective reflected in the partial information in s , from the origin, given the path. In the canonical, spatial case, accessible situations will include a variety of sub-situations of the ground that are to either side of or along the vector modeling the line of sight. All this is a function of how human vision works, plus the physics of light, etc., and would be characterized by propositions in the frame. E.g., assuming the viewer is a human being looking out along the vector with visual access only to stuff in front of her within a certain angle on either side from the vector, the frame makes location of an object at a coordinate within the sub-space so defined a necessary condition on visual access to information about that object. Hence, the human perspectival accessibility relation over the perceived physical space—at least when it doesn't involve misperception—is an epistemic accessibility relation between situations such that the second, accessible situation is a sub-situation of the first, the ground.

Consider example (1), repeated here,

- (1) Tim and Margaret are sitting at a conference table in her psycholinguistics lab, working on a grant proposal. Tim is making notes on his Compaq-brand laptop.

Margaret: How do you like your Compaq?

Tim: It's not bad, but it's getting kind of old. I wish I had a Mac. Macs are far better for graphics, and it turns out that I'm doing a lot more graphics than I'd expected.

Margaret: [gesturing with her thumb over her shoulder, and slightly turning her head in the direction of her desk in the middle of the room]
I just got [that]_F last year.

In the direction Margaret is pointing there's a lot of stuff: a desk with a big pile of papers and a flat screen Dell computer monitor on it, past the desk an eye-tracker, past that a wall calendar.

Take the ground s' to be relatively thick situation fully characterizing the lab in spatio-temporal terms, and set up a Cartesian 3-D spatial coordinate system covering it. Take

the origin to be the coordinate where Tim-and-Margaret are located in the space (or, for the purposes of the deixis, Margaret alone, since she's the speaker and the one who sets up the vector). The vector, or path, is the line from Margaret's coordinate through her arm, extending into the space through the lab. The perspective *s* includes the area visible via a human eye located at the Tim-and-Margaret coordinate and looking out along the vector. The frame, or set of limitations on the perspective, is both physical—propositions reflecting practical limitations on the information in the perspective given by the extreme range of peripheral vision of the human eye, the walls and any other barriers, etc., and abstract—the range of phenomena that are relevant, given the question, or topic under discussion (which I'll abbreviate as the **QUD**). The actual (spatiotemporal) frame characterizes a pattern of micro-blind spots in the perspective—coordinates to which an individual located at the origin and looking out along the vector would not have perceptual access, as well as the macro-limitations given by the walls, etc. For example, although there's a wastebasket lying along the vector determined by Margaret's deictic act, as one looks along the vector from the origin the wastebasket lies behind the (opaque) desk and hence is masked off, in a blind spot. So we only have visual access to entities and to facts about their properties if their coordinates are visually evident from the *origin*. Since the wastebasket's coordinate is not visible from the origin:

- (a) The wastebasket isn't referentially salient given the perspective alone. Even if Tim knows there's a wastebasket behind the desk in the direction Margaret is pointing, she cannot refer to it with *that* or *it*, even though she's pointing directly at its coordinate. In order to refer to it she would have to say something like *the wastebasket* (if it's the only wastebasket in the lab) or (pointing toward its invisible coordinate) *that wastebasket*
- (b) Similarly, given the perspective defined over the visual space alone, the interlocutors do not have access to the fact that the desk is in front of the wastebasket.

Now consider the role of the QUD in establishing the frame of reference in (2). Along the vector suggested by Margaret's gesture lie a number of inanimate objects that could be referred to via *that*: the desk, the pile of papers, the eye-tracker, etc. But it seems clear in the context of the discussion that Margaret intends to refer to the Dell computer because the topic of discussion is about what properties the interlocutors' computers have. So in view of the QUD, the computer is the most relevant object lying along the vector. I have in mind the notion of **relevance-to-a-question** defined in Roberts (1996):

- (i) An utterance is Relevant in the context of utterance just in case it addresses the question under discussion in that context.
- (ii) An utterance *u* addresses a question *q*, iff *u* either contextually entails a partial answer to *q* (*u* is an assertion) or is part of a strategy to answer *q* (*u* is a sub-question of *q* or an imperative whose realization would help to answer *q*).

As an extension of that notion, we can roughly characterize what it is for an *object* to be Relevant (Roberts 2006):

- (iii) An entity is Relevant in the context of utterance only if the issue of what properties it has bears on the truth of one or more of the possible answers to the question under discussion.

The space of possibilities opened up by a question can be characterized as a partition over the set of worlds corresponding to the possible complete answers to the question (Groenendijk & Stokhof 1987). Roberts (1996), Ginzburg () have argued that the QUD is an especially important organizing factor in discourse. This is because it establishes a goal and calls forth particular intentions on the part of those who agree to discuss it: to sort through the alternatives and eliminate all but one, if possible. Adopting a goal leads us to focus our attention on achieving it. In the case of a question, this involves attending to the logical partition induced by the question, and to those propositions and entities that are Relevant in view of the QUD. Humans tend to filter out irrelevant information, in order to be more efficient at processing the relevant. Just so, in attending to discourse while scrutinizing a perceptually accessible situation, one tends to mask out those entities in the situation that are not Relevant to the QUD. So, though in (1) the eyetracker is actually much larger than the computer screen, and the desk chair and pile of papers are closer, Tim is unlikely to take them into consideration in determining what Margaret intends to indicate with her gesture.

Gawron (1996) argues that domain restriction is determined by the QUD. There are many kinds of examples where this seems like a reasonable description of what triggers domain restriction, as in:

- (170) A: What do you want me to do with these flowers?
B: Every stem needs to be cut before it's put in the vase.

The QUD in (170) is about what goals B would like A to adopt with regard to the indicated flowers. Hence the flowers are not only perceptually salient, A having just brought them to B's attention, but Relevant. One might say that the inquiry suggested by the question brings the state of the flowers and their future under scrutiny.

But (1) suggests that this is too narrow a conception. Certainly the QUD plays an important role, but it is not generally sufficient by itself to determine the intended domain. Similarly, in the fire station example in (1), presumably the QUD is something along the lines of *How do I get to the Clintonville farmer's market?*. But the fire station in question is not in any obvious way directly relevant to that question. Instead, it is simply a readily identifiable landmark along the path suggested by the speaker.

Since the QUD will generally put constraints on what interlocutors will attend to in a discourse, we can capture what is right about Gawron's hypothesis by making the QUD play a role in defining the frame. I would be hesitant to claim that Relevance puts an absolute constraint on referential accessibility, but I do think it tends to serve as a default condition. Let's assume that the QUD induces an order on the set of entities accessible under the remaining conditions in the frame, such that those that are more directly Relevant to the QUD—either in that they reflect facts that are Relevant or contain

Relevant entities—are more highly ranked. In the terms of Kratzer (Notional Category), this amounts to adding an Ordering Source over the (e.g., epistemic) Modal Base given by the remainder of the Frame. Then we can say that some aspect of the grounding situation is more likely to come to the attention of the addressee the more Relevant it is to the current QUD, and characterize the perspective partly in terms of being closest to the Relevant ideal. By “aspect” here, I mean a sub-situation of the ground; since thick individuals are situations in the theory proposed here, this subsumes the Relevance of both sub-situations (by virtue of the facts true in them) and the individuals in them. Since, as we will discuss below, perspectives only serve to make some individuals more accessible than others, and multiple perspectives may be simultaneously available, this role of the QUD does not make it an absolute restriction on the domain of any given operator.

One more technical point about the implications of the framework in (16) and its role in interpretation. Though I don’t believe he intended to talk about the role of partial situations in understanding perspective, Cresswell points out (1978:37, fn.4):

In dealing with ordinary language, the spatio-temporal framework...may have to be relativized to an area surrounding the speaker. For instance if someone wants to talk about the relative positions of things in an aeroplane, he is likely to ignore the fact that all these things are moving with the aeroplane. For the purpose of evaluating his sentence we assume a spatial framework in which the aeroplane is constant. Of course if the sentence relates something in the aeroplane to something outside it the speaker will have to assume a wider framework, and so on until, when discussing astronomy, we may be using the whole of space as our framework.”

The implications of this are quite interesting. For example, consider (17):

- (17) [On a trans-Atlantic flight inside Cresswell’s aeroplane, one passenger asks another if he would mind watching her bag while she goes to the restroom, promising him to hurry back. He replies:] No hurry—I’m not going anywhere.

I think from the perspective of the speaker, as well as that of the addressee, limited as it is by the walls of the plane, (17) is true. On such flights one is stuck in one place *in the plane* for hours on end. But of course, if we take a broader perspective, say from a space station circling the earth, this isn’t true at all. And certainly, if when the lady returned she asked the gentleman where he was going, he’d reply “Barcelona”, not “nowhere”:

- (18) A: Where are you going?
B: Barcelona

First, I would agree with Kratzer (1989) that it is desirable to treat propositions as persistent, so that if a proposition is true in some situation *s* in a given possible world, it is true in all supersituations of *s* in that world, including the world itself. In a framework like hers where situations, and not necessarily whole worlds, are the circumstances of

evaluation, it is basically necessary to stipulate persistence. Otherwise, the same proposition that is true in the situation in the cabin in (17) could become false in the larger trans-Atlantic situation in (18). Persistence is not a necessary feature of such a semantic framework. But if we only consider whole worlds as circumstances of evaluation, as I have in mind here, then in order to avoid the conclusion that one and the same proposition is both true and false relative to the same possible world, we must relativize the proposition expressed in examples like (17) and (18) to the perspective assumed by the speaker. I.e., we are led to assume that *I'm not going anywhere* expresses different propositions in the two contexts of utterance. Hence, on the assumption that part of what it is to be a *possible* world is to reflect a consistent information, persistence falls out of the present framework. This is in itself neither good nor bad. But it has other possible ramifications that might weigh in favor of one or the other approach, and I'll mention a couple of these in the final section.

Second, Cresswell's observation underlines how multiple perspectives may be simultaneously available in the same circumstance. In each of the utterances on the plane, the speaker is in the same location at roughly the same temporal interval, with pretty much the same macro-intentions—sitting still in one place long enough to get from NY to Madrid so he can make his connection to Barcelona. The origin and perspective Relevant for the question the lady asks in (17) are both limited to the inside of the cabin. Adopting this perspective,¹² the gentleman answers with the true proposition 'I am going nowhere relative to other objects inside of the plane'. The grounding situation and perspective Relevant for her subsequent question in (18) are external to the plane, including the larger path of the journey the plane trip is part of, from which perspective the man is going to Barcelona.

I think this is one of the most striking arguments for the importance of situations in interpretation that I have seen. This doesn't mean that we have to use situations as the circumstances of *evaluation*, but strongly argues that they are truth-conditionally relevant for modeling how context influences interpretation. The information we take into account has a bearing on what we mean by a given utterance. In this light, one could explore the possibility that the argument for situation *semantics* in Kratzer (1989) is itself, actually, an argument for a role for situations in *pragmatics*: For example, the central argument for her situations in that paper lies in the role that they play in defining the notion of *Lumping*, which is in turn used to restrict the domains of counterfactual modals (and negation), resulting in a more adequate theory of counterfactual accessibility than that of Lewis (). If the thesis of the present work is correct, domain restriction is essentially a pragmatic phenomenon in that, though it directly influences truth conditions, the information involved is essentially contextual in origin, not in general given by compositional semantics.¹³

¹² Note that *go* is perspectival, so arguably it conventionally presupposes an origin in a perspective; see the following section.

¹³ I would emphasize that this claim does not bear on the question of whether domain restriction always involves the binding of variables at LF in syntactic structure (Stanley & Szabo 200?). Rather, it's a question of where the information *comes from*.

There is a great deal more to be said about the notion of perspective modeled in (16), and in particular about the role and character of the frame. Some of this will come up in the discussion to follow. But I would like to point out one characteristic of perspectives as I have defined them, a characteristic they share with Kratzer's situations: They are not psychological, but informational (epistemic, or quasi-epistemic) notions. This is in contrast to the way that notions like perspective and point of view are treated in the work of Perry () and Mitchell (), and in related notions in the work on specificity of Breheny () and of Kamp & Bende-Farkas (). In that literature, as I understand it, notions like propositional attitudes and perspective are characterized in terms of relations between a situation and an agent's internal *representation* of it, a sort of mental model. The present approach is a-representational. Certainly, one might take the information accessible to an agent at the origin of a perspective to be somehow reflective of that agent's psychological state. But as a semanticist, that would be only conjecture on my part. I am merely attempting to model various kinds of constraints on the flow of information in a discourse, as these bear on compositional, truth conditional interpretation. Insofar as there is empirical evidence *in the linguistic record* that the kinds of information in a perspective, as modeled here, regularly come to bear on interpretation, this is a linguistically interesting analysis.

§3 Perspectival Anchoring

In this section I will briefly explore some of the ways that the notion of perspective just modeled comes to bear in the course of interpretation in a natural language like English. This is by no means a full survey of the kinds of effects or the nuances that arise in their interaction with other elements in an utterance. I have selected a few that I think illustrate some of the most interesting facets of this web of phenomena.

In what follows, I will characterize as **perspectival** an expression or utterance of an expression that presupposes a perspective, in the sense of perspective defined in the preceding section. Some feature or features of the perspective then serve either to restrict the domain of the expression (if it's an operator) or to give its implicit argument(s) (if it's a functor). Thus, the role of perspective in interpretation is fundamentally contextual, pragmatic. Like the antecedents of anaphoric presuppositions, the intended perspective must be retrieved in order to determine the proposition expressed on a given occasion of utterance. But there is no need to treat perspectives as some distinguished aspect of context, or as a unique feature of a dynamically involving conversational record, like Lewis' scoreboard or Partee's (1984) Reference Time. Compare perspectives, and situations more generally, to anaphoric antecedents in the sense of Roberts (2003). The latter are neither NPs (or other linguistic expressions) nor entities in the world, but discourse referents in the sense of Karttunen, Kamp (reference markers), or Heim. They are informational entities. For example, if two interlocutors know that there is a dog in the hallway, we might model this by saying that their Common Ground has a particular kind of existential entailment: in all worlds in the corresponding Context Set, there is at least one entity that has the property of being a dog and the property of being in the relevant hallway. We can then think of this information as a familiar discourse referent,

and its existence in the Common Ground—however that information was introduced—licenses the use of expressions that presuppose familiarity, like *the dog* or (if sufficiently salient) *it*. Similarly, perspectives are kinds of contextual information that interlocutors have access to—situations, distinguished sets of propositions. Like a discourse referent, a situation, can be familiar to the interlocutors, and various familiar situations can be more or less Relevant, more or less salient, at different points in a conversation; hence, like other discourse referents they license anaphora (e.g. to *it*, *this*, or *that*). Finally, although perspectival presuppositions are sometimes conventionally triggered—e.g. inevitably triggered by a conventionally perspectival lexical item, they can also arise as non-conventionally triggered speakers’ presuppositions (Roberts 1995). I will illustrate some of this in what follows.

§3.1 Perspectival presuppositions

There are a large number of lexical items whose use seems to conventionally assume an intended perspective. One of the best-known is *local* (Mitchell 1986, Partee 19??). (19) has at least three distinct truth conditional interpretations, depending on the perspective that is taken to anchor *local*. Three different perspectives for the interpretation of (19) are suggested in (a) – (c):

- (19) After the game, every fan celebrated at a local bar.
- a. in close proximity to the location of any given fan
 - b. in close proximity to the location of the game
 - c. in close proximity to the location of the speaker

Suppose that the game is a national play-off. There are fans everywhere. Then reason tells us that the probable reading is (a). If instead it’s a game between farm-league baseball teams, not even broadcast locally, then reading (b) is probable. One has to work a bit to make reading (c) the most likely, without assuming that the speaker is at the ballpark, in which case we cannot distinguish it from reading (b). Assume that both interlocutors both live in the same part of Columbus, German Village. They know that all the city league teams play in a park built on sanitary landfill, Berliner Park, a couple miles to the south on the edge of the city by the freeway, with nothing near-by. German Village has lots of bars and restaurants and is off the freeway on the way back up into town. Then I think we can get reading (a). Take (19) to be the answer to a QUD *What happened after the game last night?* (without the then-redundant temporal adverbial). This QUD by itself wouldn’t suffice to derive the intended perspective in any of the scenarios developed. Other factors come into play, in a complex reasoning process involving plausibility on the basis of world knowledge.

In keeping with these glosses, take *local* to mean something like ‘in close proximity to the origin of’. This is an irreducibly relational interpretation. In order for an addressee to understand the proposition that a speaker intends to express in uttering (19), she must understand the perspective from which the bar in question is local. One can express the relatum, the origin of the intended perspective, via an explicit *to* prepositional phrase

argument as in (19a-c), or via a variety of other restrictive or appositive modifiers; e.g. *a local bar down the street from the stadium*. But when the relatum is not explicit, then the missing semantic argument must be retrieved in order to interpret the utterance. One can argue, thus, that when the perspectival argument of *local* is not made explicit, this amounts to a speaker's presupposition that the origin is contextually retrievable.¹⁴

Besides explicitly giving the perspectival argument, one can indicate the intended perspective via a frame adverbial:

- (20) Whenever we went to German Village, we found a local bar that had good jazz.
- (21) Shortly after the fans poured out of the stadium, they were all celebrating in a local bar.

Even here, with the adverbials giving clear hints about the intended perspective so that the missing origin argument of *local* can be retrieved, Relevance plus practical reasoning are involved in determining the intended perspective and its origin. We saw this in the original explication of how (19) might come to have its attested interpretations in different kinds of context. In fact, Relevance and reason are generally the principal guides to retrieving elided material in most uses of perspectival expressions, as with elision more generally. *Local* seems to require an argument that is presupposed to be the origin of a relevant perspective. In examples (150) and (160) in the previous section, the uses of *across* under Cresswell's analysis both require an implicit path, hence also seem to presuppose a Relevant perspective. *Through* presupposes a space (*coord*), as well as a path, but is not anchored in an origin the way *local* and *across* are. Hence, different perspectival expressions can presuppose different features of some salient perspective.

If we want a satisfactory theory of how perspective bears on interpretation in all these examples, it behooves us to say something about when and how a given perspective is made available to interpretation.

Some theories characterize domain restriction (Barwise & Perry, Gawron, Kratzer '05 *inter alia*) or perspectival anchoring (Partee) as a relatively global phenomenon: A particular resource situation, QUD or topic situation suggests domain restriction. Or a "richly structured presupposed spatial/motional context" (Partee) is available for binding the implicit variables in an entire clause, including the implicit origin in bound occurrences of *local* like that in (19); this context is shifted dynamically through the course of interpretation, like a richer cousin of the notion of the Reference Time in Partee (1984) or Lewis' () conversational scoreboard.

¹⁴ Again, see Roberts (2006), for extended discussion of how elision leads to RETRIEVABILITY presuppositions (the notion capitalized to indicate that it has a theory-internal definition). I would not want to claim at this point that *local* is itself a conventional presupposition trigger. The argument in that paper is that any semantic elision signals a *speaker's* presupposition of RETRIEVABILITY, on the assumption that the speaker is cooperative and rational. What's conventional about the presupposition arising with *local* and several other dedicated perspectival expressions is that they take a perspectival argument—origin, path. This, in turn, presupposes that there is a perspective of which the argument is a component. But Simons () has argued that this type of presupposition is really a speaker's presupposition arising from a lexical entailment. I won't attempt to resolve that here.

It is sometimes unacceptable to anchor elements in the same utterance to distinct origins, and *prima facie* these might be taken to argue for the global perspective. Here is a nice example from Partee (1989):

- (26) In all my travels, whenever I have called for a doctor, one has arrived/set out/*departed within an hour.

I agree with Partee's acceptability judgments here, but don't think the example argues for a global view of the role of perspective. My take on the unacceptability of *departed* goes as follows: (a) the frame adverbial *in all my travels* establishes the speaker as the individual at the origin of a journey, or path through space-time; otherwise, it would be difficult to understand how this adverbial could be relevant to what follows. (b) In any event instantiating the temporal frame *whenever I called for a doctor*, being called for conventionally establishes that the doctor is distal relative to the caller; hence, (c) the doctor is not at the origin. Therefore, the doctor may *arrive*—i.e. come to the origin, or *set out* (a description of departure which is neutral with respect to whether the agent is at the origin), but he may not *depart* from his distal location at the time of the call, because *depart* is an origin-oriented perspectival expression: Its agent is conventionally anchored to the origin of a presupposed perspective. The speaker is at the origin and the doctor leaves from a location that is distal to the speaker; hence the frame adverbials tell us that the doctor is not at the origin. But since he is the agent of the departure, *depart* effectively presupposes that he *is* at the origin. This is a conventionally triggered pragmatic contradiction, explaining the anomaly. The anomaly in this utterance thus seems to depend on the conventional nature of the various perspectival contributions: The special role of frame adverbials establishes a perspective which by Relevance we take to bear on the interpretation of the main clause. Then *departed* conventionally presupposes a perspective, and entails that the agent of the departure is anchored to the origin of that perspective. There doesn't seem to be any reasonable way to yield a coherent interpretation of the utterance, including the adverbials, while assuming that distinct origins ground *called for* and *departed*.

In many examples multiple perspectives can come to bear on interpreting a single utterance. In (27), the perspectival expressions in boldface can be anchored independently to one of three individuals, as indicated:¹⁵

- (27) John thinks his boss prefers a **certain** type of **local** beverage.
- a) Can the epistemic origin of *certain* be anchored to:
- the speaker? yes. The speaker's epistemic ground is always available (even in counterfactual contexts).
 - *John*? Yes, since John's epistemic ground is made relevant by *thinks*. The speaker might not know any identifying features of this beverage, just presumes that John does.

¹⁵ This is similar to an example examined in detail in Breheny, who claims it has 9 construals. But I cannot figure out what all of them are supposed to be.

- *his boss?* no. *certain* is inherently epistemic, so must be anchored to a locally salient epistemic ground. There are two in this context, the speaker's and John's, as above. Preference is bulletic (pertaining to wishes or desires), not doxastic/epistemic, so the predicate does not make the boss' epistemic ground salient.
- b) Can *local* be anchored to:
- the speaker? yes
 - John? yes
 - his boss? yes. This is easier to see if we imagine that John, an American, works for a Taiwanese company in their Athens office, home of ouzo. His boss lives and works in Taiwan, where one can obtain a special tea. The speaker is reporting this from his home in Scotland, home of scotch liquor. Now there's the greater pragmatic plausibility of three distinct locations—the speaker's, John's, and the boss's. *local* is not epistemically anchored, but anchored to one of these three pragmatically salient spatial locations.
- c) Can *local* be anchored independently of *certain*?: I think so. Suppose that the American John has noticed that when the Taiwanese boss visits Athens, he seems to favor a particular tea-colored beverage, but John doesn't know what it is or where it's from. The speaker in Scotland knows that it's scotch. So the intended perspectival anchor for *certain* would be John's (for though the speaker knows the identity of the liquid, too, he intends to convey as well that John has a particular liquid in mind); but the anchor for *local* would be the speaker's.

Though there is a tendency for a single perspective to dominate at any given point in the discourse, especially with frame adverbials making it relatively explicit, there may be multiple Relevant perspectives in a relatively short stretch of discourse (as we might expect, given that they are fundamentally epistemic accessibility relations, and many such relations typically show up in a single discourse). It seems, for example, that the concrete situation of utterance is always salient and available to serve as a perspectival ground, and this perspective can intrude readily into utterances where another is otherwise dominant. This is the case on the reading of (27) where *certain* is anchored to John, but *local* to the speaker.

Further, the present theory differs from the others cited in taking what grounds domain restriction to be not just a situation (or a particular type of "resource situation", a context of utterance), but a richer notion, a perspective. Examples like the following, from Soames (1986) and McConnell-Ginet (2005) may be taken to argue for the richer notion:

- (28) Sleep scientist enters lab where her assistants are bustling around taking readings from the research participants:
Great. **Everyone** is sleeping.

In this example, the domain restriction is a function of the frame, in this case the QUD, which has to do with the state of the experiment, hence with its subjects. The alertness of the scientist and her assistants is not relevant to the QUD, hence they are masked by the role of the QUD (Relevance) in establishing the frame for the perspective adopted.

- (29) Host at dinner table in posh establishment around which many (hungry and unfed) staff are working to keep plates and glasses full for those seated at the table:
A last, everyone is eating.

As with example (28), here the host's concern is with her guests. The servers are like the little guys in black in Kabuki Theater—basically invisible, except insofar as their work bears on the guests' comfort and contentment. Another way of putting this is that the narrative of interest does not include the staff, even though they're physically present.

McConnell-Ginet (2005, p.c.) also emphasizes that *everyone* quite often is interpreted in such a way that the speaker is not part of its domain, even though the speaker is very much in the context of utterance:

- (30) *Everyone's looking at me.*

Assume that the speaker is at the origin in her situation. From the origin, the origin is not perceptually evident: E.g., if I'm at the origin and looking out, I'm not in my line of vision. Hence, in the typical perceptual or perceptually-based perspective, the speaker is not relevant/included.

McConnell-Ginet (2005) reminded her audience that though the idea of a topic situation may be useful and appealing in some respects, one would not want to assume that it brought about a global Restriction over a discourse or even over a sentence:

- (29) Sweden is a funny place. Every tennis player looks like Björn Borg, and more men than women watch tennis on TV. But most people really dislike **foreign** tennis players.

Westerstahl, cited in von Fintel 1998

Here, *foreign* is conventionally perspectival. If we assume that it takes as argument the origin of a presupposed perspective, meaning 'foreign to', this conventionally triggered presupposition would take precedence over the tendency in this sentence to presume that all Relevant entities are Swedish, established pragmatically by the first utterance. This again illustrates how, though prior context and frame adverbials tend to foster a principal relevant perspective, the perspective against which a given expression is interpreted may be retrieved via complex reasoning based on lexical meanings, context, and world knowledge.

Moreover, even when there is a clearly dominant perspective in discourse, interlocutors still have access to the full range of information about circumstances beyond that perspective—for example, other circumstances and individuals about which they have

information in their Common Ground. They can expect their interlocutors to pick up intended reference to entities not in the grounding situation *so long as they give a sufficiently rich description* which makes it clear what they intend. Because of this, domain restriction cannot consist in simply (temporarily) shrinking the domain of the model itself in the course of interpretation. Lewis (1979) makes that point with the following example:

- (i) The cat is in the carton. The cat will never meet our other cat, because our other cat lives in New Zealand. Our New Zealand cat lives with the Cresswells. And there he'll stay, because Miriam would be sad if the cat went away.

In the first sentence, the speaker is evidently talking about a present situation, the default being the situation of utterance (though there are other possibilities, where a current situation is being described at a distance, etc.). Out of the blue,¹⁶ the use of definite descriptions in the first sentence implicates that there is a unique relevant cat and a unique relevant carton in that portion of the situation under discussion that is perspectively accessible from the presumed origin. For example, take the situation of utterance to be *s*, and *s'* to reflect all the information perceptually accessible to the interlocutors from their more-or-less-common origin (and perhaps optimally Relevant, given the QUD); hence *s'* and all the entities in *s'* are familiar and perceptually salient at this point in the discourse. The first utterance in (i) is felicitous just in case there is exactly one cat and exactly one carton in *s'*. But as the description and following discourse make clear, *our other cat* does not refer to an entity that is within that situation, let alone perspectively accessible. If the ground situation for the speaker were in New Zealand, then presumably the *because*-clause would not be a reasonable explanation for why the two cats will never meet. Moreover, it seems clear that the only reasonable antecedent for the demonstrative *there* is the sole location mentioned, *New Zealand*; and since *there* entails that the location referred to is distal to the origin, the speaker must not be in New Zealand. Moreover, the New Zealand cat is referred to repeatedly as *our other cat* or *our New Zealand cat* throughout the second and third sentences, the richer description required to provide adequate differentiation from the unique perspectively accessible cat. But in the conditional complement to *because* something shifts. In particular, if one knows that *Miriam* refers to Professor Cresswell's daughter, Mariam Meyerhoff, then in the *if*-clause *the cat* must refer not to the cat in the carton, but to the speaker's other cat, the one who lives with the Cresswells in New Zealand. Hence, though neither cat is unique in the world, both can be referred to with the simple description *the cat* within the space of a four-sentence discourse.

What licenses the use of *the cat* at the end to refer to a different animal than the one referred to at the outset? The key is that Miriam's sadness invites the addressee to consider the matter empathically from her point of view, where presumably what would make her sad is not the carton-cat going away, but the cat who lives with her. As indicated by the subjunctive mood, the perspective is from a hypothetical point of view—that of Miriam in New Zealand under the irrealis circumstance that the cat who lives with

¹⁶I.e., in the absence of other contextual means from prior discourse for framing out other, irrelevant cats and cartons—see discussion in Roberts (2003) and further below.

her has gone away. But notice that the perspective of the speaker is still evident, both in the hypotheticality of the mood and in what the choice of demonstratives would have to be: *Miriam would be sad if the cat wasn't there*, etc.¹⁷

Lewis' point was this: To make it felicitous to refer to the carton cat with a definite description that seems to entail its being the unique cat, we cannot actually restrict the domain of the model so as to temporarily eliminate all other cats. If we did the latter, we could not then refer to *our other cat* in the same sentence, let alone later refer to the latter as *the cat*. The approach to domain restriction suggested here does nothing to shrink or restrict the domain of the model or the domain of familiar discourse referents. Insofar as interpretation is concerned, interlocutors can interpret relative to entire worlds and, as the example shows, keep in mind all that they know about the larger world and the entities therein, which may include other cats, familiar or unfamiliar, outside the relevant situation under discussion. Instead of literally restricting the domain, the speaker suggests a perspective from which only certain entities are evident and relevant, and expects the addressee cooperatively adopting that perspective to be able to retrieve the suggested perspectival, pragmatic domain. When a definite's descriptive content requires us to consider entities outside that perspective (as does *other* in (i)) or the dominant perspective shifts (as to that of Miriam in New Zealand), then other familiar entities are referentially accessible.

This reminds us that what comes to mind most readily, what we can refer to most easily and naturally, is what is evident in the space to which we attend—both actual space/time and the metaphorical space of possible candidates for the way the world is that are Relevant to the QUD. It really isn't that the domain of the model or the Common Ground is restricted, so much as that the perspectival domain—that of the dominant perspective—is highlighted by its being evident and salient. The role of the QUD is natural, given our commitment to pursuing its resolution and the fact that we will naturally attend to what is relevant to pursuing an intention. That brings focus and clarity to our actions and perceptions (literal and extended), and to the (individual and common) information we are drawing on and updating. One's attitude (in the original sense) to the world, thus guides the organization of information about that world and thereby both facilitates and (relatively, temporarily) limits or dampens access to such information.

¹⁷ The shift is facilitated by the fact that the third sentence shifts the topic of conversation (or Center) from the carton cat to the New Zealand cat, but I don't believe that this would suffice to make the latter sufficiently more salient to permit the simpler description. Consider my own interpretation of this sentence when I first encountered the example as a graduate student back in the 1980s. I didn't know that *Miriam* referred to Professor Cresswell's daughter, and in fact guessed that it must be a reference to Mrs. Lewis. In that case, assuming Mrs. Lewis lived with Mr. Lewis, and hence not in New Zealand, I got a completely different interpretation. The reader is invited to consider that possibility herself before proceeding to consider my own interpretation.

I took it that it was the carton cat whose hypothetical departure would make Mrs. Lewis sad. Now what would that have to do with the New Zealand cat? Having had some cats myself, I assumed that the carton cat didn't welcome other felines into their home, and that accordingly they had not brought the New Zealand cat back with them from sabbatical because Mrs. Lewis didn't want their New Jersey cat to be unhappy and run away from home. To me, this all made sense, even though it involved a Centering Shift back from the New Zealand cat to the carton cat. What I couldn't figure out was what the example was telling us about domain restriction, since the first and last *the cat* seemed to refer to the same animal.

§3.3 The attitudes and interpretation *de se*

From where I sit, I can only see a very limited range of my own properties. I know I have on a grey skirt and a black blouse, that I sit at a certain height behind the desk and that my fingernails on the keyboard need trimming. But I can't see my face. In fact, I have never seen my face and I never will. I can see its reflection, pictures of it, and I find these things fascinating in a covert sort of way. But rays of light have never bounced off the surface of the front of my skull directly onto the back of my retina.

There are lots of other things I don't know about myself. I wish I could read this paper with another's eyes, to see where I've been clear and where the assumptions I've made or the way I'm leading the reader through the material is confusing. Sometimes one can read something one has written years later and wonder at what one wrote, read it afresh without remembering exactly what's coming next. But then at that point one is no longer the person who wrote this article.

There are two kinds of properties about myself that I don't directly know: Those that I am aware I don't know, and those that I'm not aware that I don't know. (I don't know whether this type of awareness constitutes knowledge. I sort of doubt it—more a kind of attunement to a pattern of (lack of) information. But I'm not an epistemologist, so I'll try to be neutral in my description.) I'm not aware that someone has stuck a post-it note on my back that says "Kick me!"; I'm not aware that the paragraph I wrote an hour ago is confusing to you; and I'm not aware that I'm located 10,531 miles from the most recently born zebra on earth. But I'm aware in an on-going way of my ignorance about some other things: I'm aware that I don't know what my face or my left shoulder-blade look like. I'm aware that I don't know the shape or size of my shadow when I face east at dawn (though I can conjecture or take a digital photo). I'm aware, partly because it sounds funny on a tape-recording, that I do not know what my voice sounds like to others. And given my experience navigating in the world, I'm aware that I do not have immediate perceptual access to whatever is behind that opaque chair over there (though I'm pretty sure that's where I put the wastebasket).

This latter kind of (informed) ignorance is very important. It is essential to our awareness of location. A big part of what characterizes me to me is this epistemic zero at the origin of my perceptual field as I move through various situations. I maneuver in space and time in terms of the pattern-of-lack of information and related patterns-of-distortion that, when used in conjunction with a richer ground of information, give me perspective, a sense of direction, and spatial orientation. As I'm driving through the Nevada desert, I don't have direct perceptual evidence of myself, or hence of where I am in the larger scheme of things. But I'm aware of the pattern of my ignorance (without necessarily noticing it as such); I see the reflection of the mountains behind in the mirror, the wide screen of the automobile windshield ahead, and through the windshield what I believe (on other grounds) to be parallel edges on the highway ahead appearing to get closer to each other at a certain rate. I know in the midst of these clues and systematic distortions that

I'm the entity I cannot see. These things help me to gauge where I am in space relative to the farthest point I can see on the road, and together with the information in a road map, they help me to situate myself in Nevada. If I want to go to Ely, it doesn't dismay me that I cannot see it, because I know that given where I hypothesize myself to be located, that mountain where the road veers off to the left on the windshield screen would be between me and Ely to the northwest, and hence I *shouldn't* have direct perceptual access. So the patterned lack of information in this perceptual perspective is helpful.

Hawthorne & Scala (2000:199) talk about a creature who is orientationally more challenged than humans:

A spherical being floats through space, exercising its capacity to perceive and think about the world. Its entire surface is a visual sensor. In an obvious way, then, it enjoys an advantage over more familiar perceivers. It can at a moment see itself as encircled by predators. Familiar binocular vision provides no such capacity. However, as we are conceiving it, our fictitious—yet possible—creature is in other respects deficient vis-à-vis more familiar perceivers. While being able to perceptually represent various distance-theoretic and topological facts, it is wholly incapable of perceptually representing direction-theoretic facts. It can see a spider as close or distant. It can see itself as being in between two spiders or as encircled by spiders. But it cannot see a spider as being to the left or to the right, up or down, front or back.

This creature has a point of view but no unique visual vector to orient itself in the space in which it occurs. So there are things it can see “behind”, but it doesn't know which way to go to get to them! When I'm driving, I can differentiate between what's accessible through the windshield and what through the mirror, and hence tell where I'm going and where I've come from.

Here's another way of saying how we use a map to self-locate: Being perceptually aware of one's location requires patterned ignorance—a certain kind of lack of direct information, and awareness of that pattern. But this does not suffice to know where you are in a richer, less subjective sense. The latter involves the knowing juxtaposition of two perspectives on the same ground of experience: the first, a rather thicker propositional map of the ground situation (perhaps, in the case of an omniscient being like Ranger Gary, as thick as the ground itself), overlain with the second, a thinner perceptual perspective on the same ground, with—crucially—an awareness that these are both perspectives on the same ground, so that the same coordinate system obtains in the two perspectives as in the ground. Having epistemic access to both and knowing that they both have the same ground, one uses the pattern of ignorance in the perceptual perspective, in conjunction with the less direct information in the prepositional map, drawing inferences about their relationship that permit one to apply the you-are-here arrow to the map.

In fact, it is just this patterned ignorance that David Lewis' (1979) two omniscient gods are lacking! Remember?: One lives on the highest mountain and throws down

thunderbolts, while the other lives on the sunniest mountain and sends down manna, or something like that. Both have full propositional knowledge about their worlds. But apparently that's all, and precisely because of this, they cannot tell themselves apart: Neither knows which of the two gods he is because he lacks the *ignorance* associated with a perceptual perspective.

I think a related problem is at the heart of Rudolph Lingens' dilemma too. Perry (1977) tells us that Lingens is an amnesiac who doesn't know who he is or where he is. He's actually in aisle five, floor six, of the Main Library, Stanford. It's a wonderful library, and there he has full access to all the propositional knowledge one might want, including full information about a fellow named Lingens who's located in aisle five, floor six, of the Main Library, Stanford on that date one thousand nine hundred and seventy-seven years after the purported birthdate of the prophet Jesus. Moreover, Lingens has a detailed map of the library and is fully aware of how one could, in principle, exit from aisle five, floor six. But still, like Lewis' gods, Lingens is disoriented. He may be propositionally omniscient about himself and the circumstances in which he is located, and assuming that he's perceptually competent, he does have access to and awareness of the patterned ignorance associated with being *perceptually located at an origin* in the concrete situation in which he finds himself. But he does not know that the two perspectives have the same ground, floor six of the Stanford Main Library. What he really needs is access to a computer terminal where he could link up to Google Earth and find a direct feed from a panoramic webcam located on aisle five, floor six of the Stanford Library, one moreover *labeled as such*, so that it could be correlated with other kinds of information in his propositional map about that location. Given the labeled webcam image, an ability to match that information with the pattern of his immediate perceptual experience, and the map he's put together from his propositional information, he could figure out where he* was.

One might argue that all Lingens needs is to be able to pattern-match the labelled webcam images to his perceptual experience, and that then he would know that the label applied to his location, as well. Since there are different senses in which one might be said to know where one is, I could concede this to be one of them. But insofar as knowledge-of is empowering, just knowing the name of one's location wouldn't be a very satisfying, or deep kind of knowledge-that because it wouldn't help Lingens to understand in the relevant way *who* he is. The knowledge that would be really useful for Lingens, the antidote for his condition, would require correlating his perceptual map with his propositional knowledge, so that he could figure out that the origin of his current perceptual experience must be at the present location of that guy Lingens about whom he has so much propositional information, and since there's nobody else at or near his perceptual origin, that he himself must be Lingens.

Same for another poor amnesiac, Ernie Banks, the former baseball great now fascinated with the story of the great Ernie Banks (Morgan 19??). But in this case, the pattern of ignorance Banks lacks awareness of is temporally dynamic: his present location at a coordinate is at one end of a kind of connected journey, or *path*, through space-time known as a life-story, with a shifting origin (experientially an epistemic space-time hole),

periodic awareness of being addressed (oriented to and frontally interacted with) as “Ernie Banks”, etc.

De se interpretations of attitude reports are understood to be self-ascriptive on the part of the holder of the attitude (Casteñeda 19??, Lewis 1979). When we report that:

(40) Ernie Banks knows that Ernie Banks was an out-fielder.

this is pretty clearly true. But (41) strikes us as unclear:

(41) Ernie Banks knows that he was an out-fielder.

Insofar as *he* is co-referential with *Ernie Banks* and (40) is true, there’s a sense in which (41) is true. But of course what Ernie doesn’t do is self-ascribe having been an out-fielder, and it seems to us that (41) has such an interpretation, as well. Hence, there’s a sense in which (41) is false.

In keeping with all the above, I think that self-ascription is a species of self-location, and hence is inherently perspectival. Because self-location is about ignorance, I think that we can model it in a possible world semantics with partial objects. Lewis himself was sure that the key to characterizing attitudes *de se* lies with partiality, and accordingly argued that the objects of the attitudes are properties, not propositions. But I think we can still take the objects of the attitudes to be propositions—at least when expressed by finite clauses¹⁸—and yet characterize the difference between a *de se* and a non-*de se* interpretation, by using perspectival domain restriction. Just as in the case of Cresswell’s man on the aeroplane, where the proposition expressed differs according to the perspective he adopts, so we can take the proposition expressed by the sentential complement of an attitude predicate to involve contextually-given perspectival domain restriction. The domain restricted will be that of the modal implicit in the lexical semantics of the predicate (Hintikka, et al.; see also Roberts (1998) on domain restriction and modal subordination in such predicates). The restriction will be given via an epistemic modal accessibility relation based on perspectival situations, to capture the notion of partial information that’s part of what is needed to characterize a *de se* attitude. Insofar as self-ascription is about ignorance, or lack of access to certain kinds of information at a certain locale—the origin, that ignorance can be reflected in the thickness and in the *pattern of thinness*, of the accessible situations and the individuals in them. Let me roughly sketch how this works:

Compare the two quasi-logical forms (a) and (b) for (40), (40b) containing Casteñeda’s *he** to indicate that it’s to reflect the self-ascriptive interpretation:¹⁹

- (40) Lingens knows that he is in the Stanford Library.
(a) Lingens knows that he is in the Stanford Library. (non-*de se* interpretation)

¹⁸ See Roberts (2006) for arguments that infinitival complements are quite different.

¹⁹ I’m not assuming that there is necessarily a difference in the syntactic Logical Form for the two interpretations. This is just a convenience to distinguish the interpretations of interest.

(b) Lingens knows that he* is in the Stanford Library. (*de se* interpretation)

Let us take as our grounding situation the thick situation consisting of all the true propositions pertaining to the 6th floor of the Stanford Library. Then consider the following two circumstances characterizing Lingens' knowledge:

Circumstance(a): Lingens is not aware that his perceptual origin is on floor 6

Circumstance(b): Lingens is aware of being at the origin on floor 6

In both circumstances Lingens has the same perceptual perspective, involving seeing the immediate concrete situation around him, with patterned ignorance; let's call this relatively thin set of propositions $s_{perspec}$. Also, in both circumstances Lingens has the same total access to propositional information about the ground, without ignorance; let's call this thick set of propositions s_{map} . In Circumstance (b), Lingens knows that the two perspectives are grounded in the same situation, hence have a common coord. He hence knows that the origin of his perceptual experience is somewhere on floor six on his propositional map. This means that in all the worlds w' that are Lingens-epistemically accessible from w there is a complete individual $Lingens_{w'}$ such that:

$$origin_{perspective} \subseteq Lingens_{w'} \text{ and } Lingens_{map} \subseteq Lingens_{w'}$$

Hence both partial individuals grow into the same complete individual in all Lingens' epistemically accessible worlds, so that he knows that the individuals playing a role in the different perspectives he has on the situation are one and the same. On the other hand, in Circumstance (a), Lingens does not have this information about the relationship of the two perspectives. Hence, there are worlds w' Lingens-epistemically accessible from w where there is no individual $Lingens_{w'}$ such that $origin_{perspective} \subseteq Lingens_{w'}$ and $Lingens_{map} \subseteq Lingens_{w'}$. I.e., there are epistemically accessible worlds in which the two partial individuals do not correlate with a single individual. Hence, although Lingens might admit that it was *possible* that he* was Lingens, he doesn't know so.

On the present proposal, situations are informational entities, sets of propositions, as are thick individuals—individuals about whose properties we have information. The difference between circumstances (a) and (b) in (40) is a difference in whether one knows that two partial thick individuals are actually the same. The two partial individuals corresponding to the perceptual *origin* and the man known in the books as *Lingens* may turn out to be the same in all the epistemically accessible worlds for such an agent, or they may not. I.e., the two perspectives may map via epistemic access onto the same grounding situation with Lingens at the perspectival origin, or not. The joint mapping to a single individual is required for the truth of (40).

Hence, the account is one in which in the two different circumstances (a) and (b) different propositions are expressed by the complement *that he is in the Stanford Library*, in virtue of different domain restrictions over the universal modal lexically inherent in the perspectival attitude predicate *know*. Then, though the content of the sentential complement to the attitude is the same when uttered in the two circumstances, that

content does not by itself determine a proposition. Like modals in Kratzer (1977, 19??), the implicit (quasi-)epistemic modal in an attitude predicate has its domain restricted by a contextually-given modal accessibility relation. This is a form of semantic contextualism.

This differs from the “Austinian view” promoted by Barwise & Perry (198?) and Kratzer (1998, 2005): For these authors, an attitude is a three-way relation between an agent, a proposition and a resource situation or topic situation. So it could be true that an agent believed a proposition relative to one resource situation, but false that she believes that same proposition relative to a different resource situation. This seems to be a covert form of semantic relativism, and I would argue that the notion of proposition that we get out of the present proposal is more intuitive and coherent, not to mention more persistent.

Finally, the fundamental claim here is that with the attitudes, the modal accessibility relation is perspectival, and in fact self-ascriptive (again, see Roberts 2006 APA talk). And the possibility of *de se* interpretation arises out of the distinguished pattern of lack-of-information associated with the awareness that some individual in the eventuality toward which one has an attitude is at the origin of one’s own perceptual perspective brought to bear in restricting the modal domain.

§4 Descriptions

§4.1 Incomplete

§4.2 Specific

§5 Conclusion

inter alia:

Ranger Gary: Both his general epistemic perspective and his visual perspective might come into play in interpreting a single given utterance. This bears on epistemic certainty: Are you sure? Well, yes and no. Also, the different kinds of standards of certainty we bring to bear on a given judgment, depending on the relevant goals, could be modeled as whether or not those goals play a role in the perspective we bring to bear on the judgment.

What does the inherent persistence of this framework buy us?: Cf. predicates of personal taste.

References

Acknowledge: The work of John Lyons, David Lewis, Barwise & Perry, Jonathon Mitchell, Angelika Kratzer. Discussions with Sally McConnell-Ginet, Louise McNally, Rich Thomason, Paul Portner, Thony Gilles, Kai von Fintel, and Wm Taschek; Alan and Clara Crockett and Ann Silverman.