Craige Roberts
May, 2010

Retrievability and Incomplete Descriptions*

0. Introduction

Whenever something is missing in an utterance—there is an anaphoric presupposition to be satisfied or an ellipsis to be interpreted—in order for the utterance to be meaningful$_{nn}$ (in Grice's sense) and felicitous in the context of utterance, the missing material must be Retrievable in that it must be reasonable for the speaker to expect that the addressee can grasp the speaker's intended meaning$_{nn}$ in so-uttering in that context. The theory of discourse context in Roberts (1996,2004) permits us to make more precise the sense in which content is Retrievable. These shed light on a number of *prima facie* diverse phenomena, including ellipsis (Roberts, in preparation a) and prosodic deaccentuation (Roberts, in preparation b), as well as on presupposition resolution generally. Here I focus on the ramifications for the interpretation of definite descriptions and pronouns, and in particular those whose descriptive content would not suffice to pick out a unique denotatum.

In previous work (Roberts 2003) I argued for a revision of the classical Russellian treatment of definite descriptions, proposing instead that they conventionally trigger two presuppositions, one of *weak familiarity* (a form of anaphoricity) and a second I called *informational uniqueness*. These are the informational counterparts of Russellian existence and uniqueness, respectively. In other work, I argued that these same presuppositions are central to the meaning of pronouns (Roberts 2004) and demonstratives (Roberts 2002). Now I show that the general Gricean view of discourse sketched here permits a simplification of that theory: The uniqueness *effect* observed in certain contexts follows from Retrievability, with no need to stipulate even informational uniqueness. The resulting theory stands in contrast to a number of prominent recent treatments of definite noun phrases which treat them as syntactically and/or semantically elliptical: the descriptions as incomplete (Neale, Recanati), the pronouns as E-type or D-type implicit descriptions (Evans, Neale, Heim 1990, Elbourne, *inter alia*). I argue that the so-called problem of incomplete descriptions is actually a non-problem, stemming from the attempt to treat a fundamentally presuppositional phenomenon in purely semantic, proffered terms. But still, the analysis rests on the conventional content of the definite article and pronouns—their status as triggers of anaphoric presuppositions. So unlike Reimer (1998) and Lepore (2005), it does not rest on a distinction between the conventional content of an utterance (false, when a description therein is incomplete) and the proposition conveyed (still possibly true in such a case).

The paper is organized as follows: In section 1, I lay out the problem and some data. In section 2, I introduce the notion of Retrievability and relate it to a theory of context which helps to gives it more content. In 3, I apply that theory to incomplete definite noun phrases, arguing that it

* I am grateful to many people who have given me criticism, comments and suggestions on earlier versions of this paper. In particular, thanks to Barbara Abbott, Michael Glanzburg, Hans Kamp, Chris Kennedy, Ezra Keshet, Yusuke Kubota, David Manley, Sarah Moss, Richard Samuels, Kevin Scharp, Stewart Shapiro, E. Allyn Smith, Jason Stanley, and Richmond Thomason; to participants in the workshop on Ellipsis and Philosophy of Language at the University of Western Ontario in spring 2006, the 2009 University of Michigan Workshop on Anaphora, the OSU Working Group in Pragmatics, my autumn 2009 seminar in the Dept. of Philosophy at the University of Michigan, and Stewart Shapiro's spring 2010 Seminar on Contextualism and Relativism in the Dept. of Philosophy at OSU; and to Colloquium audiences at The OSU Dept. of Philosophy, Michigan State University, and the University of Chicago.

provides a satisfactory way of accounting for their occurrence and interpretations in context. In 4, I compare this approach favorably to that of two other contemporary theories of incomplete definites, Neale (1990) and Elbourne (2005). And in 5 I draw conclusions and consider prospects for further work

1. Anaphora and Anaphora Resolution: Incomplete content in context

Anaphoric elements in human language are those in which something is missing which the speaker presumes can be found. The Greek root of the word means 'to bring back': Accordingly, an anaphoric element carries a conventional presupposition to the effect that in a discourse context in which it is felicitously uttered its intended interpretation can be readily retrieved from information antecedently available.

In general, anaphora does not require a *linguistic* antecedent—in the case of definite NPs, a preceding coreferential NP. A definite "brings back" some entity which was already under consideration. Though prior mention of such an entity with an explicit NP suffices for it to have been under consideration, that is certainly not necessary, as is illustrated by examples like (1) – (6) below. The anaphoric elements of interest are underlined.

(1) [Tim and Margaret are sitting at a conference table in her psycholinguistics lab, working on a grant proposal for some eye-tracking experiments. Tim is making notes on his laptop computer, not a MacIntosh brand.]
Margaret: How do you like your laptop?
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, in the direction of her desk in the middle of the room]
I just got [<u>that</u>]$_F$ 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 Mac computer monitor on it, past the desk an eye-tracker, past that a wall calendar.

Margaret's gesture is vague, in that:
(a) even from her perspective it doesn't uniquely pick out the monitor, and
(b) Tim has a slightly different perspective than hers anyway, because he is sitting two feet to the south, so that he couldn't say exactly which thing is at the end of the trajectory from Margaret's shoulder through her index finger to the desk-area.
Still, Tim has no difficulty picking out the monitor as the entity demonstrated—and hence the computer of which it is part as the intended referent—because the question Margaret asked at the outset made the computer the only entity in the direction indicated that was relevant to their discussion.

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

(3) [naturally occurring example (1992): Ed Keenan is giving a colloquium at the University of Amsterdam. During the talk, he twirls his reading glasses by the stem and the screw holding the stem to the frame falls out, so that the glasses drop to the table. He puts the

stem down and continues his talk. At the break, he begins searching intently under the papers on the table.]
Craige: It's probably on the floor.

Though there is no mention of the missing screw in (3), use of the pronoun is perfectly felicitous. Something comparable is going on in (4), involving the anaphoric presupposition triggered by *too*, in the elliptical constructions in (5) and (6):

(4) [Two women are standing at a bus stop on a rainy day. A car drives by, through a puddle, splashing one of the women with muddy water.]
To the splashed woman: One splashed ME this morning, too.[1]

(5) [A policeman on the beat turns a corner and sees a young kid with a rock, poised to throw it through the window of a school. The kid doesn't see him yet:]
Policeman: I wouldn't, if I were you.

(6) [A guy is standing in front of a make-shift monument on the sidewalk in front of a house where two neighborhood kids carried out a suicide pact earlier that week. Another neighbor comes up to stand beside him, also silently gazing at the monument. After awhile, one of them says:] I just cannot fathom why.

Despite the fact that most authors in the literature on *too*, VP ellipsis and Sluicing assume that generally an explicit antecedent constituent is required in all these cases, (4), (5) and (6) show that on occasion that is not the case.

I think we can tell a very simple story about how we retrieve the intended content of incomplete or elliptical expressions, however they are licensed:

(a) A speaker can only be reasonably assured that the use of such descriptively impoverished expressions will be successful if she can assume that at the time of utterance her interlocutors will be attending to the intended referent, or to something closely related to it in some way.
(b) At any given time, what we intend is the central factor in determining what we attend to. Our goals and commitments thus direct our attention.
(c) The exchange of information in a discourse is organized around questions for discussion. The goal of conversation at any given point is to address the question currently under discussion, and there are rational constraints on which questions can be felicitously raised, given what previous questions remain unresolved. The upshot of this is that in making an utterance the speaker can reasonably expect that a competent, cooperative addressee who is engaged in the discussion will be attending to the resolution of the question under discussion.
(d) Hence, in a rational interchange the retrieval of what has been omitted is necessarily dependent either on the question under discussion or, if the addressee's attention is elsewhere, on whatever her evident immediate goals lead her to attend to in that extra-linguistic sphere.

This story intuitively respects and accounts for why explicit antecedents are required in cases where the descriptive content of an anaphoric element is relatively impoverished (as with a

[1] The capitalization of *me* indicates that it bears the nuclear accent, or greatest prosodic prominence in the utterance.

pronoun, *too*, or syntactic ellipsis), but also leaves room in certain contexts for felicity despite a lack of explicit antecedent constituent. It is quite clear how this pans out in an example like (1). Margaret's question establishes the parameters of relevance for the conversation. In answering the question, Tim raises one consideration bearing on his level of satisfaction with his computer, that of its age. This criterion makes implicitly salient a more general issue about age and obsolescence in one's professional equipment, and in turn this makes it relevant for Margaret to contrast her own computer with Tim's in this respect. Seen through the filter of relevance to this line of inquiry, the indicated still-life—desk, eye-tracker, calendar, etc.—contains only one object directly relevant to the question that frames the discussion. Assuming that Tim was attending to the line of inquiry, hence intended to address the questions raised by Margaret, Margaret can assume that in looking at the desk Tim will be looking for the most relevant potential demonstratum. Hence, there is no need for more descriptive content in the demonstrative NP. Margaret could have said *I just got that computer last year*, but she didn't need to. And one of the points of elliptical expression is to be economical where feasible.

Lest there be any question about whether it is the currency of the question that counts—so that it has to be the *one question immediately under discussion* (QUD), consider the variant in (1′). Recall that the grant that Margaret and Tim are working on is for a series of eye-tracking experiments. Assume that both know that the eye-tracker located on the other side of her desk is nearly obsolete, compared with the new models, and this has been hindering their research:

(1′)	Margaret:	How do you like your laptop?
	Tim:	It's not bad, but it's getting kind of old. Anyway, let's get back to business. What do you think of this budget?
	Margaret:	[gesturing with over her shoulder, in the same direction as in (1)] Could we add some money to replace <u>that</u>?

Note that in (1′) the demonstrative *that* is as close in actual word-count as was its counterpart in (1) to both the question Margaret originally asked about Tim's computer, and to Tim's last mention of a computer. Yet here, because the question under discussion has switched, one naturally takes the demonstrative to refer not to the computer, but to the as-yet-unmentioned eye-tracker. It is the latter that is relevant for the proposed budget, and hence relevant to the question Tim has raised.[2]

In (2), the definite description *the fire station* does not uniquely denote. Since the speaker is giving directions to a stranger, we cannot even assume that there is a fire station which is somehow unique in their mutual experience (say, 'the neighborhood fire station', or the like). Yet not only does the example seem felicitous, but the definite description displays the uniqueness effect typically associated with the Russellian interpretation of definites: The addressee can reasonably assume that when she gets four blocks from her turn, not only will there be a fire station, there will be no more than one of them. Yet the noun phrase *the fire station* is incomplete in that it does not have sufficient descriptive content to uniquely pick out any given fire station, among all the thousands in the world. I have argued (Roberts 2003) that this definite is felicitous in (2) because directions are designed for use in a particular context, the one the addressee will find himself in when he is intent on following the directions. Given this intention, it is reasonable to expect that he will be closely attending to the situation along the directed path. As he does, when he gets to the corner four blocks after the light, he will expect to see one and only one fire

[2] In general, when one interlocutor has posed a question for discussion, as Tim just has in (1′), the respondent will be taken to be addressing that question unless s/he explicitly rejects it. So we take Margaret to have accepted his proposal for the QUD at the time of her turn.

station, marking where he should turn. If this is the case, then the directions will have proven accurate; otherwise, not. Abstracting away from particularities, this is just the kind of circumstance that justified the vague demonstration accompanying (1)—assumptions about what the addressee can be expected to attend to at the time of interpretation justify the use of a descriptively minimal definite, whether demonstrative or definite description.

Now it should be clear that what licensed the use of *it* in (3) was very much the same kind of assumption on my part about what Ed Keenan was attending to at the moment—I reasonably assumed that he was searching for the missing screw for his glasses. Since some kind of representation of that familiar object was what he was clearly attending to, and he could reasonably assume that that would be obvious to an audience member who had observed his mishap, my utterance was felicitous. Note that in form, (3) is perfectly parallel with the well-known marble examples in (7) and (8), due to Barbara Partee (p.c. to Heim 1982):

(7) I dropped ten marbles and found all of them, except for one.
It is probably under the sofa.

(8) I dropped ten marbles and found only nine of them.
?It is probably under the sofa.

The first sentences of these two examples are logically equivalent, both entailing that there's exactly one missing marble. Yet, at least when the examples are uttered more or less out of the blue, pronominal anaphora to the missing marble seems less felicitous in (8), where the missing marble wasn't directly mentioned, than in (7), where it was. Roberts (2003) argued that the difference lies in the lower salience of the missing marble when it hasn't been mentioned; pronouns require very high salience of their antecedents, due to their extremely impoverished descriptive content. In this, they contrast with definite descriptions, as we see reflected in the fact that (9), with a definite description instead of a pronoun, is impeccable when uttered out of the blue:

(9) I dropped ten marbles and found only nine of them. [Roberts 2004]
The missing marble is probably under the sofa..[3]

Here, again, the definite description is descriptively incomplete, in that there are surely many missing marbles in this world. But the situation under discussion to which the addressee is attending contains only one missing marble so far as we know, so there is only one *relevant* missing marble, and the example is, in this respect, like (3). Again, (3′) illustrates that what is important is not what precedes the moment of utterance so much as what we take the addressee to be attending to:

(3′) [The same colloquium scenario, up to the break. But at that point, Keenan absent-mindedly picks up the stem of the glasses and plays with it while he chats with an audience member about his talk.]
Craige: #<u>It</u>'s probably on the floor./√<u>The screw</u> is probably on the floor.

Keenan isn't attending to the missing screw in (3′), so the salience condition on

[3] Default nuclear accent on the subject of the second sentence will be on *missing*, so that *marble* bears no accent. This is consistent with the description of what licenses use of the definite description in (9) in conjunction with the theory of de-accentuation of Schwarzschild (1999) (under the assumptions here about weak familiarity, as discussed below). See Roberts (in preparation b) for discussion of how Retrievability bears on deaccentuation.

pronominal anaphora is not satisfied. But in the situation of utterance there is only one obviously missing screw, so a hearer could make sense of the incomplete definite *the screw*.

The remaining examples are of the same character; in each, the incompleteness is justified by the perceived attention of the addressee to the intended missing argument, property, or proposition. (4) contains both the nominal pro-form *one* and *too*, the latter triggering a presupposition whose content, 'a car splashed you this morning', is hinted at by the focal structure of the clause to which it attaches. Two types of syntactic ellipsis are illustrated by (5) and (6), VP ellipsis and Sluicing. There is a long-standing debate over the character of such constructions and how to explain their distribution and restrictions, but everyone agrees that there is something missing in their syntactic structures, as in a range of other elliptical phenomena, like gapping and null arguments (*pro*-drop). In (5), VP ellipsis is licensed by the mutual recognition on the part of the policeman and the boy of the latter's evident intentions. And in (6), Sluicing (deletion of the clausal complement of a *wh*-complementizer like *why*) is licensed out of the blue by recognition on the speaker's part of what the addressee was evidently attending to, the sad event of earlier in the week.

In the following section, I will develop the story about how incomplete content is retrieved in more detail, in the context of a more general theory of the structure of discourse context which I've argued for extensively elsewhere, and against the backdrop of a principle of *Retrievability*, which I'll argue follows without additional stipulation from the Gricean view of meaning in discourse. I want to claim that in a variety of anaphoric domains, including those illustrated in (1)-(6) and others besides, there are *prima facie* unrelated puzzles which can be addressed satisfactorily by a clear theory of Retrievability in conjunction with a theory of the intentional structure of discourse.

2. Retrievability

I assume that something close to Grice's notion of utterance meaning should form the foundation for any semantic account: [4]

> "U meant_{nn} something by uttering x" is true iff, for some audience A, U uttered x intending:
> (1) A to produce a particular response *r*
> (2) A to think (recognize) that U intends (1)
> (3) A to fulfill (1) on the basis of his fulfillment of (2). (Grice 1957)

Note that meanings_{nn} are conveyed by speakers engaged in a certain kind of action, that of making utterances:

> **utterance**: The use of a linguistic constituent *s* in a given context *c*. Formally modeled as <s,c>

[4] A later definition in Grice (1969:99-100) is intended to handle various problems via a prohibition of covert intentions, adding the following to the definition cited in the text:

> [Moreover,] there is no inference-element E such that U uttered x intending both (1') that A's determination of *r* should rely on E and (2') that A should think U to intend that (1') be false.

I take it that whatever a context is, it should at least contain information about the identity of the agent of the utterance, and whatever other else one might deem important to individuating utterances.

It's only utterances that have speakers' meanings$_{nn}$. Linguistic constituents themselves, including sentences, have only conventional meanings (sometimes called *conventional content*), and in the general case, these conventional meanings do not by themselves determine propositions. This last point is clearest, perhaps, with utterances containing anaphoric and indexical constituents, whose intended meaning is thus essentially a function of the context of utterance. But as Grice makes clear, in some cases what is meant is more than what is explicitly said, as in conversational implicature, or else there is something obviously missing in the utterance—a conventionally triggered anaphoric presupposition to be satisfied or an ellipsis to be interpreted.

On this Gricean view, to which I would generally subscribe, in order to grasp what a speaker *means$_{nn}$* in making an utterance, the addressee(s) must be able to recognize the speaker's semantic intentions therein, presumably in part on the basis of the speech signal or written code utilized in making the utterance. Of course, this requires that the addressee speak the language in which the utterance is made, that the performance of the utterance (articulation or penmanship) be competent, that acoustic conditions (or lighting, e.g. for sign language) be adequate, etc. But even when these obvious conditions are met, this is not sufficient to guarantee that the intended meaning can be understood.[5]

Note that *pace* von Rooij,[6] even rivals must cooperate in the following sense if they're to mean something by what they say: They must behave in such a way as to permit the competent, attentive addressee to recognize their semantic intentions. This being the case, it follows that the retrievability of speakers' meanings (under reasonable assumptions) is one of the hallmarks of rationally cooperative discourse. Formulating this as a pragmatic principle:

Retrievability: In order for an utterance to be rationally cooperative in a discourse interaction *D*, it must be reasonable for the speaker to expect that the addressee can grasp the speaker's intended meaning in so-uttering in *D*.

I take it that Retrievability is what Neale (2005) has in mind in the following:

> If, as Grice suggests, what *A* meant by uttering *X* on a given occasion is determined by certain interpreter-directed intentions, then assuming he is being co-operative *A* cannot mean that *p* by uttering some sentence *X* if he believes it is impossible for his audience *B* (or at least any rational, reasonably well-informed interpreter in *B*'s shoes) to construe him as meaning that *p*. (77)

> [Moreover,] *A* cannot (intend to) refer to some particular individual α by *X* if he believes it is impossible for his audience *B* (or at least some rational, reasonably well-informed interpreter in *B*'s shoes) to construe him as (intending to) refer to α by *X*. (81)

Assume that the intended response in clause (i) of Grice's definition of meaning$_{nn}$, say for an assertion, is for the addressee to recognize that the speaker proposes that the proposition expressed is true (and thereby proposes that it should be added to the Common Ground, following

[5] Since we are solely concerned with linguistic meaning, in what follows I use *mean(t)* and *meaning* as shorthand for *mean(t)$_{nn}$* and *meaning$_{nn}$*.
[6] Most recently in Franke, de Jager & van Rooij (to appear).

Stalnaker 1979).[7] In order for the addressee to produce this response, he must grasp the proposition that the speaker intends to express. I.e., it must be Retrievable. Retrievability is, then, a constraint on rational cooperation in discourse that follows from the canonical goal of communication, to convey meaning$_{nn}$. Assuming that a speaker is rational, if she means$_{nn}$ something by what she says, it follows that she assumes that what she means$_{nn}$ is (so far as she knows) Retrievable.

But what would make it reasonable for a speaker to take what she means, her unique semantic intentions on that occasion, to be Retrievable? We have to account for how clause (2) of Grice's definition can be satisfied in cases involving anaphora, syntactic ellipsis, and other non-explicit facets of the intended meaning:

> What is required in order for it to be reasonable to intend that one's audience will recognize that one intends them both to grasp the proposition one intends to express and to recognize that one proposes that this proposition is true?

I would maintain that **the structure of a discourse interaction is designed to help satisfy this constraint**, facilitating the conveyance of both intentions—the Retrievability of the intended proposition and the recognition that the speaker proposes that it is true. This is not by convention, but by virtue of the structure of discourse being functionally optimal from the perspective of the conveyance of meaning.

Discourse is an activity, a task-oriented interaction between interlocutors. Elsewhere, I have characterized the general task in discourse as the conduct of an inquiry, the attempt to share information in the interest of resolving certain questions. We engage in these inquiries because we have other goals—perhaps we want to learn how to perform heart surgery or to drive from Ann Arbor to Chicago. Hence, the goals involved in inquiry are typically subordinated to other goals, the whole structured in a way that respects the relationships between the associated intentions. We could think of the intentional structure of a discourse interaction roughly in the following terms:

[7] Grice (1957) assumed a stronger response: creating a belief in the addressee that the indicative is true. But in Grice (1968:230) he recognized that this was too strong, and modified the intended response to be that "the hearer should *think that the utterer believes* something". I think my formulation here is a reasonable modification of this view, on the assumption that the speaker *is* trying to *convey information* to the addressee via the assumed truth of the proposition expressed, albeit perhaps indirectly, and that insofar as the speaker is cooperative and observes the maxim of Quality, this will implicate that the speaker herself takes the utterance to be true.

Scoreboard of a rational discourse interaction *D*: [Roberts 1996,2004][8]

At any given point *t* in *D*, the information shared by the interlocutors is structured as follows:

I, the set of interlocutors at *t*

G, a set of sets of goals in effect at *t*, such that

for all $i \in$ I, there is a (possibly empty) $G_i \in G$ which is the set of goals which *i* is committed at *t* to trying to achieve, and

$G = \{ G_i \mid i \in I \}$.

$G_{com} = \{g \mid \forall i \in I: g \in G_i\}$, the set of the interlocutors' common goals at *t*.

M, the set of moves made by interlocutors up to *t*, with distinguished sub-sets:

$A \subseteq M$, the set of assertions

$Q \subseteq M$, the set of questions

$S \subseteq M$, the set of suggestions

$Acc \subseteq M$, the set of accepted moves

$<$ is a total order on M, the order of utterance

CG, the common ground, the set of propositions treated as if true by all $i \in I$ at *t*

DR, the set of discourse referents, corresponding to entities entailed to exist in CG

$QUD \subseteq Q \cap Acc$, the ordered set of questions under discussion at *t*, s.t.

for all $Q \in QUD$ there is a $g \in G_{com}$ such that g is the goal of answering Q, and

for all $Q \in QUD$, it is not the case that CG entails an answer to Q

For all $i \in$ I, if *i* is a sincere, competent and cooperative interlocutor in *D*, we can characterize two kinds of goals held by *i* at *t*:

Discourse Goals = QUD

Domain Goals of *i* = $G_i \backslash Q$

The information on the scoreboard at a given time *t* constitutes the **context** of utterance for any given utterance made at *t*. In an idealized model of how utterances have meaning, the context is explicit, as if fully accessible to all interlocutors. It is useful to work with such idealized models in order to develop a theory of what the relevant parameters of discourse are, and how they interact. This permits us to make clear predictions about *what s would mean if uttered in D*, despite the fact that we know that in fact contexts are generally not fully accessible to all the participants, and may often be defective in some sense, in that some of the interlocutors may have incorrect information about what is assumed by the others. We'll have more to say about DR, the set of discourse referents, in the following section. Here I will focus on the Common Ground and the set of questions under discussion.

It is important to note that as in Stalnaker's conception of the Common Ground, information in CG needn't be introduced linguistically. Information that is available through common experience, culture or perception in the situation in which the discourse takes place has the same status in CG as the propositions expressed by utterances in the discourse in question. Moreover,

[8] The scoreboard metaphor is, of course, due to Lewis (1979), the notion of inquiry to Stalnaker (1979, etc.), the idea that discourse games are structured around questions to Carlson (1983), and the role of intentions to Planning Theory in Artificial Intelligence (see especially early applications to discourse in Grosz & Sidner (1986), Perrault (1990), Thomason (1990), and other references cited therein). This is the notion I called the *Information Structure* of the discourse in the earlier work cited. I have reverted to the scoreboard metaphor because so many others use the term *information structure* in a very different sense—a level of syntactic structure, or the like. I meant something very different, an actual structure on information.

CG includes information about the discourse itself, including the previous and current states of the scoreboard.

Sarah Moss (p.c.) notes that insofar as all the information about the moves made and whether they were accepted might be expected to be reflected propositionally in the CG, one might argue that a distinct QUD is otiose. However, I would argue that although from the CG as described we can retrieve the QUD at any given point in the discourse, these two elements of the scoreboard are functionally distinct in that (a) they play different roles in constraining felicity and driving scoreboard update, and (b) they themselves display different patterns of change over the course of update. With respect to the (b), CG is (ideally) monotonic, while QUD is not—as the Pragmatics of Questions suggests, ideally questions are ultimately taken off the QUD stack when answered (or given up as practically unanswerable). But point (a) is, I think, even more important, and bears on the fundamental insight behind this approach to pragmatic theory.

Taking certain kinds of communicative intentions to be the core of Gricean meaning$_{nn}$, it is perhaps not surprising that the interlocutors' joint communicative goals, in the form of the QUD, constitute a central feature of the quasi-public record of a conversation. Moreover, I argue that the intentions which serve those goals are both the motivating force for further conversation and the central constraint on cooperative behavior in such an interchange and, hence, on felicity in discourse. Let's assume that the goals a rational agent holds at a given time are ideally consistent, in view of the fact that they're being simultaneously pursued. Then the discourse goals to which the agent subscribes, inquiring into particular questions, must in turn be compatible with the agents' domain goals. In fact, typically discourse goals subserve domain goals. This plays an important role, for example, in the way that agents' assumed domain goals can play a role in the generation of conversational implicatures. In other words, the whole scoreboard reflects a game which is fundamentally structured and constrained by rationally pursued intentions.

This characterization of context is intended to capture a conception of discourse as a game in which the central goal is to share information, as in Stalnaker's *Inquiry*. But what's missing in Stalnaker, I would argue, is recognition of the role of goals and intentions in the pursuit of shared information. An inquiry is pursued not in a haphazard way, with interlocutors flinging out propositions for addition to the CG, but via an orderly investigation of chosen issues, or questions. The interlocutors adopt particular discourse goals, i.e. goals to address particular questions. Cooperative interlocutors adopt the intention to pursue the particular goals mutually adopted.

They do so via making three central kinds of moves, or utterances, in discourse—assertions, questions and suggestions. Each of these kinds of moves affects the score, i.e. leads to the update of the scoreboard for the utterance, doing so in a way particular to that kind of move:[9]

Pragmatics of Assertion: (following Stalnaker 1979)
If an assertion of $.\alpha$ is accepted by the interlocutors in a discourse D, $|\alpha|^D$ is added to CG on the scoreboard at that point in D.

[9] In these definitions, for any linguistic constituent κ, $|\kappa|^D$ is the compositional semantic interpretation of κ relative to *D*. $?\kappa$ indicates that κ is in interrogative mood, $!\kappa$ indicates that κ is in imperative mood, and $.\kappa$ indicates that κ is in indicative mood.

Pragmatics of Questions: (Roberts 1996)
If a question posed by **?**α is accepted by the interlocutors in a discourse D, then $|?\alpha|^D$, a set of propositions, is added to the QUD in D. A question is removed from QUD iff either its answer is entailed by the common ground CG or it is determined to be unanswerable.

Pragmatics of Suggestions: (Roberts 2004; cf. Portner 2005, 2007)
If a suggestion posed by !P, P a one-place predicate, is accepted by the intended addressee *i* in a discourse *D*, $|P|^D$ is added to G_i, the set of *i*'s goals in *D*, and |intend(i,[P(i)])| is added to CG.

All of these canonical (types of) actions in discourse are guided and constrained by the goals reflected in the QUD. To see what I mean by this, consider how goals and the intentions to achieve them are related to action. What is it to intend (to do) something? Bratman (1987:17) argues convincingly that holding an intention involves certain characteristic dispositions on the part of a rational agent: *Inter alia*, an intention to act in a certain way involves "a disposition to retain this intention without reconsideration, and a disposition to reason from this retained intention to yet further intentions, and to constrain other intentions in light of this intention." I.e., intentions "are conduct-controlling pro-attitudes, they have inertia, and they serve as inputs into further practical reasoning" (27). An intention always involves a goal, and having a goal involves having the intention to achieve it. Given a goal and the intention to achieve it, we reason to form a (partial) plan to do so. Hence, a rational agent's intentions are intrinsically bound up with her plans for action.

Given this, we can see that the following principle[10] is a consequence of what it is to behave rationally in the kind of intentional interaction modeled in the conception of a discourse game I have proposed, involving the scoreboard sketched above and the three canonical kinds of moves:

Rational Cooperation in a Discourse *D*: Make your utterance one which promotes your current intentions in *D*. [cf. Grice 1967, Bratman (1987), Thomason 1990]

We can be even more specific about what this entails: Given that among the current intentions of all the interlocutors in *D* is the intention to resolve the set of questions under discussion, and given what it means to hold an intention, then any rational, cooperative move should address those questions:

A move *m* is **Relevant** at a given point in a discourse *D* iff it addresses the question under discussion at that point in *D*, where:

> An utterance *u* **addresses** a question *q*, iff either *u* is an assertion which contextually entails a partial answer to *q*, or *u* poses a question which is part of a strategy to answer *q* (*u* is a sub-question of *q*), or *u* is an imperative whose acceptance and enactment might help to answer *q* (*What's in the closet? Open the door!*). [Roberts 1996/2004]

.
If we take the goal of a cooperative interlocutor to be to address the question under discussion at that point in discourse, then though I agree that Relevance is a central principle of discourse, we needn't stipulate that it is required for felicity. It simply follows that a cooperative interlocutor will endeavor to make her utterances relevant in this sense. Among other things, this effectively puts a logical constraint on the QUD stack, so that adding a question is only permitted if it's a

[10] This is clearly a version of Grice's Cooperative Principle: **"**Make your conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged." (Grice 1967)

sub-question of the question previously on top. I.e., the new question must be one whose complete answer contextually entails a partial answer to that question already on the stack; cf. the sub-question relation defined in Groenendijk & Stokhof 1984, and the discussion and illustrations in Roberts 1996. Thus, all the questions on the QUD stack are required to contextually entail all those beneath them on the stack; and if you offer a partial answer to the topmost question, you have thereby offered a partial answer to all the others.[11]

More generally, we can define what it is for an action to be Relevant to a goal:[12] This is so in a given context just in case the action in question furthers the achievement of the goal in that context. Insofar as discourse goals are subordinated to domain goals, then the QUD must further those domain goals, and any way of furthering the pursuit of a question under discussion is potentially a way of furthering the domain goal to which it is Relevant. This is intended to suggest, following Grice, that the general character of meaning, and of the discourse structure which subsumes it, is not inherently linguistic. Both are functionally grounded in what it is to hold an intention, and in what is required to make such an intention efficiently evident to another agent in the course of a collaborative interaction.[13]

Returning to the question of the relationship between CG and QUD, I hope it is clear that the two play very different roles in constraining interlocutors' actions and supporting their intentions in a discourse. The CG is just a repository of shared information, useful but in itself inert. It has sometimes been suggested that the content of a given CG does play a role in constraining discourse felicity. For instance, Stalnaker[14] proposes a principle according to which one cannot assert a proposition which is already entailed by the CG. But this is questionable. For example, we often assert propositions which are already known, not with a view to asserting their mere truth, but because *insofar as they bear (directly or indirectly) on the particular question under discussion*; bringing them back to the attention of the interlocutors *in this connection* may help to resolve the QUD. We might say that one asserts them *as Relevant*.

In contrast, the QUD both drives and essentially constrains felicitous play in the discourse game, by virtue of the status of its constitutive questions as common goals of the interlocutors and of the way that goals influence rational action, as encoded in Relevance. Kevin Scharp (p.c.) observes that use of an independent QUD on the scoreboard results in a finer-grained theory, one which permits a more perspicuous characterization of discourse moves, coherence and felicity. For example, though as Moss suggests we can derive the QUD from CG, CG alone is too blunt an instrument to permit ready formulation of the pragmatics of questions. Insofar as Relevance to the QUD *does* play a central role in driving and constraining discourse, a scoreboard with QUD

[11] In conversation, Ben Caplan and David Sanson pointed out that licensing by existence entailments is problematic: Which entailments can we reasonably expect an interlocutor to draw? But, of course, this is a general problem, not specific to the present proposal. Inference is arguably the central engine driving a Gricean theory of interpretation. Clearly, it must be constrained: We need recourse to a notion of *practical reasoning*, the subject of a great deal of study in artificial intelligence. See Bratman (1987) and a great deal of subsequent work on plans and practical reasoning.

[12] This was suggested by me by Kevin Scharp (p.c.).

[13] My concern with Relevance and with the efficient conveyance of more than what-is-said are shared by Sperber & Wilson (1985) and others who work on Relevance Theory in that tradition. However, in many central respects the theory proposed here differs significantly from theirs. A comparison is beyond the scope of the present paper. But do note that what it is to rationally hold an intention plays no role in their definition of Relevance.

[14] E.g. Stalnaker (1979:325): "A proposition asserted is always true in some but not all of the possible worlds in the context set."

distinguishes that aspect of the information shared by the interlocutors which plays a central and distinct function in driving and constraining discourse.

It turns out that when combined with the principle of Rational Cooperation, Relevance imposes a very strong constraint on felicity in discourse (as explored in great detail in Sperber & Wilson 1985 and related work in that tradition). For one thing, Relevance arguably represents a core requirement on discourse coherence—succeeding utterances tend to be related via their Relevance to the QUD. This is not to suggest that this is all there is to say about coherence. As I discussed in Roberts (2004), there are many ways to address a given question, which I call various **strategies of inquiry**, a given strategy often involving a sequence of utterances. Theories of rhetorical relations and constraints on coherence proposed by Hobbs et al. (1993), Asher & Lascarides (1994,1998,etc.), and Kehler (2002), among others, can be taken to capture features of the micro-structure of a well-formed strategy of inquiry.

But just as important, Relevance in the sense defined here—constrained by the QUD and the logical constraints on relations between questions in the QUD—offers a very strong and useful constraint on the inferences which must be drawn to assure that we have Retrieved the speaker's intended meaning. In particular, I want to argue that Relevance to the QUD radically restricts the search space for the antecedents for anaphoric expressions, facilitating their Retrievability. This perspective sheds light on examples like (1) – (6), considered in section 1, and has implications for the general theory of anaphora, for domain restriction, and for other issues besides. To explore this in full detail is not possible here. Instead, we'll focus on one particular case: the Retrieval of the intended meaning of incomplete definite NPs. I turn to this in the following section.

3. Case Study in Relevance and Retrievability: Incomplete Definite NPs

For the better part of the 20th century, philosophers and linguists have debated the meaning of the English definite article and definite descriptions. One puzzle has been the problem of so-called *incomplete descriptions*—those whose descriptive content is by itself insufficient to uniquely identify the NP's intended denotation. This is illustrated by examples like (2), repeated from above:

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

Though the speaker clearly has a particular fire station in mind, that entity cannot be determined on the basis of *the fire station* alone. Neale characterizes the general problem:

> Our question…is, 'How are we to explain the incontrovertible fact that a speaker can use a description "the φ" in an utterance of the simple form "the φ is ψ" and thereby perform a perfectly felicitous speech act, indeed *say something true*, even though he and his hearer both know that φ(x) is true of more than one thing?' Neale (2004:107)

This is certainly a problem if one assumes that part of the conventional content of a definite description itself is that there is a unique entity in the world (or situation of evaluation) which bears the property corresponding to the descriptive content of the definite. Call the uniqueness involved on these accounts *semantic uniqueness*. Such theories may either be like the theory of Russell (1905) in claiming that uniqueness-under-the-descriptive-content is part of the proffered

content of the definite,[15] or like that of Frege (1891) or Strawson (1950) in taking such uniqueness to be presupposed. Recently, a spate of work has sought to extend the semantics of definite descriptions to the analysis of the meaning of certain unbound, non-demonstrative pronouns in discourse. Though the proposals vary in how this is to be done, as we'll discuss in the following section, in all these theories, the overt descriptive content of the pronominal which is a disguised definite description is incomplete in the sense given above. But even in those theories in which uniqueness is not semantic, one owes an account of how the meaning clearly conveyed in (2) pertains to a particular fire station.

In a series of papers (Roberts 2002,2003,2005), I have proposed a familiarity theory of definite descriptions, pronouns, and demonstratives in discourse, a theory which does not posit Russellian uniqueness as part of the semantic content of these definites. In this theory, following the earlier familiarity-based theory of Heim (1982),[16] definite descriptions, pronouns and demonstratives carry a presupposition of weak familiarity. That is, in a sense which I will define below, all definites are anaphoric. In that earlier work I adopted a weaker notion of uniqueness than the Russellian, semantic variety, which I argued does a better job than Russellian uniqueness of predicting when uniqueness effects arise in the use of definite NPs. But here I want to argue that if we understand Retrievability in the sense outlined in the previous section, we do not need to posit even this weaker uniqueness. It turns out that the following two questions are intimately related:

a. How can we explain the uniqueness effects for definite descriptions, and thereby deal with the problem of incomplete descriptions?
b. How do addressees retrieve the intended antecedent for an anaphoric element?

When we understand the interpretive effects of Retrievability in conjunction with an anaphoric theory of definites, there is no need to stipulate uniqueness for any of these kinds of NPs. The general requirement of Retrievability of anaphoric antecedents will suffice to account for uniqueness effects, when those arise.

3.1. A Weak Familiarity Theory of Definite NPs

Here I give a brief summary of the theory of definiteness of Roberts (2002,2003,2005):

English Definite NPs: definite descriptions, personal pronouns, demonstrative descriptions, demonstrative pronouns, and proper names

This constitutes a class by virtue of the fact that all these kinds of NPs display definiteness:

Definiteness: An NP is definite just in case it carries a particular kind of anaphoric presupposition, a presupposition of weak familiarity.

Anaphora is often conceived of rather narrowly, as involving a pronominal element whose interpretation is to be grasped by finding a coreferential antecedent NP in prior discourse. I propose that we derive more illuminating, as well as more empirically adequate theories of anaphora, if we aim at a more general target:

[15] Roberts (1996) suggested the term *proffered content* as a cover term for a constituent's contribution to asserted content in an indicative, the question in an interrogative, or the suggestion in an imperative.
[16] Heim herself has abandoned the familiarity approach, in Heim (1990), Heim & Kratzer (1998).

Anaphora: A type of presupposition wherein a constituent in an utterance "harkens back" to information antecedently available to the interlocutors—in terms of the theory of context proposed above, information available on the Scoreboard.

In most types of anaphora, the Retrieval of this antecedent information is crucial to grasping the speaker's intended meaning for the utterance. This is certainly the case with the anaphora triggered by use of a pronoun, where the descriptive content of the pronoun itself is rarely sufficient to determine the intended antecedent. In canonical pronominal anaphora with an overt, coreferential NP antecedent, we can think of the antecedent of a pronoun as the individual whose prior mention licenses the use of the pronoun and gives its intended interpretation. VP ellipsis is also anaphoric, and there the antecedent is a salient property; in the case of (5), this is the property of throwing a rock through the window of the school. It is often claimed that the antecedent property for VP ellipsis must be one denoted by an explicit prior VP; and in fact cases like (5), without an overt VP antecedent, are rather rare. Lack of overt antecedent is much more common with pronominal anaphora. Although some theorists use the term *anaphora* to cover only those cases with overt antecedents, I use it more broadly, requiring for felicitous use of a definite only *weak familiarity*:

Weak familiarity: A presupposition of weak familiarity requires that the existence of the intended entity be entailed by the local context of utterance, as reflected on the scoreboard in DR.

Assuming that indefinite NPs entail existence, then if a definite is coreferential with a prior indefinite NP, the NP *licenses* the use of the definite, since the indefinite's use introduces the existence entailment that satisfies the definite's familiarity presupposition. But weak familiarity does not require a preceding coreferential NP. In (3) above, the interlocutors' experience of watching Keenan's glasses fall off the stem licensed the default assumption that there was a screw which was holding them together and had come loose, and it is this existence implication that satisfies the weak familiarity presupposition triggered by use of *it*.

As has been explored in great detail in the literature on familiarity-based approaches to anaphora since Kamp (1981) and Heim (1982), the context in which the presupposition is to be satisfied must be local, not necessarily the global context of the CG, since the local context may include information which is hypothetical or provisional in some sense, thus not necessarily assumed to be true. This is illustrated by the famous donkey sentences, like (10), and also by examples where there is an extended irrealis context, as in (11):

(10) Every man that owns a zebra feeds it oats.
#I thought it looked miserable yesterday.

(11) Suppose there was a zebra walking down High Street.
It wouldn't be very friendly, since they're very reserved in public.
cf.: #It wasn't very friendly….

In (10) there is no particular zebra under discussion, but under the scope of the universal *every* use of *it* is felicitous. As Chierchia (1995) points out, *every* is a conservative generalized quantifier, i.e. *every*(φ,ψ) is equivalent to *every*(φ,φ&ψ). Thus, if φ has an existence entailment, then that is available to satisfy presuppositions triggered in the constituent whose denotation is ψ. In utterance of the second sentence of (11), the pronoun *it* would be taken by a competent speaker to mean something like 'the zebra walking down High Street'. But of course, there is no zebra on High Street in Columbus, probably never has been nor ever will be. Still, we intuitively

understand the felicity of the anaphora here via the assumption that in the type of hypothetical situation invoked by the first utterance there *is* a zebra. The modal auxiliary *would* in the second sentence then crucially facilitates continued discussion of that irrealis scenario (modal subordination, Roberts 1989), and since we know there is a zebra there, and that zebras can be referred to in English using neuter pronouns, and since the hypothetical zebra is the only salient entity so referred to, we can use that information about the hypothetical zebra to Retrieve the intended interpretation of the pronoun. We see that these contexts are local—only obtaining under the scope of *every* or *would*, in the infelicity of anaphora to the zebra beyond the scope of *every* or without the *would*, as in the infelicitous follow-ups marked with '#'.

Karttunen (1976) was concerned with phenomena along these lines when he pointed out that we cannot say in general that what licenses anaphora is prior reference to a thing in the actual world, since we very commonly find anaphora to fictional entities or to entities which were only inferred to exist in intensional contexts, such as we saw in (11). We may also keep track of information about someone we believe to be real, not knowing whether this individual is the same as some other with whom we are familiar, perhaps later merging information about the two. Note that on this characterization, anaphora is not about coreference, and familiarity isn't at all the same thing as acquaintance. One can be familiar with an entity in this sense without even knowing whether it exists, let alone having any (direct or indirect) contact with it. One may also have a discourse referent which is not about a particular entity in the world: Definites are tools for extended "attribution" and not for "reference". Consider:

(12) There was at least one dog in my garden while I was at work yesterday.
If its owner was with it, it wasn't well-supervised.
It trampled the pansies, left a calling card on the lawn, and pee'd on the poor barberry bush at the corner of the lot.
I think it also chased my cat, because when I came home, the cat was stuck up a tree.

The speaker of (12) clearly has no particular dog in mind, so the uses of *it* throughout cannot be taken as referential. Instead, these anaphoric pronouns are used to attribute further predicates to the dog entailed to exist by *at least one dog*. The speaker intends us to keep track of a bundle of information about a single entity, whoever it may be. Note that one or more of the later instances of *it* may be replaced by *the dog*, with no change in felicity or sense, so far as I can tell. So pronouns and definite descriptions are alike in this respect.

To distinguish such bundles of information from actual entities in the world, Karttunen called them *discourse referents*. These are the entities "on file" for all the interlocutors in Heim's (1982) File Change Semantics. Heim equivocates to some extent about what it takes to license the introduction of a discourse referent, but in my work on definites I claim that they are licensed by local existence entailments, whether or not introduced by an overt NP antecedent. Following Karttunen and Heim, I also call such bundles of information *discourse referents*:

Discourse referent: the entification of an existence entailment in the interlocutors' local context of utterance.

Heim (1982) and Kamp (1981) gave technical expression to this notion by modeling discourse referents as constraints on the assignment functions interlocutors can use to interpret variables at a given point in discourse. We assign an "address" to each bundle of information we take to be about a single (actual or hypothetical) individual whose existence is entailed by CG, the address being a particular variable, say x_{28}. Then we require that any assignment function we use to interpret utterances at that point in the discourse, including those with free variables, be such that

whatever it assigns as the value of x_{28} be an individual which satisfies all the information we have about that (actual or hypothetical) individual discourse referent. Hence, instead of assignment functions being arbitrarily chosen, they are used to encode the information about familiar discourse referents, permitting a continuity of reference. The information associated with the existence entailment in CG constrains how we can assign a value to any nominal which takes that DR-address as its antecedent.[17] But it doesn't uniquely determine a referent in the world, since the information one has to date may fail to identify a unique entity which satisfies all the information we have about that discourse referent. If it turns out for (12) that *multiple* dogs were in the speaker's yard on the day in question, each of which did all the things predicated of the arbitrary dog there (trampling, pooping, peeing, chasing), then what the speaker has said does not strike us *false* or even infelicitous.

In view of the fact, illustrated by some of the examples in the first section, that not all definites require overt antecedents, I take it as crucial to the success of a general familiarity-based theory of nominal anaphora to take the antecedent which satisfies the anaphoric presuppositions triggered by a definite to be not a noun phrase (NP) uttered in prior context, but a discourse referent representing information shared by the interlocutors, information assumed to be about a single (real or imagined) individual:[18]

The **antecedent** of an anaphoric nominal is the discourse referent which satisfies its anaphoric presupposition. The information in the local context about that discourse referent must be compatible with the descriptive content of the anaphoric expression.

I hasten to add that adopting the notion of a discourse referent from Heim (or the closely-related notion of a *reference marker* from Kamp 1981) does *not* commit one to assuming that the interpretation of definites involves unselective binding, which I take to be very problematic.[19] Assuming that definites are anaphoric in the sense just sketched merely commits one to the claim that they carry a certain sort of presupposition. This is perfectly compatible with the assumption that the proffered content of a definite (say, its contribution to the truth conditions of an utterance in which it occurs) is parallel to that of an indefinite, e.g. a type of existential generalized quantifier:

(13) proffered content of *the zebra*: $\lambda P_{<e,t>}.\exists x[\text{zebra}(x) \wedge P(x) \wedge x = y]$

(14) proffered content of *it*: $\lambda P_{<e,t>}.\exists x[\text{atomic}(x) \wedge \text{neuter}(x) \wedge P(x) \wedge x = y]$

where *atomic(x)* and *neuter(x)* in (14) are shorthand for requirements that the entity in question can be appropriately *described as* atomic and neuter. The last conjunct in each of these formulae, $\wedge x = y$, introduces a free variable *y*. It is this free variable which triggers the anaphoric familiarity presupposition, the assumption that *y* is the value of a familiar discourse referent in DR, thereby distinguishing these definites from indefinites. Then (13) is the set of properties which hold of some zebra or other *which is the same entity as the value of y under an assignment function admissible in the local context*. This denotation differs from that of *a zebra* only in the last conjunct, the portion of the paraphrase in italics. Since there may be more than one

[17] A full discussion of the technical details of proposals for modeling the dynamics of interpretation in discourse is beyond the purview of the current paper. See the references cited, from Kamp and Heim through more recent work on compositional dynamic interpretation due to Muskens (1996) and many others.

[18] This is also important in extending the account to one of demonstratives, as in Roberts (2002).

[19] For a good overview of the problems involved with unselective binding, including the so-called *proportion problem*, see Chierchia (1995).

admissible assignment function, the denotation isn't referential. (14) is similar, but with much less descriptive content

Across discourse, the DR, in its technical realization as the set of permissible assignment functions in the global context, is managed dynamically to reflect the way that information changes across discourse: If we add more information about a particular discourse referent—e.g. through predicating properties of anaphoric NPs that take that discourse referent as antecedent, the permissible assignment functions will be further constrained to reflect that information about the discourse referent in question. If a discourse referent is introduced in the course of instantiation of a generalization or is hypothetical, triggered by information under the scope of a modal or other operator, then the new information is only added provisionally, as it were—constructing a merely local context; once the hypothetical mode has been closed off, outside the scope of that operator, the discourse referent is no longer available for reference, and accordingly, the relevant constraints on assignment functions no longer hold.[20] As in the earlier theories, this is the key to interpretation of donkey pronouns like *it* in (1). Apparent exceptions to this generalization, as in (11), have been discussed extensively elsewhere (Roberts 1989 and much subsequent literature), where the general conclusion is that they involve some sort of domain restriction on the subsequent operator (here, *would*), which serves to continue the suppositional local context involving existence of a zebra.

Just as we saw that in principle the QUD is a dispensable component of the Scoreboard for a discourse, so we might argue that the set of discourse referents DR is dispensable, since the CG (or its local modification) carries information about what is said to exist in any given situation discussed. But again, the entification represented by DR is at least convenient, and arguably yields a more perspicuous account of anaphora. As with the QUD, DR has a distinguished role in discourse, in the licensing and resolution of anaphora via familiarity. As we just saw, like QUD but under different circumstances, DR is non-monotonic under update: discourse referents may be added in the local context under the scope of an operator, but be no longer available once we leave that local context.

Moreover, distinguishing the DR from the corresponding existential propositions in the CG permits us to keep track of relative salience in discourse via a contextually-driven partial order on DR. As we will see shortly, this is also crucial to appreciating how a speaker guarantees Retrievability of the intended antecedent.

Summarizing to this point, I have claimed that English definite NPs carry a presupposition of weak familiarity. Note that this is the presuppositional counterpart to the existence implication associated with the Russellian treatment of the definite description. One of the virtues of treating the familiarity presupposition as weak, instead of strong, is that obviates treating examples of anaphora without overt NP antecedents as involving accommodation (in the sense of Lewis 1979). Accommodation does sometimes play a role in satisfying familiarity, as in other types of presuppositions, but on the present account its role is far smaller and more restricted than is often assumed because the notion of familiarity is *weak*, only requiring existential entailment and not explicit prior mention. So, for example, no accommodation is required in any of the examples (1) – (6) at the beginning of this paper, since the context arguably makes evident—introduces into the CG and DR—the intended discourse referent antecedent for the definite.

[20] It is natural to think of this as parallel to hypothetical assumption in natural deduction.

We have, then, a presuppositional treatment of existence. What then of the Russellian uniqueness implication?

Roberts (2003) is an extended exploration of uniqueness in definite descriptions and pronouns, and Roberts (2005) considers yet more data pertaining to pronouns. I concluded there that assuming Russellian uniqueness cannot adequately account for the distribution of the attested uniqueness effects.[21] First, to the extent that we do observe uniqueness effects, they are clearly presupposed, and not part of the proffered content of the utterance in which they occur. We see this in the following:

(15) [Talking about a NYC apartment building:] A pleasant young boy lives upstairs. Also, a strange man lives here. If the man sees a cat, he screams. (after Kadmon 1987)

If *the man* in the conditional carried a uniqueness entailment, this would yield the wrong truth conditions for the conditional, wherein its truth depends in part on whether the strange man in question is the unique (strange) man (who lives in the location in question). Intuitively, the uniqueness of the man (even under generous assumptions about its descriptive content, in the parentheses just preceding) is not at issue in the conditional. But worse, the Russellian story, even if presuppositional, predicts uniqueness where it does not arise. One place where this is evident is in (16), and in its counterpart with a definite description in (17):

(16) Everybody who bought a sage plant here bought eight others along with it. (Heim 1982)
(17) Everyone who bought a sage plant or a rosemary planted the sage plant with extra bone-meal or the rosemary in a well-limed soil, (and if it was a sage plant, bought eight others along with it). (Roberts 2003)

It is clear that requiring that the antecedent of *it* in (16) or *the sage plant* in (17) would lead to a contradiction where intuitively there should be none.

However, I concluded that there are uniqueness effects in several kinds of examples involving definite descriptions, and that this should be captured by any adequate theory of the English definite article. Two kinds of cases involve descriptions used as titles, as in (18) and (19), and relational descriptions, as in those in implicit relation with the quantificational subjects ("bridging") in (20)-(22):

(18) The Ohio State University is in Columbus.
(19) The Queen of England had a bad year in 1993.
(20) Every car has a statue on the dashboard.
(21) Every car had a puncture in the tire.
(22) Every unicycle had a spoke missing from the wheel.

In (20) and (22), the definite description picks out the unique dashboard per car or wheel per unicycle; most speakers find (21) less felicitous, presumably because there is no obvious indication of how to pick out a unique tire per car.

In view of these and other examples discussed in detail in Roberts (2003), I argued that along with a presupposition of weak familiarity, what we needed to account for the attested uniqueness

[21] The reader is referred to Roberts (1983) for a more detailed argument against Russellian uniqueness. Here I can only offer highlights. See that paper as well for more detailed discussion of a wider variety of uniqueness effects and how those are accounted for on the theory I propose.

effects was not semantic uniqueness à la Russell, but what I called *informational uniqueness* relative to the property denoted by the NP's descriptive content:[22]

> **Informational uniqueness of a discourse referent relative to a property *P***: Given a discourse D with Common Ground CG and discourse referents DR, a discourse referent *d* ∈ DR is informationally unique in DR with respect to a property *P* iff *d* is the unique element of DR which is entailed by CG to have property *P*.

In the cases just considered, informational uniqueness suffices to explain the uniqueness effect associated with the definite descriptions. For titles and institutions, we understand that they are *designed* to denote only one entity at a time. In the bridging examples, since we know that typically cars have exactly one dashboard, unicycles a single wheel, we can assume that for any instance *x* of the quantificational subject in (20) or (22) this default information entails the existence of exactly one dashboard/wheel, and hence the corresponding discourse referent satisfies both weak familiarity and informational uniqueness under the descriptive content *dashboard of x*/*wheel of x*. Since we know that most cars have five tires (including the spare), informational uniqueness under the description *tire of x* fails for any car *x* in (21), explaining the reduced felicity in that example.

Of course, on this account, the problem of incomplete descriptions persists. The relational character of the descriptive content in (20)-(22) is merely implicit, just as for that of *the man* in (15).

3.2. Uniqueness via Retrievability

Recall that on the understanding of *anaphora* just sketched, both English definite descriptions and personal pronouns are anaphoric, in that they carry a presupposition of weak familiarity. Then since by hypothesis Retrievability constrains the felicitous use of all anaphoric expressions, use of such definites is subject to Retrievability as well. In this section, I argue that for definite NPs, Retrievability under anaphora entails informational uniqueness, so that such uniqueness need not be stipulated.

Consider again the definition of Retrievability, repeated here:

> **Retrievability**: In order for an utterance to be rationally cooperative in a discourse interaction *D*, it must be reasonable for the speaker to expect that the addressee can grasp the speaker's intended meaning in so-uttering in *D*.

Note that what must be reasonable for the speaker to expect is that the addressee can grasp *the* intended meaning. I.e., this presupposes that there *is* a single semantic intention in making a given utterance. This rules out intentional ambiguity on the part of a cooperative speaker.[23]

---

[22] Here I give a simplified definition, only considering uniqueness with respect to DR and CG. A generalization to cover non-global, i.e. local contexts, as discussed above, is straightforward if one adopts a particular theory of dynamic interpretation.

[23] Chris Kennedy (p.c.) raised the question of whether this claim can be maintained for vague predicates and vague modal meanings (e.g., those under investigation in von Fintel & Gillies 2008), and in cases involving under-determined quantifier scope. In each of these, an argument might be made that the speaker's semantic intentions themselves are rather ill-defined. With respect to vagueness, if the truth conditions as intended by the speaker are vague, as is arguably the case with modals quite generally, then

If Retrievability presupposes that the speaker has a unique semantic intention in making her utterance, and Retrievability is required for felicity in anaphora, then the felicitous use of a definite, and hence anaphoric NP presupposes that in the context of use the speaker can reasonably expect the addressee to grasp her uniquely intended meaning for the NP via Retrieving its intended antecedent, a weakly familiar discourse referent. I.e., one can assume that the definite's descriptive content suffices to pick that particular discourse referent out as the intended antecedent. If there is more than one suitable candidate for antecedent, then the use of the definite is ambiguous and, so, irRetrievable. I.e., on this account **informational uniqueness follows from Retrievability**, and hence we expect it to arise in interpreting *any* anaphoric expression: Anaphora resolution is meaning$_{nn}$ RETRIEVAL, and that always requires a unique resolution of the speaker's semantic intentions.

Note that this aspect of Retrievability is not particular to (pro)nominal anaphora, but is common with various kinds of ellipsis and with other arguably anaphoric presuppositions (Beaver & Zeevat 2007), like that of *too*. In VP ellipsis, a particular property must be Retrieved (Merchant 2005; see Roberts, in preparation, for discussion with respect to Retrievability). Kripke (2009) makes the point that *too* also requires a particular antecedent. E.g., for examples like (23), general knowledge to the effect that lots of people had dinner in NY the previous evening will not license felicitous utterance of this example; instead, there should be some particular individual such that her/his having had dinner in NY the previous evening is familiar from prior context:

(23) [John]$_F$ had dinner in NY last night, too.

In order to make this a convincing proposal, we need to explore the question of how Retrieval is accomplished: In a given discourse, there may be many, many Discourse Referents—entities which are weakly familiar given the CG. So how do we generally Retrieve the intended antecedent for a definite? This question is especially pressing for pronouns and for definite descriptions whose descriptive content is minimal. I will argue that Retrieval generally depends upon the structure of the discourse as reflected in the Scoreboard described in section 2.

First, note that informational uniqueness is not a problem for an anaphoric expression whose descriptive content uniquely identifies a single individual in the world, i.e. a definite which by virtue of that content displays Russellian uniqueness. If the definite's anaphoric presupposition is satisfied by a familiar discourse referent antecedent in DR, then the fact that in the CG it is known that this DR represents an entity which is unique in the world with respect to the NP's descriptive content suffices to guarantee that this will be the only discourse referent which is a candidate antecedent: We only have one moon, commonly familiar to all and known by all to be unique in being our moon, so it's no trouble to Retrieve the unique intended discourse referent for the definite description *the moon*. That is, Russellian uniqueness guarantees informational uniqueness relative to the descriptive content of the definite itself, and hence guarantees

---

what is Retrieved is itself vague, and rightly so. In this connection, Kratzer (1981) argues that since the use of modals is often vague with respect to the intended domain (as given by her contextually Retrieved modal base and ordering source), the criterion of adequacy for a semantic theory of modals is not lack of vagueness, but clarity about where the vagueness lies and how it arises. If we suggest an intended modal restriction with *In view of what the law requires*, since what the law requires is such a contentious issue we cannot be faulted for interpreting the utterance as a vague statement. Moreover, just because truth conditions are vague or fuzzy, this doesn't mean that the meaning in question isn't uniquely intended. Put another way: Being compatible with one or more precisification doesn't make the intended meaning *ambiguous*.

Retrievability. This tends to be the case for superlative NPs as well, like *the tallest mountain in West Virginia*.

Moreover, the descriptive content of the Definite NP itself is not the only clue to its intended discourse referent antecedent. As Barbara Abbott (p.c.) reminded me, Evans (1977) noted that the relevant property for uniqueness may be partly given by whatever is predicated of the NP, as well, as illustrated with examples like the following:[24]

(24) The administration denied the protesters a permit because they were violent.
(25) The administration denied the protesters a permit because they were afraid of violence.

In these cases, we might say that plausibility of the given predicate plays a role in restricting one's choice of antecedent for *they*. Protesters are stereotypically more violent than administrators, the latter stereotypically more fearful of violence by virtue of representing the status quo. It isn't surprising that plausibility plays a role in Retrieving antecedents: After all, plausibility arguably plays a central role in lexical disambiguation: It's very unlikely that you'll be depositing money at the river bank.

But in case neither the definite's descriptive content nor what is predicated of it are sufficient to guarantee Russellian, semantic uniqueness, then competent use of a definite requires that there be some other way to reasonably guarantee Retrievability.

If a speaker could somehow reasonably assume in advance that the addressee's attention could be restricted to some subset of DR, a subset in which the definite NP's intended familiar discourse referent antecedent *was* informationally unique, then this would guarantee that that antecedent was readily Retrievable. Call this the **hypothesis of the attentional restriction of the domain for Retrievability**.

There is a very large literature on attention in psychology which supports the general picture I assume about the relationship of intentions to attention. For example, in the famous experiments on *inattentional blindness* of Simons & Chabris (1999), when experimental subjects were directed to attend to groups of people passing a basketball in a video, over 50% completely missed a woman in a gorilla suit walking through the scene. This has subsequently been replicated in other carefully designed experiments. There is also good evidence, discussed in detail in Bloom (2000), that children first learn the meanings of words through tracking what their caregivers visually attend to, ultimately using that information about attention, in conjunction with their emerging theory of mind, to help Retrieve the caregiver's semantic intentions. Also relevant here is the extensive body of work due to Tanenhaus and his associates, using eye-tracking to explore on-line language processing in adult subjects. Tracking a subject's attention across a visual field provides strong evidence that the subject's goals serve as a guide to the resolution of both ambiguity and anaphora. See Tanenhaus, Spivey-Knowlton, Eberhard & Sedivy (1995); Chambers, Tanenhaus, Eberhard, Filip & Carlson (2002); Hanna & Tanenhaus (2004); and many other papers from this group. Thus, I take it as psychologically plausible that as speakers and addressees we exploit this basic perceptual mechanism—the limitation of the field of attention via intention—in order to make reference more efficient and successful in discourse.

In fact, I believe that something like this is very likely the fundamental functional motivation for our use of something like the QUD to organize discourse. I characterized the QUD as the

[24] According to Abbott, examples like these were in Winograd (1972), but I haven't verified that.

interlocutors' discourse goals, and claimed that it is subordinated to their domain goals. Then one of the ways in which Relevance to the QUD is used in discourse is to mask the attentional field for anaphora resolution: As speakers, we can be reasonably sure that a cooperative interlocutor, one who is attending to the QUD in keeping with the normal intentions of cooperative interlocutors in such a task-driven interaction, will principally pay attention to those discourse referents which are Relevant. In other terms, this is just what Barbara Grosz said about anaphora back in the 1970s (Grosz 1977; Grosz & Sidner 1986): The intentional structure of discourse determines the attentional structure of discourse—what we sincerely intend restricts what we attend to—and it is quite generally this restriction of the attention of the interlocutors which permits a competent speaker to felicitously use a definite description or pronoun whose descriptive content alone does not guarantee informational uniqueness.

In the Gricean theory of context outlined in section 2, Relevance to the question under discussion in discourse requires that interlocutors always address the QUD. Hence, in making an utterance the speaker can reasonably expect that a competent, cooperative addressee who is engaged in the discussion will be attending to the resolution of the QUD (a semantic entity—a set of alternatives). This entails that one who holds such an intention will have the QUD itself foremost among those entities being attended to. But more, we can expect that the attentive addressee will also be attending to any discourse referents which are relevant to the QUD:[25]

The set of Relevant discourse referents: In a discourse with scoreboard *S*, discourse referent *d* is Relevant to the QUD *q* just in case for some property *P*, the question of whether *d* has *P* is Relevant to *q*.

When the addressee is attending to what is Relevant, the set of Relevant discourse referents can always reasonably be assumed to be in the interlocutors' attentional field. Thus, attending to the QUD in obedience to Relevance restricts the domain of entities immediately under consideration, effectively limiting the search space for an antecedent.

Note that claiming that Relevance restricts the search space for incomplete descriptions should not be taken to suggest that we can only felicitously refer anaphorically to Relevant familiar discourse referents. This is not about *global* restriction of the domain for reference. When we have a definite description with a rich descriptive content, it can harken back to entities which are not at all obviously Relevant in the immediate discourse: Consider NPs like *that Bulgarian guy we were talking to last month at Stauff's*. Rather, what I am proposing is that when the descriptive content of a definite does *not* guarantee informational uniqueness, the addressee's attention will naturally first be drawn to those candidates for antecedent that are Relevant.

But, of course, even among the Relevant discourse referents, it is sometimes the case that two or more would satisfy the descriptive content of a definite NP. In such cases, the definite will only be felicitous, its intended antecedent Retrievable, in case the antecedent is more salient than any of its competitors. The more impoverished the descriptive content of the definite, the more likely that non-informational uniqueness within the field of Relevant discourse referents will arise. Hence, it is not surprising that pronouns tend to require much greater salience of their antecedents than do definite descriptions.

In Roberts (2003) I argued that pronouns differ from definite descriptions in requiring a higher degree of salience for the discourse referent satisfying the weak familiarity presupposition. This

[25] Of course, the question of *which* entities are Relevant is a question for *practical inference*, and the general problem of how to restrain this persists.

explains the greater felicity of (9), with a definite description, than Partee's marble example (8), repeated here:

(7) I dropped ten marbles and found all of them, except for one.
It is probably under the sofa.
(8) I dropped ten marbles and found only nine of them.
?It is probably under the sofa.
(9) I dropped ten marbles and found only nine of them.
The missing marble is probably under the sofa.

And (3) shows that Relevance to a task at hand (a domain goal) suffices to make the pronoun's antecedent sufficiently salient that it is Retrievable, while lack of such Relevance makes (3′) infelicitous:

(3) [naturally occurring example (1992): Ed Keenan is giving a colloquium at the University of Amsterdam. During the talk, he twirls his reading glasses by the stem and the screw holding the stem to the frame falls out, so that the glasses drop to the table. He puts the stem down and continues his talk. At the break, he begins searching intently under the papers on the table.]
Craige: It's probably on the floor.
(3′) [The same colloquium scenario, up to the break. But at that point, Keenan absent-mindedly picks up the stem of the glasses and plays with it while he chats with an audience member about his talk.]
Craige: #It's probably on the floor./√The screw is probably on the floor.

The different salience requirement is evident in the pattern of acceptability of pronouns vs. that of definite descriptions in the following, from Roberts (2003):

(26) A woman entered from stage left.
Another woman entered from stage right.
#The woman/√The FIRST woman/√The SECOND woman was carrying a basket of flowers.
(27) A woman entered from stage left.
Another woman entered from stage right.
SHE was carrying a basket of flowers, while #the woman/√the FIRST woman/#the SECOND woman led a goat.
(28) A woman entered from stage left.
There was a basket of flowers in the middle of the stage.
She picked it up.
(29) A woman entered from stage left.
Later in the act, a woman entered from stage right.
√?The woman/√She was carrying a basket of flowers.

In that earlier paper, I proposed that pronouns carry an extra presupposition of maximal salience of the antecedent. But in light of the present discussion, it becomes clear that this needn't be stipulated. It is merely a reflection of a functional constraint on salience: The more descriptively impoverished the trigger, the higher the degree of salience usually required to guarantee Retrievability of the intended antecedent. A given NP only requires higher salience as a function of its descriptive incompleteness relative to the field of Relevant discourse referents. Since pronouns are very descriptively impoverished, they typically require higher salience.

There are many things one could say about salience, but that discussion goes beyond the purview of the current paper. Only two points: First, Terken & Hirschberg (1994) discuss experimental evidence which argues that for anaphora resolution, recency of mention is not particularly significant in assessing salience. Instead, maintenance of surface order (first) and of thematic role (second) are much stronger indicators of relative salience. I suspect that this is a reflection of the fact that what guides anaphora resolution is *the relation of the utterance in question to those that precede it within a given discourse segment*. Relevance then would pose a general constraint on salience, but more local, micro-relations (perhaps some of them pertaining to rhetorical structures that reflect strategies of inquiry) play a role as well.

Second, in support of the centrality of the intentional structure of discourse, and its relation to Relevance and Retrievability, note that relatively abstract entities relevant to resolving the QUD may be *more* salient than visually evident entities in the situation of utterance. We see this in (30):

(30) [Context: You and I are sitting in a café discussing how to understand Sperber & Wilson's (1985) definition of Relevance, and I say:]
I see it now!
[Even though I'm holding a coffee mug by the handle right under your nose and shaking it for emphasis, you don't take *it* to refer to the concrete mug.]

More often, discourse goals and other means of increasing salience interact to define the attentional field—the Relevant discourse referents, making it possible for a hearer to Retrieve the intended discourse referent for an anaphoric element. We saw this in the intentional minimal pair in (1) and (1′), repeated here, where the deixis combined with the QUD to permit the addressee to Retrieve the intended antecedent for *that*:

(1) [Tim and Margaret are sitting at a conference table in her psycholinguistics lab, working on a grant proposal for some eye-tracking experiments. Tim is making notes on his laptop computer, not a MacIntosh brand.]
Margaret: How do you like your laptop?
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, in the direction of her desk in the middle of the room]
I just got [<u>that</u>]$_F$ 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 Mac computer monitor on it, past the desk an eye-tracker, past that a wall calendar.

(1′) Margaret: How do you like your laptop?
Tim: It's not bad, but it's getting kind of old. Anyway, let's get back to business. What do you think of this budget?
Margaret: [gesturing with over her shoulder, in the same direction as in (1)]
Could we add some money to replace <u>that</u>?

I would argue that all this follows from the simple assumption that definite descriptions are anaphoric, and hence presuppose a weakly familiar discourse referent. One need not stipulate that they also presuppose informational uniqueness, as that follows under the assumption of Retrievability. Nor need one stipulate the higher salience requirement of pronouns. Definite descriptions are typically less descriptively impoverished than pronouns (not always: consider *the*

entity), so the intended antecedent discourse referent for a description can be Retrieved in a broader intentional field, so long as the discourse referent in question is informationally unique in bearing the descriptive content of the description. Something similar can be said for the difference between demonstrative descriptions and pronominal descriptions; see Roberts (2002). None of this requires more than informational uniqueness, and that is built into Retrievability, so Russellian uniqueness is otiose.

Summarizing, anaphora resolution is constrained by the following factors:

1. the antecedent must satisfy the definite NP's descriptive content
2. the resolution must be plausible in view of what's predicated of the trigger
3. the antecedent must be the (uniquely) Retrievable discourse referent under the plausible descriptive content, either
    a. in virtue of that descriptive content alone (in which case it's Russellian unique) or
    b. if 3a is not the case:
        i. the discourse referent must be Relevant to the QUD, and
        ii. if there's more than one Relevant entity which satisfies the plausibly enriched descriptive content, the intended antecedent must be more salient than any other Relevant discourse referents which satisfy the first two constraints.

Arguably, the first constraint is the most important, and plausibility and Relevance are equally strong constraints. From a production point of view, anaphora resolution probably privileges Relevance. That is:

- the speaker knows that the attention of a cooperative addressee is restricted, via
    - QUD
    - relative salience
- Hence, the speaker uses that NP whose descriptive content most efficiently picks out *the* intended antecedent, a weakly familiar discourse referent.
    - if the antecedent is Relevant and sufficiently salient relative to competitors, the speaker uses a pronoun.
    - if only Relevant, she uses the poorest descriptive content which will make the antecedent informationally unique among the Relevant discourse referents,
    - if not Relevant, she uses sufficient richness of description to make the antecedent informationally unique in CG.

Notice that in this theory the question of incompleteness as discussed by Neale does not arise, because:

a) Definites in general neither presuppose nor entail semantic uniqueness. Hence, there is no conventional uniqueness implication associated with a definite description.
b) Pronouns are not disguised definite descriptions, unlike in the E-type theories.

All that's required is sufficient descriptive content to make the intended anaphoric antecedent Retrievable.

4. Comparison with proposals involving semantic uniqueness

The majority of those who work on definites in the modern literature on semantics and the philosophy of language treat them as involving semantic uniqueness, whether presupposed or proffered. In this section, I compare the theory described in the preceding section to those of two

prominent proposals involving semantic uniqueness: Neale (1990,2005) and Elbourne (2001,2005). These are the most detailed and successful theories along these lines that I am familiar with, and, as we shall see, they display certain differences that are of interest as well.

One goal these two proposals have in common is one they share with most contemporary approaches, including the proposal in the preceding section: This is the desire to capture the semantic similarity between definite descriptions and personal pronouns; while these differ in some respects, the two classes of NPs share very similar distributions, often yielding identical truth conditions under substitution. This is captured in the theory proposed here by the assumption that both types of NP are anaphoric to a weakly familiar discourse referent. Most theories involving semantic uniqueness take a different tack, treating at least one class of pronouns as *disguised definite descriptions*: those with an apparent quantificational NP antecedent which does not stand in a position to bind the pronoun. Evans (1980) calls such pronouns *E-type*, others, including Neale (1990), call them *D-type*. Authors who take this general approach include:

Quine (1960)	Kadmon (1987,1990)	van der Does (1993)
Parsons (1978)	Lappin (1989)	Lappin & Francez (1994)
Cooper (1979)	Neale (1990)	Chierchia (1995)
Evans (1980,1982)	Heim (1990)	Elbourne (2005)
Engdahl (1986)		

The technical details of their proposals differ considerably: E.g., some authors treat the pronouns in question as having the semantics of a Russellian definite whose impoverished descriptive content receives contextual enrichment; others as a denoting a variable over functions whose contextually given value (via presupposition resolution) is truth conditionally equivalent to that of the Russellian definite description; others treat the pronouns as syntactically elliptical for a full definite description. In the present section, I want to argue that the account proposed in the previous sections is superior to these E-type/D-type theories, or to any account positing lexical uniqueness (whether proffered/"Russellian" or presupposed/"Fregean") for pronouns. For detailed discussion and critique of more of these theories, including empirical problems beyond those considered here, see Roberts (2005) (which was written before I knew of Elbourne's work).

Neale's (1990) treatment of definite descriptions as restricted quantifiers in (31) is "definitionally equivalent" to the standard Russellian truth conditions in (32). His parallel treatment of D-type pronouns is shown in (33) (his (P5)):

(31) '[the x: Fx] (Gx)' is true iff |F-G| = 0 and |F| = 1[26]

(32) $(\exists x)((\forall y(Fy \equiv y=x) \,\&\, Gx)$

(33) If *x* is a pronoun that is anaphoric on, but not c-commanded by, a quantifier '[*Dx*: *Fx*]' that occurs in an antecedent clause '[*Dx*: *Fx*](*Gx*)', then *x* is interpreted as the most "impoverished" definite description directly recoverable from the antecedent clause that denotes everything that is both *F* and *G*.

Hence, Neale takes both the descriptive content of an antecedent NP and whatever is predicated of that NP to be the understood descriptive content of the Russellian definite which is the interpretation of the D-type pronoun.

[26] This is the case where *F* is singular. Where *F* is plural, Neale proposes to replace "|F| = 1" with "|F| > 1".

In Neale (2005) he extends the treatment of definite descriptions to one which is intended to cover both attributive and referential uses, using what he calls *Gödelian completions*. Effectively, these involve adding to the restriction in the LF in (31) the condition that $x = a$, where a "stands for" the individual referred to. Hence, (31) is intended as a general theory of the semantics of definite descriptions and D-type pronouns.

Neale (2005) makes it clear that where incomplete descriptions are concerned, this theory does not assume syntactic ellipsis—there is nothing syntactically incomplete about these NPs. Rather, we are driven to infer the speaker's intended meaning—*what she said* in the Gricean sense, on the basis of the conventional content of the utterance and (presumably, by Gricean Quality) the assumption that the cooperative speaker intends a meaning of the definite which is complete, yielding a true utterance. He offers a comment which is very similar in spirit to the principle of Retrievability:

> If, as Grice suggests, what *A* meant by uttering *X* on a given occasion is determined by certain interpreter-directed intentions, then assuming he is being co-operative *A* cannot mean that *p* by uttering some sentence *X* if he believes it is impossible for his audience *B* (or at least any rational, reasonably well-informed interpreter in *B*'s shoes) to construe him as meaning that *p*. (77)

In bridging the gap between incomplete explicit content and what was felicitously meant:

> The slack is taken up by a *pragmatic* theory… [I]dentifying what *A* said involves processing not only the semantic information encoded in a sentence's form, but accessing and processing information that must be picked up by listening, watching, remembering, hypothesizing, and inferring, essentially the capacities exercised in identifying what *A implied*. To this extent, then, identifying what is said is a *pragmatic* as well as a semantic matter. It involves *pragmatic inference* as well as *linguistic decoding*. (86-87)

Without giving details about how the required information comes to bear on interpretation, he claims that we identify what the speaker said via "knowledge of the linguistic meaning of the sentence uttered, pragmatic knowledge about the way rational, cooperative beings operate, knowledge about the speaker, knowledge of context, and just about anything else he can get his hands on" (91). He calls this the *explicit approach* to utterance ellipsis generally and to incomplete definites in particular, contrasted with the *implicit approach* favoring global domain restriction of the sort proposed by Barwise & Perry (1983).[27] On the explicit approach, "all of the work is done by pragmatic inference" (123). One needs to infer in what respect the pronoun should be understood as "shorthand for, elliptical for, an abbreviation of at least one richer nominal the speaker *could have* used and *could* produce if asked to be more explicit" (121) I.e., this is a "modal" theory of pragmatically recovered incompleteness.

Elbourne (2005) argues for a "Fregean" analysis of definite descriptions, according to which existence and semantic uniqueness is presupposed, rather than proffered. Here is his (simple, extensional) semantics for the definite article (in a Heim & Kratzer (1998)-style LF) (p.98):[28]

---

[27] I won't discuss theories like Barwise & Perry's here. See Soames (1986) and Kratzer (2007/rev.2009) for discussion of their treatment of incomplete descriptions, and the latter for discussion of its relationship to a situation semantics approach like that in Elbourne (2005).

[28] Elbourne ultimately argues for what he calls the "add-a-variable-to-the determiner" theory of *the* (p.113), in which a second argument is added, an index used to capture relational interpretations of definite

(34) Elbourne's semantics for *the*, extensional version:
λf: $\underline{f \in D_{<e,t>} \ \& \ \exists!x\ f(x) = 1}$.$\iota x\ f(x) = 1$

This is a function from sets of entities (functions in <e,t>) to individuals; the argument is the denotation of the NP complement to the determiner *the*.[29] The underlined portion is a presupposition: 'exactly one thing is *f*', while the proffered content, the value when the NP complement is saturated, is that unique individual which has the property *f*. The Russellian existence and uniqueness entailed by the proffered content is redundant of the Fregean presupposed content, but there's no harm in that: They just guarantee that the unique individual presupposed to exist *is* the denotation of the DP.

Elbourne then proposes an NP-deletion theory of D-type anaphora. Pronouns are definite articles with syntactically present but phonologically null NP complements. Here is his denotation of *he* as a Fregean definite article (extensional version):[30]

(35) Elbourne's semantics for *it*, extensional version:
λf: $\underline{f \in D_{<e,t>} \ \& \ \exists!x\ f(x) = 1}$.$\iota x\ f(x) = 1$

As in non-pronominal *the*, the denotation of *he* takes a function to yield an individual, carrying a presupposition to the effect that there is exactly one individual which has the property corresponding to the function, and its proffered content is the unique individual which has that property. As in ellipsis generally, the addressee must rely on contextual information to determine the elided NP argument.

To deal with the problem of incomplete descriptions in a theory of semantic uniqueness, Elbourne borrows a tool suggested by Heim (1990): the use of minimal situations to limit the domain of the definite. The stage for this account is set in the intensional version of the semantics for *it* in (36), which takes an intension to be a function not from entire worlds, but from situations to denotations of the appropriate extensional type:

(36) Elbourne's semantics for *it*, intensional version:
$||it||^g = \lambda f_{<<s,e>,<s,t>>}.\lambda s{:}\underline{\exists!x\ f(\lambda s'.x)(s) = 1}.\iota x\ f(\lambda s'.x)(s) = 1$

it takes a function which is the denotation of the implicit NP complement, to yield a function from situations to individuals. A function from situations to individuals is an individual concept (in the sense of Montague 1973); hence, the full DP denotes an individual concept. The first, NP argument itself takes an individual concept to yield a proposition—a function from situations to truth values. The underlined portion of (36) tells us that *it* presupposes that there is exactly one individual *x* s.t. the rigid individual concept yielding *x* as its value for all situations has the property denoted by the NP in the situation of evaluation *s*. The proffered content, after the underlined portion, is the unique individual such that it satisfies this functional description at *s* (i.e., this is the presupposed individual).

---

descriptions under quantification. While important for developing a general theory of descriptions, it isn't directly relevant here, and it is subject to the same problems I point out for the simpler theory.

[29] For Elbourne, the entire definite description is a DP, headed by the determiner *the*. The complement of the determiner, an NP, denotes the DP's descriptive content. Thus we have [$_{DP}$ the [$_{NP}$ man]].

[30] Elbourne assumes that English pronouns may include a gender restriction on their denotations (e.g. Chapter 2:44), but we ignore these here for simplicity.

This treatment of definite descriptions and pronouns is then embedded in a situation semantics utilizing minimal situations (following Berman 1987, in a situation semantics along the lines of Kratzer 1989), to guarantee that we don't get too much uniqueness. Without going into unnecessary technical detail here, the classical donkey sentence in (37) gets the interpretation in (38) (p.63):[31]

(37) Every man who owns a donkey beats it.

(38) λs_4. for every individual *y*
for every minimal situation s_5 such that
$s_5 \leq s_4$ and *y* is a man in s_5 and there is an individual *x* and a situation s_2 such that s_2 is a minimal situation such that $s_2 \leq s_5$ and *x* is a donkey in s_2, such that there is a situation s_3 such that $s_3 \leq s_5$ and s_3 is a minimal situation such that $s_2 \leq s_3$ and *y* owns *x* in s_3,
there is a situation s_6 such that
$s_6 \leq s_4$ and s_6 is a minimal situation such that $s_5 \leq s_6$ and *y* beats in s_6 ιz *z* is a donkey in s_6

Even though a given man may own several donkeys, (37) can be true, and the use of *it* to refer to *the donkey the man owns* be felicitous, because we are only quantifying over universally quantifying over *minimal* situations containing a man and a donkey he owns (s_5) (and so, no other entities). To verify (37) under the interpretation in (38), for each such minimal s_5 we only need to find *some* super-situation s_6 which includes the man and the unique donkey in s_5 and in which the man beats the donkey. Hence, the presupposed uniqueness implication doesn't cause problems for a number of kinds of cases which Heim (1982) had argued weigh against semantic uniqueness in the Russellian treatment of definites. The proposed reading is intended to have the same truth conditions for donkey sentences that we find in Kamp (1981) and Heim (1982), wherein we're essentially saying that any given farmer beats *all* his donkeys; but here instead of *every* unselectively binding a variable corresponding to *a donkey*, it does so indirectly by quantifying over minimal situations *containing* a donkey the man owns.

Thus, though Neale and Elbourne both assume semantic uniqueness, both aim to account for the semantics of a broad class of pronouns, and neither bases his theory of definites on familiarity, the two theories display some important differences:

- Neale assumes that the elliptical character of incomplete definites is essentially pragmatic, without any necessary syntactic reflexes, but Elbourne proposes that at least some incomplete definites—the pronouns—involve syntactic ellipsis.

[31] Elbourne also assumes non-persistence: Just because a proposition is true in situation *s*, that doesn't guarantee that it's true in all supersituations *s'*. This assumption is crucial for his treatment of examples involving apparent violations of uniqueness, including the indistinguishable-participants examples. There are a number of potential difficulties for our understanding of (38) which arise from non-persistence. For one thing, just because it's specified that *y* is a man in s_5, does it follow that *y* is a man in the super-situation s_6? Absent persistence, I don't see why, unless there are certain metaphysical principles (e.g.: something that's of kind *k* in a given situation *s* must be of kind *k* in any synchronous super-situation *s'*) which always apply to situations ordered by inclusion. It may seem obvious *prima facie* that that's desirable, but there are lots of kinds of properties which might cause problems: Just because John is the biggest football player in s_1, does that mean that he's the biggest football player in the larger situation s_2? There are other such worries when we add in times as well, but this gives the flavor of my concern.

- On Neale's account, semantic uniqueness is proffered, while on Elbourne's it is presupposed.
- Neale leaves the resolution of incompleteness to pragmatics, invoking a variety of pragmatic processes which lead to "enrichment" of the descriptive content of an incomplete definite. But for a very prominent sub-class of the definites, the D-type pronouns, Elbourne builds a resolution for incompleteness into the syntactic LF, using minimal situations and non-persistence to guarantee completeness

In what follows, I list some of the central problems with one or both of Neale's and Elbourne's accounts, and, where they have acknowledged these problems, their responses to these issues:[32]

4.1 Problems getting the right truth conditions

Russellian definites. by virtue of the uniqueness clause, are superficially too strong in a great many uses. At the extreme, if you take the descriptive content of the pronoun *he* to be 'is a singular male', then it's hard to see how uniqueness could ever be satisfied; but things aren't much better for the very common *the man*. This is, of course, the heart of the "problem of incomplete descriptions". E-type or D-type theories owe an account of how this problem is overcome. Neale's (1990,2005) response, as we just saw, is to claim that the same kinds of processes that yield pragmatic domain restriction generally account for the pragmatic enrichment of descriptive content which ameliorates this problem. Insofar as this is essentially a pragmatic account of the incompleteness problem, it is far from explicit how he thinks this resolution would proceed in particular cases. In that respect, the present proposal would appear to be more clearly falsifiable, hence stronger.

Elbourne follows Heim (1990) in overcoming the problem of incomplete descriptions by building in severe domain restriction in the LF of definite descriptions, including that of D-type pronouns, restriction in the form of minimal situations. Hence, Elbourne's iota operator is effectively much weaker than it usually is with respect to Russellian uniqueness: Although something is unique-under-its-description in the minimal situation of interpretation, this doesn't make it unique in the intuitive Russellian sense, i.e. in the whole world. So in a way Elbourne gets rid of the problem of incomplete descriptions by denying uniqueness. One might wonder what advantage it has over theories which don't posit uniqueness to begin with.

The theory would appear to make a clear difference in predicted felicity and truth conditions in cases where the D-type pronoun is not subject to the quasi-unselective binding that leads to quantification over minimal situations. (12), repeated from above, would appear to be such a case:

(12) There was at least one dog in my garden while I was at work yesterday.
If its owner was with it, it wasn't well-supervised.
It trampled the pansies, left a calling card on the lawn, and pee'd on the poor barberry bush at the corner of the lot.
I think it also chased my cat, because when I came home, the cat was stuck up a tree.

On Elbourne's account, presumably each occurrence of *it* takes the preceding NP *dog in my garden (while I was at work yesterday)* as antecedent for its elided complement (with possibly further descriptive content enriching that in subsequent occurrences). But since quantification

[32] A number of other problems for Neale's treatment of D-type pronouns, and for other E-type or D-type theories of pronoun interpretation are discussed in Roberts (2005).

over minimal situations doesn't come into play here (or only irrelevantly so in the conditional), this predicts that we get the full force of uniqueness in the interpretation of each instance: So *it chased my cat* presupposes that there was a unique dog in my garden yesterday, and proffers 'the unique dog in my garden yesterday chased my cat'. This is incompatible with our understanding of the example, under *at least one dog…* in the first sentence.

Beyond the problem of incomplete descriptions, the accounts of donkey anaphora in both theories are problematic. Consider the truth conditions Elbourne predicts, in (38) above. Elbourne (2005:23) points out that these show how the minimalism in his theory is "too strict": "There is a problem [for D-type theories] in that the most straightforward situation semantics used in conjunction with the D-type theory predicts only strong readings [of donkey sentences]. But this is likely to be a technical problem with that situation semantics, perhaps having to do with too rigorous an insistence on strictly *minimal* situations being quantified over." I.e., Elbourne predicts that donkey sentences will only have "strong" readings of the sort usually associated with (37), and never "weak" readings like that of Pelletier & Schubert's (1989) (39) (and see Chierchia 1995):

(37) Every man who owns a donkey beats it.
(39) Every man who had a quarter in his pocket put it in the meter.

This is because the universal quantifier *every* in (37) is quantifying not over farmers alone, but over farmer-minimal situation pairs, getting some of the effects of unselective binding, as we noted. Moreover, for this very reason it seems to me (though Elbourne doesn't discuss this) that extending his proposed truth conditions to proportional quantifiers like *most* or *usually* will lead to the same "proportion problem" which was a serious issue for Kamp (1981) and Heim (1982). This is illustrated by examples like:

(40) Most farmers who own a donkey beat it.

Consider a situation with ten donkeys, nine of whom have one donkey each and do not beat it, one of whom owns ninety-one donkeys and beats them all. Unselective binding, and, so far as I can tell, Elbourne's account, would incorrectly predict that (40) was true in such a case. Heim (1990) proposes a fix for her minimal-situations theory, which Elbourne (2005:140-141) criticizes, for it fails to extend to examples with indistinguishable participants; but he fails to make an alternative proposal.

Neale, too, predicts only the strong readings of donkey sentences, never the weak, though he doesn't encounter the proportion problem (see his discussion, Chapters 5 and 6). Roberts (2005) proposes that it is the *weak* readings of donkey sentences, like that in (39), which are basic, and that the strong reading displayed in (38) is derived by pragmatic strengthening.

4.2 The use of anaphoricity in Elbourne's account

As we just saw, the use of minimal situations is crucial in permitting Elbourne to avoid some of the problems with semantic uniqueness. But he points out (2005:59) that this makes (41) come "dangerously close to meaning [(42)]":

(41) Every man likes the woman.
(42) Every man likes a woman.

Under his truth conditions, the individual denoted by *the woman* need only be unique in some extension of a minimal situation containing one of the men in the restriction, leaving open the possibility that there are many such extensions. Hence, the conditions for uttering (41) turn out to be no more stringent with respect to the *semantic* uniqueness of the denoted woman than are those for (42) with the indefinite! And Elbourne has no general way of differentiating definites from indefinites.

To address this serious issue, Elbourne adopts the view he attributes to van der Sandt (1992), Zeevat (1992) and Beaver (1997), that the presuppositions triggered by definites are *anaphoric*, so that "the necessary information must somehow be contextually salient". In his very brief discussion of this matter, Elbourne fails to mention the familiarity theories of Kamp (1981) and Heim (1982), though van der Sandt's DRT theory is effectively a *familiarity* theory of anaphora, wherein a definite presupposes an accessible *discourse referent antecedent*.[33] Elbourne says that this anaphoric presupposition has to be accommodated on his account, though he doesn't explain why or how. He only points out that this proposal leaves open the possibility that the existence presupposition can be locally satisfied, e.g. in classical donkey sentences (p.63) and even in (41), so that the value for *the woman* might "co-vary" with instantiations of the domain of men. This is "a pragmatic matter".

Hence, Elbourne ends up with familiarity, much like Heim's 1982 theory, but with too much uniqueness in cases where a pronoun isn't under quantification, and no pragmatic account.

4.3 Failure to capture the scope constraint on anaphora

Most theories of definite descriptions and E-type/D-type pronouns fail to capture a generalization over their distribution which follows straightforwardly from most familiarity-based theories of definites, like those of Kamp and Heim.[34] Roberts (1996) called it the *Scope Constraint*, and I offer it here in modified form:

> **Scope Constraint on Anaphoric Relations**: A discourse referent introduced by an NP *x* is anaphorically accessible to NP *y* just in case any operators that have scope over *x* have scope over *y* as well.

As Roberts discusses in some detail, the Scope Constraint is an empirical generalization, capturing the fact that anaphora isn't possible in cases like (10) and (11), repeated from above, and (43) (contrasted with (44), involving narrow scope of the potential antecedent with respect to the intensional predicate *doubt*:

(10) Every man that owns a zebra feeds it oats.
#I thought it looked miserable yesterday.

(11) Suppose there was a zebra walking down High Street.
#It wasn't very friendly.

[33] OK, it's a DRT *reference marker*, but those are effectively discourse referents, treated via parallel mechanisms to those proposed for discourse referents in Heim (1982). There are differences between van der Sandt's theory and Heim's, but they are not relevant here.

[34] Although in some expositions of DRT the notion of anaphoric accessibility is treated as if configurationally defined over Discourse Representations, Chierchia & Rooth (1984) show that the defined relations between DRs come down to semantic scope relations.

(43) I very much doubt that John has a girlfriend. #She's good looking.
(44) John has a girlfriend. She's good looking.

The constraint follows without stipulation in familiarity-based accounts of anaphora with the usual semantics for operator scope, where the scope of a higher operator limits narrow scope existential entailments, effectively introducing anaphoric islands around any indefinites introduced in the arguments of the operator. Note that anaphoric *in*definites like *one*, which do not require familiarity, are not subject to the Scope Constraint, as we see in (45), in contrast with the definite *it* in (46):[35]

(45) Either John has a new car, or else Mary has one.
(46) Either John has a new car, or else #Mary has it.

Disjunctions are well-known scope islands. Thus, in a theory in which the satisfaction of a familiarity presupposition is the key to felicitous anaphora, and in which that satisfaction, like presupposition satisfaction generally, is subject to the scope constraint, you predict the infelicity of (46).

The Scope Constraint is a problem for Neale's account, as given in his (33) above, since that account merely requires an antecedent quantificational NP from which the definite's descriptive content is to be *borrowed*: There is no anaphoric relation between the two NPs, hence no scope constraint. Of course, Neale could elaborate his theory to try to capture the constraint. But insofar as the *explanation* for the scope constraint lies in the anaphoric relation between the definite and the discourse referent introduced by the preceding NP, then it seems likely that any satisfactory story would end up being anaphoric in a very similar sense to that proposed here and off-handedly adopted by Elbourne.

Though he doesn't seem to notice this,[36] it is precisely because Elbourne adopts an anaphoric distinction between definites and indefinites like that in DRT that he effectively predicts that definites will display the scope constraint on anaphoric accessibility. But at this juncture, we might ask whether, if Elbourne assumes familiarity, he still needs "Fregean" existence plus uniqueness? I would argue that his existence plus uniqueness presupposition isn't necessary if we have (weak) familiarity plus Retrievability, and in fact, the latter does a better job at predicting the attested occurrence of semantic uniqueness effects. Then the same question applies to any modification of Neale's theory which correctly predicts the Scope Constraint: Do we still need Russellian uniqueness?

[35] I suspect that *one*-anaphora is a species of anaphora to kinds, as discussed in Roberts (2005). If this is the case, then the true antecedent of *one* in (45) is the kind *new cars*, which is only brought to salience by the licensing NP *a new car*.

[36] His discussion ignores the scope constraint except to briefly mention (i), his (110), about which he says, "we can leave aside…the problem of how *the fish* in the second sentence of [(i)] comes to talk about the putative fish that John may or may not catch."
(i) John wants to catch a fish. He hopes I will grill it for him.
See Roberts (1989,1996) for extensive discussion of such cases of modal subordination, and how they do, in fact, respect the scope constraint.

4.4 Pronominal ambiguity

E-/D-type theories end up taking pronouns to be ambiguous. For Elbourne, when pronouns are "bound or referential", the phonologically null NP which is their argument is a referential index, whereas in the D-type cases, there is elided descriptive content.[37]

Neale (1990) claims that there are three possible interpretations of pronouns:

i. Pronouns are *referring expressions* when they are:
   a) non-anaphoric (in his sense), i.e. fail to have an NP antecedent, or are
   b) anaphoric on referring expressions (e.g. demonstratives or proper names)
ii. Pronouns are interpreted as *bound variables* when they are anaphoric on c-commanding quantifiers.
iii. Pronouns are interpreted as Russellian ("*D-type*") quantifiers when they are anaphoric on non-c-commanding quantifiers.

The theory of pronominal definites proposed here and in the other papers by Roberts cited posits no ambiguity, and can capture all the kinds of readings discussed by Neale and Elbourne, plus some they don't consider (e.g., the generic interpretations discussed in Roberts 2005).

For these and other reasons, I conclude that an approach to definites which includes semantic uniqueness, whether presupposed or proffered, is:

- Not necessary: The semantics proposed are too rigid with respect to uniqueness, unless (in cases under quantification) uniqueness is so weakened by the use of minimal situations as to amount to no condition at all. In contrast, weak familiarity plus the informational uniqueness that follows for free from Retrievability yields the full range of uses of definites, including the attested range of uniqueness effects, without Russellian uniqueness.
- Not sufficient: Since these are purely semantic, and not anaphoric (informational) theories of uniqueness, they are inadequate to guarantee Retrievability, which is triggered by anaphora. Insofar as Retrievability seems to offer a satisfactory explanation of the distribution of uniqueness effects, E-/D-type theories would have to be enriched by something like anaphoric familiarity (as in Elbourne's theory) to suffice.
- Not extendable: They are too particular to definite descriptions and particular uses of pronouns to be extended to include the full class of definites, including demonstratives, non-nominal presuppositional triggers like *too* and *only*, ellipses, and prosodic deaccentuation. Though discussion of such cases is beyond the scope of the present paper, I claim that Retrievability is independently required for all these phenomena, since they are all in fact anaphoric. Roberts (in preparation a, in preparation b) argues that in each of these cases, Retrievability permits a more perspicuous and empirically satisfactory theory.

Though details of other proposals in the recent literature differ from those in the theories discussed here, those of, e.g., Lepore (2005) and Recanati (2005) both share the assumption that the basic semantics of the definite article involves semantic uniqueness. Both therefore encounter several of the problems noted here, including the problem of too much uniqueness—and more generally, an inability to predict the attested range of uniqueness effects. Neither considers the type of anaphoric account proposed here, and hence both fail to predict the scope constraint.

[37] Elbourne talks at some points as if his is a non-ambiguity theory of pronouns, but at certain junctures tries to escape from problems by claiming that the pronouns in those cases are "referential" (e.g., p.85). Moreover, his D-type account requires an overt antecedent, so there are clearly cases he can't treat as D-type. Taking "referential indices" to be a species of descriptive content does not strike me as intuitive.

5. Conclusions and Prospects: Context, Reference, and Domain Restriction

The discussion of the theory of definite NPs in sections 2 and 3 is abbreviated here, of necessity, and in particular there was not space to offer detailed consideration of the wide range of types of examples of anaphora which I believe support the present proposal. Some of this material is discussed in Roberts (1989,1996b), and the papers in which the present theory was first developed (Roberts 2002,2003,2005); such examples are crucial to appreciating the flexibility and scope of this theory. What I hope to have made plausible here is that there is no need to assume that informational uniqueness is part of the (presupposed) content of definites themselves; instead, it follows from their anaphoric character—the presupposition of weak familiarity—and the general Gricean Retrievability condition on felicitous anaphora. And that the present theory makes correct predictions about felicity and meaning in a wider range of cases than either of the other two theories of definite descriptions and pronouns in discourse with which it was compared.

What is key to understanding the contribution to meaning of an utterance of a definite description or pronoun is the way that its familiarity presupposition is intended by the speaker to be resolved. This doesn't require any particular theory of logical form, nor does it require syntactic ellipsis. The definites themselves (*the*, the personal pronouns) simply trigger the familiarity presupposition. Though this gives the effect of domain restriction, it needn't involve variables over a presupposed domain, as in von Fintel (1994), or global domain restriction, as in Barwise & Perry (1983), nor does it make recourse to situations in interpreting predicates generally, as in Recanati (2005). The anaphora which is central to the account is triggered by the anaphoric definite lexical items—*the*, *he*, *that*—and from that and general pragmatic principles, all the rest follows.

Though the present analysis includes a strong pragmatic component, it does not follow authors who argue that utterances involving incomplete definite descriptions are false, but that one may thereby use them to convey something true (e.g. Reimer 1998, Lepore 2005)—in parallel to the kind of pragmatic re-interpretation Kripke (1977) proposed to yield the so-called *referential* uses of definite descriptions. The definite article is the most frequently used word in the English language.[38] I am willing to bet dinner at The French Laundry restaurant[39] that the majority of the definite descriptions in the large, representative Oxford English Corpus are incomplete in the sense of interest. I myself would be uncomfortable saying that all the utterances containing these descriptions are therefore literally false, and only pragmatically convey the information intended.

On the theory of definites proposed here, no matter how poor the descriptive content of a definite, there is no sense in which, in felicitous use, its description is incomplete. The theory suggests an alternative characterization of the role of descriptive content in definite NPs: It is not about picking out some entity in the world which is unique in virtue of bearing the property corresponding to the descriptive content. Rather, it need only give the addressee just enough information to uniquely pick out the intended antecedent for the definite *among those entities weakly familiar in the discourse*, *given what is Relevant, plausible and salient*. The characterizations of the contextual scoreboard and Retrievability are independently motivated,

[38] From the Wikipedia entry on "Most common words in English": "The list below of most common words in English. . .is based on an analysis of the Oxford English Corpus of over a billion words, and represents one study done by Oxford Online, associated with the Oxford English Dictionary." See the Ask-Oxford.com site for details.

[39] http://www.frenchlaundry.com/. I dare you.

and they suggest in a psychologically plausible and natural way how the search-space for anaphoric resolution for a particular definite is constrained: The addressee can reasonably be assumed to be attending to certain entities, those relevant to the QUD, and it is not semantic uniqueness but informational uniqueness relative to that set (or the entire DR, if applicable) which is required for Retrievability. Thus, uniqueness arises from the general Retrievability condition on anaphora, and not from the semantics of the definite article itself.

Hence, the "problem" of incomplete descriptions is really an epiphenomenon of the basic assumption underlying theories based on the Russellian truth conditions for *the*: that one needs to account for the uniqueness effects sometimes observed with definite descriptions directly in terms of the lexical semantics of the definite article.

There are three kinds of issues raised in the present discussion which call out for further consideration:

The first is how the present notion of context, embodied in the conversational scoreboard presented in section 2, bears on discussions of the role of context in interpretation and of contextualism and relativism in the philosophy of language. When I read that work, I often feel that different authors are talking past each other in adopting rather different, albeit often rather vaguely defined notions of context. Some seem to assume that a context of utterance is the actual, physical circumstance in which the utterance occurs, a situation say; others that it is a set of indices, privileging a few types of information from such a situation and essential for indexicality in a fairly narrow sense; others that it is something like the common ground, shared propositional information crucial for the consideration of how presupposition works. Though this makes it tricky to compare the various proposals, I do not want to suggest that there is any one notion of context which is *the* notion relevant for natural language semantics and pragmatics. But I do think that a more nuanced, explicit characterization of the kinds of information that interlocutors evidently track in discourse and of the relationships between these different kinds of information is very useful for developing detailed analyses of particular expressions in particular discourse contexts. And that an approach to modeling discourse context based on the intentional structure of agents' interactions can shed light in detailed and *predictive* ways on the interplay between conventional semantic content and pragmatic principles in yielding the attested interpretations of particular expressions of a language. The present essay is intended as an illustration of that claim. I think that only when we have more analyses at this level of detail can we really take stock of the role of context in interpretation, and hence of contextualism.

The second issue which bears further investigation in light of the preceding discussion is the nature of linguistic reference. I would agree with those who argue that it is not linguistic expressions like noun phrases which refer, but speakers—those who use such constituents—as a function of those speakers' referential intentions. Questions arise: What do speakers intend to refer to? and How do linguistic expressions lend themselves to this purpose? In most discussions of reference in the philosophical literature, the interpretation of definite noun phrases is central; e.g., in a recent overview, Reimer (2003) focuses almost entirely on proper names, indexicals (demonstratives among them), and definite descriptions, all of these definite by linguists' criteria (Abbott 2004). For at least three of these types of expressions—demonstratives, definite descriptions and pronouns— I have argued that insofar as they *are* used by speakers to refer to entities in the world, there is reason to think that they do so indirectly, via presupposing familiar discourse referents in the interlocutors' common information. See also Geurts (1997) and Cumming (2008) for relevant recent work on proper names. This suggests that while speakers may on occasion *intend* to refer to entities in the world via the use of a definite, this may be accomplished only indirectly, since the noun phrases they use to do so are presuppositionally

anchored not directly to entities in the world, but to information *about such entities*, i.e. to the discourse referents shared with other interlocutors in the conversation. These questions, of course, go well beyond the present discussion.

Finally, and most specifically, the present discussion of incomplete descriptions raises the more general issue of the way in which the domains of various operators are restricted implicitly in discourse. This has been a hot topic in the past few years, discussed by von Fintel (1994), Roberts (1995), Gawron (1996), Recanati (1996), Bach (1999), Stanley & Szabo (2000), and many others besides. But that, too, will have to wait for another occasion.

Bibliography

Bibliography:

Abbott, Barbara (2004) Definiteness and indefiniteness. In Laurence R. Horn & Gregory Ward (eds.) *Handbook of Pragmatics*. Oxford: Blackwell, 122-149.

Asher, Nicholas, and Lascarides, Alex (1994) Intentions and information in discourse. Paper presented at *32nd Annual Meeting of the Association of Computational Linguistics*, Las Cruces, NM.

Asher, Nicholas, and Lascarides, Alex (1998) Questions in dialogue. *Linguistics and Philosophy* 21:237-309.

Bach, Kent (1999) The semantics pragmatics distinction: What it is and why it matters. In K. Turner (ed.) *The semantics-pragmatics interface from different points of view*. Elsevier, Oxford, 64-84.

Barwise, Jon & John Perry (1983) *Situations and Attitudes*. MIT Press, Cambridge, MA.

Beaver, David & Henk Zeevat (2007). Accommodation. In G. Ramchand and C. Reiss (eds.), *The Oxford Handbook of Linguistic Interfaces*, 533–538. Oxford University Press.

Berman, Stephen (1987) Situation-based semantics for adverbs of quantification. In James Blevins & Anne Vainikka (eds.) *Studies in Semantics*:46-68. UMass Occasional Papers in Linguistics, Volume 12. GLSA, University of Massachusetts at Amherst.

Bloom, Paul (2000) *How Children Learn the Meanings of Words*, MIT Press.

Bratman, Michael E. (1987) *Intention, Plans, and Practical Reason.* Harvard U. Press.

Carlson, Lauri (1983) *Dialogue Games: An Approach to Discourse Analysis*. Reidel, Dordrecht.

Chambers, C.G., Tanenhaus, M.K, Eberhard, K.M., Filip, H. & Carlson, G.N. (2002). Circumscribing referential domains in real-time sentence comprehension. *Journal of Memory and Language* 47:30-49.

Chierchia, Gennaro (1995) *Dynamics of Meaning: Anaphora, presupposition, and the meaning of grammar*. University of Chicago Press.

Chierchia, Gennaro & Mats Rooth (1984) Configurational notions in Discourse Representation Theory. In Charles Jones & Peter Sells (eds.) *Proceedings of NELS 14*, UMass/Amherst.

Cooper, Robin (1979) The interpretation of pronouns. In Frank Heny and Helmut S. Schnelle (eds.), *Syntax and Semantics10: Selections from the Third Groningen Round Table*. Academic Press, New York, 61-92.

Cumming, Samuel (2008) Variabilism. *Philosophical Review* 117.4:525-554

Elbourne, Paul D. (2001) E-type anaphora as NP-deletion. *Natural Language Semantics* 9:241-288.

Elbourne, Paul D. (2005) *Situations and Individuals*. MIT Press.

Engdahl, Elisabet (1986) *Constituent Questions*. Reidel, Dordrecht.

Evans, Gareth (1977) Pronouns, quantifiers and relative clauses (I). *Canadian Journal of Philosophy* 7:467-536. Reprinted in Mark Platts (ed.), *Reference, Truth and Reality: Essays on the Philosophy of Language.* Routledge and Kegan Paul, 255-317.

Evans, Gareth (1980) Pronouns. *Linguistic Inquiry* 11.2:337-362.
von Fintel, Kai (1994) *Restrictions on Quantifier Domains*. Ph.D. dissertation. University of Massachusetts at Amherst.
von Fintel, Kai & Anthony S. Gillies (2008) CIA Leaks. *The Philosophical Review* 25(1-2):139-201.
Franke, Michael, Tikitu de Jager & Robert van Rooij (to appear) Relevance in cooperation and conflict. *Journal of Logic and Computation*.
Frege, Gottlob (1891) Function and concept. Translated in M. Black & P. Geach (1960) *Translations from the Philosophical Writings of Gottlob Frege*. Blackwell, Oxford:56-78.
Gawron, Jean Marc (1996) Quantification, quantificational domains, and dynamic logic. In Shalom Lappin (ed.) *The Handbook of Contemporary Semantic Theory*. Blackwell, Cambridge, MA.
Geurts, Bart (1997) Good news about the description theory of names. *Journal of Semantics* 14:319-348.
Grice, H.P. (1957) Meaning. *The Philosophical Review* 66. Reprinted in Grice 1989.
Grice, H. P. (1969) Utterer's Meaning and Intentions. *The Philosophical Review* 78. Reprinted in Grice 1989.
Grice, H.P. (1989) *Studies in the Way of Words*. Harvard University Press.
Groenendijk, Jeroen & Martin Stokhof (1984) *Studies on the Semantics of Questions and the Pragmatics of Answers*. Ph.D. dissertation, University of Amsterdam
Groenendijk, Jeroen & Martin Stokhof (1997) Questions. In Johan van Benthem & Alice ter Meulen (eds.) *Handbook of Logic and Language*, Elsevier Science/MIT Press, Chapter 19, 1055-1124.
Grosz, Barbara J. (1977) *The Representation and Use of Focus in Dialogue Understanding*. Ph.D. Dissertation, University of California, Berkeley.
Grosz, Barbara, & Candice Sidner (1986) Attention, intentions, and the structure of discourse. *Computational Linguistics* 12:175-204.
Hanna, J.E. & Tanenhaus, M.K. (2004). Pragmatic effects on reference resolution in a collaborative task: evidence from eye movements. *Cognitive Science* 28:105-115.
Heim, Irene (1982) *The Semantics of Definite and Indefinite Noun Phrases*. Ph.D. dissertation, University of Massachusetts, Amherst.
Heim, Irene (1990) E-Type Pronouns and Donkey Anaphora. *Linguistics and Philosophy* 13:137-177.
Heim, Irene & Angelika Kratzer (1998) *Semantics in Generative Grammar*. Blackwell, Oxford.
Hobbs, Jerry R., Stickel, Mark E., Appelt, Douglas E., and Martin, Paul (1993) Interpretation as abduction. *Artificial Intelligence* 63:69-142.
Kadmon, Nirit (1987) *On Unique and Non-Unique Reference and Asymmetric Quantification*. Ph.D. dissertation. UMass, Amherst.
Kadmon, Nirit (1990) Uniqueness. *Linguistics and Philosophy* 13:273-324.
Kamp, Hans (1981) A Theory of Truth and Semantic Representation. In Jeroen Groenendijk, Theo M. V. Janssen and Martin Stokhof (eds.) *Formal Methods in the Study of Language, Vol. I*, Mathematische Centrum, Amsterdam. Reprinted in Groenendijk, Janssen and Stokhof (eds.) (1984) *Truth, Interpretation and Information*, Foris, Dordrecht.
Karttunen, Lauri (1976) Discourse referents. In James D. McCawley (ed.) *Syntax and Semantics 7: Notes from the Linguistic Underground*. Academic Press, New York.
Kehler, Andrew (2002) *Coherence, Reference, and the Theory of Grammar*. CSLI Press, Stanford, CA.
Kratzer, Angelika (1981) The notional category of modality. In H.J. Eikmeyer & H. Rieser (eds.) *Words, Worlds and Contexts: New approaches to word semantics*. de Gruyter, Berlin. Reprinted in P. Portner & B.H. Partee (eds.) (2002) *Formal Semantics: The essential readings*.

Kratzer, Angelika (1989) An investigation of the lumps of thought. *Linguistics and Philosophy* 12:607-653.
Kratzer, Angelika (2007/rev.2009) Situations in Natural Language Semantics. *Stanford Encyclopedia of Philosophy*. http://plato.stanford.edu/entries/situations-semantics/
Kripke, Saul A. (1977) Speaker's reference and semantic reference. In P.A. French, T.E. Uehling & H.K. Wetttein (eds.) *Contemporary Perspectives in the Philosophy of Language*. University of Minnesota Press, Minneapolis.
Kripke, Saul A. (2009) Presupposition and anaphora: Remarks on the formulation of the projection problem. *Linguistic Inquiry* 40.3:367–386.
Lappin, Shalom, and Nissim Francez (1994) E-type Pronouns, I-sums, and Donkey Anaphora. *Linguistics and Philosophy* 17:391-428.
Lepore, Ernie (2005) An abuse of context in semantics: The case of incomplete definite descriptions. In M. Reimer & A. Bezuidenhout (eds.) *Descriptions and Beyond*. Oxford University Press, 41-67.
Lewis, David (1979) Score-keeping in a language game. In Rainer Bauerle, Urs Egli & Arnim von Stechow (eds.) *Semantics from a Different Point of View*. Berlin: Springer.
Merchant, Jason (2005) Revisiting syntactic identity conditions. Handout from a workshop ellipsis, UC Berkeley, available on Merchant's website.
Montague, Richard (1973) The proper treatment of quantification in ordinary English. In Jaakko Hintikka, Julius Moravcsik & Patrick Suppes (eds.) *Approaches to Natural Langauge: Proceedings of the 1970 Stanford Workshop on Grammar and Semantics*:221-242. Reidel, Dordrecht.
Muskens, Reinhart (1996) Combining Montague Semantics and Discourse Representation. *Linguistics and Philosophy* 19:143-186.
Neale, Stephen (1990) *Descriptions*. MIT Press (Bradford Books), Cambridge, MA.
Neale, Stephen (2005) This, that and the other. In M. Reimer & A. Bezuidenhout (eds.) *Descriptions and Beyond*. Oxford University Press, 68-182.
Parsons, Terence (1978) Pronouns as Paraphrases. Ms., University of Massachusetts, Amherst.
Perrault, C. Raymond (1990) An application of default logic to speech act theory. In P. Cohen, J. Morgan, and M. Pollack (eds.) *Intentions in Communication*. MIT Press, Cambridge, MA.
Pelletier, J., and L. Schubert (1989) Generically speaking. In G. Chierchia, B. H. Partee and R. Turner (eds.) *Properties, Types and Meaning*, Volume 2, Dordrecht: Kluwer.
Portner, Paul (2005) The Semantics of Imperatives within a Theory of Clause Types. In Kazuha Watanabe and Robert B. Young, eds., *Proceedings of Semantics and Linguistic Theory 14*. Ithaca, NY: CLC Publications. Available at semanticsarchive.net.
Portner, Paul (2007) Imperatives and modals. *Natural Language Semantics* 15.4:351-383.
Quine, W. V. (1960) *Word and Object*. Cambridge, MA: MIT Press.
Recanati, François (1996) Domains of discourse. *Linguistics and Philosophy* 19:445-75.
Recanati, François (2005) Descriptions and situations. In M. Reimer & A. Bezuidenhout (eds.) *Descriptions and Beyond*. Oxford University Press, 15-40.
Reimer, Marga (1998) Quantification and context. *Linguistics and Philosophy* 21:95-115.
Reimer, Marga (2003) Reference. *Stanford Encyclopedia of Philosophy*. http://plato.stanford.edu/entries/reference/
Roberts, Craige (1989) Modal Subordination and Pronominal Anaphora in Discourse. *Linguistics and Philosophy* 12.6:683-721. Reprinted in Javier Gutierrez-Rexach (ed.) *Semantics: Critical concepts in linguistics*, Routledge, 2003.
Roberts, Craige (1995) Domain Selection in Dynamic Semantics. In Emmon Bach, Eloise Jelinek, Angelika Kratzer, and Barbara H. Partee (eds.) *Quantification in Natural Languages*, Kluwer, Dordrecht.

Roberts, Craige (1996) Information Structure in Discourse: Towards an Integrated Formal Theory of Pragmatics. In Jae Hak Yoon and Andreas Kathol (eds.) *Ohio State University Working Papers in Linguistics* Volume 49. To appear in *Semantics and Pragmatics*. The most commonly circulated version was revised in 1998. The paper is to appear in *Semantics and Pragmatics*.

Roberts, Craige (1996b) Anaphora in intensional contexts. In Shalom Lappin (ed.) *The Handbook of Contemporary Semantic Theory*, Basil Blackwell, 1996

Roberts, Craige (2002) Demonstratives as definites. In K. van Deemter & R. Kibble (eds.) *Information Sharing: Reference and Presupposition in Language Generation and Interpretation*. CSLI Press, Stanford, 89-196, 2002.

Roberts, Craige (2003) Uniqueness in definite noun phrases. Linguistics and Philosophy 26:287-350.

Roberts, Craige (2004) Context in dynamic interpretation. In Laurence Horn and Gregory Ward (eds.) *Handbook of Contemporary Pragmatic Theory*, Blackwell.

Roberts, Craige (2005) Pronouns as definites. In M. Reimer & A. Bezuidenhout (eds.) Descriptions and Beyond. Oxford University Press.

Roberts, Craige (in preparation a) Ellipsis and Retrievability.

Roberts, Craige (in preparation b) Resolving Focus.

Rooth, Mats (1987) Noun phrase interpretation in Montague Grammar, File Change Semantics, and Situation Semantics. In Peter Gärdenfors (eds.) *Generalized Quantifiers*. Dordrecht: Reidel, 237-268.

Schwarzschild, Roger (1999) GIVENness, AVOIDF and Other Constraints on the Placement of Accent. *Natural Language Semantics* 7(2):141-177.

Simons, Daniel & Christopher Chabris (1999) Gorillas in our midst: Sustained inattentional blindness for dynamic events. *Perception* 28(9):1059-1074.

Soames, Scott (1986) Lost innocence. *Linguistics and Philosophy* 8:59-71.

Sperber, Dan, & Deirdre Wilson (1986) *Relevance: Communication and cognition*. Harvard University Press, Cambridge, MA.

Stalnaker, Robert (1978) Assertion. In Peter Cole (ed.) *Syntax and Semantics, Volume 9: Pragmatics*, 315-32. Academic Press, NY.

Stanley, Jason, & Zoltan Gendler Szabo (2000) On quantifier domain restriction. *Mind & Language* 15.2-3:219-261.

Strawson, P.F. (1950) On referring. *Mind* 59:320-344.

Tanenhaus, M.K., Spivey-Knowlton, M.J., Eberhard, K.M. & Sedivy, J.E. (1995) Integration of visual and linguistic information in spoken language comprehension. *Science* 268:1632-1634.

Terken, Jacques, & Julia Hirschberg (1994) Deaccentuation of words representing "given" information: Effects of persistence of grammatical function and surface position. *Language and Speech* 37(2):125-145.

Thomason, Richmond (1990) Accommodation, meaning, and implicature: Interdisciplinary foundations for pragmatics. In P. Cohen, J. Morgan, and M. Pollack (eds.) *Intentions in Communication*. MIT Press, Cambridge, MA.

van der Does, Jaap (1993) The Dynamics of Sophisticated Laziness. In J. Groenendijk (ed.) *Dyana Deliverable* R22A. Amsterdam: ILLC.

Winograd, Terry (1972) *Understanding Natural Language*. Academic Press, New York