Revealing Concealment
A (Neuro-)Logical Investigation of Concealed Questions

MSc Thesis (*Afstudeerscriptie*)

written by

Jesse Aron Harris
(born 20 February 1981 in Seattle, WA, USA)

under the supervision of **Paul Dekker, Liina Pylkkänen,** and **Martin Stokhof**, and submitted to the Board of Examiners in partial fulfillment of the requirements for the degree of

MSc in Logic

at the *Universiteit van Amsterdam*.

Date of the public defense:	**Members of the Thesis Committee:**
30 August 2007	Maria Aloni
	Paul Dekker
	Peter van Emde Boas
	Liina Pylkkänen
	Martin Stokhof

Institute for Logic, Language and Computation

The Chair

A funny thing about a Chair:
You hardly ever think its *there*.
To know a Chair is really it,
You sometimes have to go and sit.
Theodore Roethke

Acknowledgements

Since this thesis took about twice as long to complete as it should have, I have twice as many people to thank had it been completed on time. Like the Roethke poem quoted above, I did not know how much I would learn from, and struggle with, a thesis of this size and scope until I sat down and wrote it. And I could not have done it without the support of the people I thank below.

First and foremost, I must thank my advisers for their guidance. Having a chance to work with Paul and Martin was the primary motivation for returning to the ILLC and I was not disappointed. I thank them for taking on this project from across the Atlantic. Liina got me started thinking about concealed questions by offering me a chance to give a guest lecture on the topic in her graduate semantics course at New York University in the spring semester of 2005. The experimental part of this thesis simply would not have been possible without her support, intellectually and financially.

Benedikt Löwe helped bring me back to ILLC and thereafter acted as my unofficial academic adviser. Tanja Kassenaar was wonderfully helpful in navigating the administrative mine-field that is UvA. Thanks to you both! Chapter 4 is extracted and modified from an article co-authored with Liina, Brian McElree and Steven Frisson. I thank all three for their contributions to the original article and especially to Steven for running the eye-movements experiment and analyzing the data. The experimental work was carried out during my two-year tenure as Lab Manager for the NYU Neurolinguistics Lab. During that time, Dr. Rudolfo Llinás generously made available his MEG lab. Kim Moran and Puma Jagow deserve special thanks for patiently guiding me through and fixing some persistent technical difficulties that occurred when trying to run subjects with the MEG instrument.

A great many people have had the good patience to talk with me about concealed questions, despite their busy schedules. In addition to my advisors, Paul Elbourne, Lyn Frazier, Kyle Johnson, Anastasia Giannakidou, Angelika Kratzer, Barbara Partee, Chris Potts, Craige Roberts, and Maribel Romero have all given me excellent feedback and encouragement. Special thanks goes to Kyle Johnson and Chris Potts who allowed me to incorporate this work into my assignments and to Maribel Romero whose critical assessment of a previous instantiation of the work convinced me that I needed to rethink just about all of the thesis. I also thank the members of my defense committee, which included my advisers and Maria Aloni and Peter van Emde Boas, for their comments. Maria was particularly generous with her time after the defense in working out a few technical problems with the defense draft.

I have been in three different departments while working toward this degree. I'm happy to say that I've met and was influenced by more people than I have space to thank. At NYU, I have to thank Eytan Zweig for showing me the ropes on the MEG and the other lab members for enduring tricky and subtle judgments on concealed questions. Fellow lab member Andrea Martin made hot New York summers a little more bearable with our weekly kvetching sessions at Think in Washington Square. At the University of Massachusetts-Amherst, many thanks go out to my movie-watching buddies – especially Maria Biezma-Garrido and Youri Zabbal who gave me a place to stay when I was finishing the defense draft – and my classmates for their fine pastry creations during our weekly Sunday synsem homework sessions. At the ILLC, I had the pleasure of staying with some of the best roommates I've yet had. Thanks to everyone at Biesboschstraat 18/3 for making me feel at home, especially Piwi, whose orange fuzziness cheered me up when I need it most.

Lastly, I thank my family in all its extensions and permutations. Mom and Paul, Dad and Janice, Ben and Jessica, and Margery: thanks for helping me see this project through!

Contents

1
Introduction

Natural language is notoriously multi-faceted and thoroughly nefarious. Even the simplest linguistic elements may turn out to support different interpretations in different contexts. And the more one examines natural language, the more disparate and complex the relationship between a word's meaning and the context in which it occurs seems to become.

This thesis is about contexts in which nominal elements are not interpreted as nominal elements typically are. To be specific, this thesis is about so-called *concealed questions*, nominal phrases (*e.g., the price of milk*) which may be paraphrased as indirect questions (*e.g., what the price of milk is*). This investigation stands at the crossroads of several lines of research, which are themselves perhaps only tenuously connected. The basic outline of the thesis can be divided into three major types of inquiry: (i) descriptive, (ii) theoretical, and (iii) experimental. It is hoped that the sorts of research discussed here complement one another, giving a richer and deeper understanding of a highly complex and ill-understood phenomenon. I briefly present the basic landscape of the thesis, highlighting the important results.

1.1 Overview of Thesis

Chapter 2 is primarily descriptive. The core, established properties of concealed questions (CQs) are critically reviewed, and evidence for two new, but revealing, properties is provided. Then, three major accounts are presented: Irene Heim's early, influential account; a string of recent articles by Maribel Romero arguing that CQs are individual concepts, and the recent dissertation of Lance Nathan, exploiting a typological correspondence between CQs and propositions.

Chapter 3 is primarily theoretical. A novel semantic account of concealed questions is developed, one that employs the pragmatic notion of conceptual covers (Aloni, 2001). Here, I argue that CQs are best expressed as relations between individual concepts and identity questions, such that the concept provides a constituent answer to the question itself. In particular, I posit a type-shifting rule $\mathscr{Q}$ which lifts the nominal to a more complex type $\langle\langle s,e\rangle,\langle s,\langle s,t\rangle\rangle\rangle$. Two different ways of viewing this relation are presented, and the cognitive naturalness of $\mathscr{Q}$ is briefly discussed and defended.

The formal framework is itself developed within a fully compositional Dynamic Montague Grammar (Groenendijk and Stokhof, 1990a) with Existential Disclosure (Dekker, 1993b). The unique binding properties of Dynamic Semantics is exploited to treat a long standing puzzle regarding an ambiguous interpretation of CQs (Heim's ambiguity). Together, the chapter gives an in depth framework for addressing some of the context sensitivity that CQs display. Throughout, the chapter discusses how CQs relate to propositions, free relatives, and identity questions.

Chapter 4 presents the first psycholinguistic experiment on concealed questions. Two separate methodologies, eyetracking and magnetoencephalography (MEG), found that that interpreting CQs taxes the language processor. An interpretation of the results which is broadly compatible with a type-shifting analysis like the one presented in the preceding chapter is subsequently argued for in detail.

1.2 Formal and Conceptual Prerequisites

For ease of reference, the most general requisite background is provided in the sections that follow. Three topics are introduced, and the formalisms briefly sketched. First, relational nouns and an important subdistinction thereof are identified. Second, a basic intensional language is introduced. Third and finally, a brief and incomplete sketch of the semantics of questions follows, focusing primarily on the partition semantics developed by Groenendijk and Stokhof (1984, 1990b, 1997).

1.2.1 Relational Nouns

As much of the work of the thesis hinges on an understanding of relational nouns, a brief overview is provided for reference here. In general, relational nouns, such as those listed in (1), are nouns that inherently denote relations between the head of the noun (*mayor*) and an extrinsic argument (*town*). Note that though the extrinsic argument may or may not be expressed overtly (1–2), genuine relational nouns are always to be interpreted relative to another argument. All this is to say that one cannot be a mayor, unless one is a mayor of something.

(1) Relational nouns without an overt argument

a. mayor
b. brother
c. price
d. cost

(2) Relational nouns with an overt argument

a. mayor of the town
b. brother of Aaron
c. price of milk
d. cost of the reform

In an account which exploits thematic relations familiar from the verbal domain, Barker and Dowty (1993) define relational nouns thusly:

> In general, a relational noun is one such that an entity qualifies for membership in the extension of the noun only by virtue of there being a specific second entity which stands in a particular relation to the first, and where that relation is determined solely by the noun's lexical meaning.

In addition, de Bruyn and Scha (1988) observe that overt expression of these extrinsic arguments is typically optional, and that 'unsaturated' relational nouns are syntactically very similar to non-relational nouns, *e.g., stick*. Given this resemblance to ordinary nouns, we will treat the exceptionality of relational nouns in the semantics. Following Dekker's (1993) dynamic framework, we will assume that a distinguished prepositional phrase of_2 introduces the extrinsic argument into the relation denoted by the relational noun head. As the account is presented in §3.4.1, we withhold the formal details until then.

An important, if ill-understood, distinction holds between (1/2a–b) and (1/2c–d). The latter have been termed *abstract* relational nouns, as they denote abstract properties holding of the extrinsic argument (see Barker and Dowty (1993) and Nathan (2006) for discussion). In contrast, the former are known as *concrete* relational nouns, and are thought to represent public positions denoting persons (Janssen, 1984). Both types of relational nouns are context sensitive, and we treat them uniformly as individual concepts, functions from indices to individuals. However, we encounter evidence that abstract and concrete relational nouns have different distributions and interpretations in intensional environments in §3.6.2.

1.2.2 Intensional Logic

In nearly all of what follows, we use a two-sorted type theory that distinguishes between basic kinds of types. As in a purely extensional type theory, we allow e, the type of individuals, and t, the type of truth values, and any combinations thereof as basic types. We further assume a type s corresponding to objects from W, the set of *possible worlds*. Let D_a and D_b be the set of objects of type a. Like extensional type theory, two-sorted type theory $T2$ forms composite types as functions from D_a to D_b, for any well-formed types a and b in the type logic. The syntax is given more formally in Definition 1.2.1, after Gamut (1991, §5.8):

Definition 1.2.1. [IL Type Theory] Let $T2$, the set of types in intensional type theory, be the smallest set such that:

1. $s, e, t \in T2$
2. If $a, b \in T2$, then $\langle a, b\rangle \in T2$

Given a type $a \in T2$, an object of type $\langle s, a\rangle$ is interpreted as a function from possible worlds W, maximal and complete ways in which the world might be, to objects of type a. Accordingly, we will speak of an object of type $\langle s, a\rangle$ as the *intension* of a. General domains of interpretation are defined as follows:

Definition 1.2.2. [Domains of interpretation] For $a \in T2$, a domain of individuals D, and a set of worlds W, we interpret a as a member of the domain $\mathbf{D}_{a,D,W}$ such that

1. $\mathbf{D}_{e,D,W} = D$
2. $\mathbf{D}_{t,D,W} = \{0, 1\}$
3. $\mathbf{D}_{\langle b,c\rangle,D,W} = \mathbf{D}_{b,D,W} \mapsto \mathbf{D}_{c,D,W}$
4. $\mathbf{D}_{\langle s,b\rangle,D,W} = W \mapsto \mathbf{D}_{b,D,W}$

Distinguished among intensional objects, for our purposes, are (i) *individual concepts*, functions from worlds s to individuals, (ii) *propositions*, functions from worlds s to truth values t, (iii) properties, functions from worlds s to sets of individuals $\langle e, t\rangle$, and (iv) *propositional concepts*, functions from worlds s to functions from worlds s' to truth values. Ignoring the subscripts D and W, the types and domains of the objects and their parts are schematized in Table 1.1.

	Name	Type	Domain of Interpretation
Extensional	individual	e	$\mathbf{D}_e$
	truth value	t	$\mathbf{D}_t$
	predicate	$\langle e,t\rangle$	$\mathbf{D}_e \mapsto \mathbf{D}_t$
Intensional	individual concept	$\langle s,e\rangle$	$W \mapsto \mathbf{D}_e$
	proposition	$\langle s,t\rangle$	$W \mapsto \mathbf{D}_t$
	property	$\langle s,\langle e,t\rangle\rangle$	$W \mapsto \mathbf{D}_{\langle e,t\rangle}$
	propositional concept	$\langle s,\langle s,t\rangle\rangle$	$W \mapsto W \mapsto \mathbf{D}_t$

Table 1.1: Important Types in IL

A complete syntax of IL is presented below Gamut (1991). Note that we list only negation, $\neg$, and conjunction, $\wedge$, in our construction rules, other standard connectives being derivable from these two:

Definition 1.2.3. [Syntax of IL] For the set of well-formed expressions WE_α of IL for $a \in T2$, the set WE is the smallest set formed from the following construction rules:

1. If $\alpha \in VAR_a$ or $\alpha \in CON_a$, then $\alpha \in WE_a$
2. If $\alpha \in WE_{\langle a,b\rangle}$ and $\beta \in WE_a$, then $(\alpha(\beta)) \in WE_b$
3. If $\phi, \psi \in WE_t$, then $\neg\phi, \phi \wedge \psi \in WE_t$
4. If $\phi \in WE_t$ and $v \in VAR_a$, then $\forall v\phi, \exists v\phi \in WE_t$
5. If $\alpha \in WE_a$ and $v \in VAR_b$, then $\lambda v.\alpha \in WE_{\langle b,a\rangle}$
6. If $\alpha, \beta \in WE_a$, then $(\alpha = \beta) \in WE_t$
7. If $\phi \in WE_t$, then $\Box\phi, \Diamond\phi \in WE_t$
8. If $\alpha \in WE_a$, then ${}^\wedge\alpha \in WE_{\langle s,a\rangle}$
9. If $\alpha \in WE_{\langle s,a\rangle}$, then ${}^\vee\alpha \in WE_a$

We let a model M be a tuple $\langle W, D, I\rangle$, where W and D are defined as above, and I is the interpretation function, mapping expressions $\alpha \in WE_a$ to objects in $\mathbf{a} \in \mathbf{D}_{a,D,W}$. We let g be an assignment function mapping variables of type a, VAR_a, to domains of interpretation,*i.e.,* $g : VAR_a \mapsto \mathbf{D}_{a,D,W}$. For a model M, a world of evaluation w, and an assignment function g the reference of $\alpha \in WE$ in w with respect to M and g is defined as $[[\alpha]]_{M,w,g}$, as usual. The complete semantics for IL is given below, following Gamut (1991):

Definition 1.2.4. [Semantics for IL]

1. If $\alpha \in CON_a$, then $[[\alpha]]_{M,w,g} = I(\alpha)(w)$
2. If $\alpha \in VAR_a$, then $[[\alpha]]_{M,w,g} = g(\alpha)$
3. If $\alpha \in WE_{\langle a,b\rangle}$ and $\beta \in WE_a$, then $[[\alpha(\beta)]]_{M,w,g} = [[\alpha]]_{M,w,g}([[\beta]]_{M,w,g})$
4. If $\phi, \psi \in WE_t$, then
   (a) $[[\neg\phi]]_{M,w,g} = 1$ iff $[[\phi]]_{M,w,g} = 0$
   (b) $[[\phi \wedge \psi]]_{M,w,g} = 1$ iff $[[\phi]]_{M,w,g} = [[\psi]]_{M,w,g} = 1$
5. if $\phi \in WE_t$ and $v \in VAR_a$, then
   (a) $[[\forall v\phi]]_{M,w,g} = 1$ iff for all $d \in D_a :$ $[[\phi]]_{Mw,g[v/d]} = 1$
   (b) $[[\exists v\phi]]_{M,w,g} = 1$ iff for some $d \in D_a :$ $[[\phi]]_{Mw,g[v/d]} = 1$
6. If $\alpha, \beta \in WE_a$, then $[[\alpha = \beta]]_{M,w,g} = 1$ iff $[[\alpha]]_{M,w,g} = [[\beta]]_{M,w,g}$
7. If $\alpha \in WE_a$ and $v \in VAR_b$, then $[[\lambda v\alpha]]_{M,w,g}$ is that function $h \in {D_a}^{D_b}$ such that for all $d \in D_b : h(d) = [[\alpha]]_{M,w,g[v/d]}$
8. If $\phi \in WE_t$, then
   (a) $[[\Box\phi]]_{M,w,g}$ =1 iff for all $w' \in W :$ $[[\phi]]_{M,w',g} = 1$
   (b) $[[\Diamond\phi]]_{M,w,g}$ =1 iff for some $w' \in W :$ $[[\phi]]_{M,w',g} = 1$
9. If $\alpha \in WE_a$, then $[[{}^\wedge\alpha]]_{M,w,g}$ is that function $h \in W \mapsto \mathbf{D}_a$ such that for all $w' \in W$: $h(w') = [[\alpha]]_{M,w',g}$
10. If $\alpha \in WE_{\langle s,a\rangle}$, then $[[{}^\vee\alpha]]_{M,w,g} = [[\alpha]]_{M,w,g}$(w)

Clauses 2 states that the *intension* of α is to be interpreted as a function from worlds to extensions of α at those worlds. Clause 3 simply states that the *extension* of $\alpha \in WE_{\langle s,a\rangle}$ at a world w is evaluated at that particular w. We will be very sloppy in the terminology below – we use ${}^\wedge\alpha$ and $\lambda w.\alpha(w)$ interchangeably, define the types for variables on the fly, and sometimes use worlds as a vast concept which covers indices in general, be it worlds, times, or information states. It is hoped that the intended meanings of all these terms will be easily gleaned from the text.

1.2.3 Semantics for Questions

An important part of this thesis involves determining the extent to which concealed questions are similar to stand-alone interrogatives. For completeness, we present a basic sketch of the kind of semantics for question that is adopted below. Far more complete overviews can be found in Dekker et al. (2005), Ginzburg (1996), and Groenendijk and Stokhof (1997).

Groenendijk and Stokhof's partition semantics of questions, along with several other theories of questions, follows the three Hamblin postulates listed below (see Hamblin (1958) and Groenendijk and Stokhof (1997) for formulation):

(3) Hamblin Postulates

A. An answer to a question is a sentence or statement

B. The possible answers to a question form an exhaustive set of mutually exclusive possibilities

C. To know the meaning of a question is to know what counts as an answer to that question

For Hamblin (1973), questions are simply sets of propositions. An answer to a question is then one (or more) of the propositions contained in the set of answers. In a closely related approach, Karttunen (1977) restricts question extensions to the set of true answers. Groenendijk and Stokhof also analyze questions in terms of their possible answers, but depart from the Hamblin/Karttunen approach in that interrogatives are given an *intensional* interpretation. In the words of Aloni (2001, 3) "the denotation of an interrogative in a given world is the proposition which expresses the complete true answer to the question in that world."

Definition 1.2.5. [Interrogatives]. For a (possible empty) sequence $\alpha_1, \ldots \alpha_n$ denoted by $\overrightarrow{\alpha}$ and a query operator ?, an interrogative $[[?\overrightarrow{x}\phi]]_{M,w,g}$ denotes the set of worlds v in which the individuals $\overrightarrow{x}$ satisfying ϕ are the same as in the world of evaluation w.

$$[[?\overrightarrow{x}\phi]]_{M,w,g} = \{v \in W \mid \forall \overrightarrow{d} \in D^n : \ [[\phi]]_{M,v,g[\overrightarrow{x}/\overrightarrow{d}]} = [[\phi]]_{M,w,g[\overrightarrow{x}/\overrightarrow{d}]}\}$$

For example, if the sequence $\overrightarrow{x}$ is empty, then the interrogative in question expresses a polar question, as in *Is it now raining?* Given that we have only two truth values, worlds can be only be differentiated in the interrogative as to whether ϕ is true in them or not. That is, the interrogative partitions logical space – the set of possible worlds – into two coarsely grained blocks of worlds:

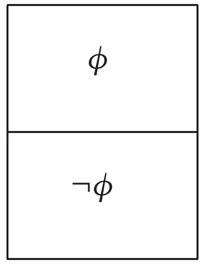

Figure 1.1: Polar questions

On the other hand, if $?\overrightarrow{x}\phi$ is a constituent question, as in *Who did Mary kiss?*, then the interrogative divides logical space into as many blocks of the partition as there are possible denotations of ϕ; it partitions logical space according to worlds which differ on the possible answers to the question posed by the interrogative.

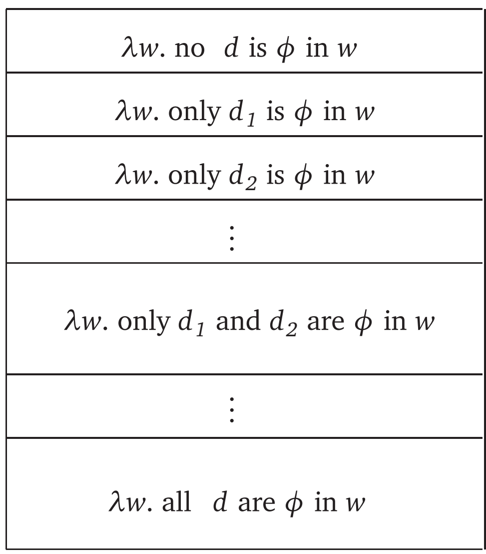

Figure 1.2: Constituent questions

By Definition 1.2.6 from Groenendijk and Stokhof (1990b, 15), a partition $\mathscr{A}$ is an exhaustive division of a set A such that each block X in a partition (i) is non-empty and (ii) does not overlap with any other block in the partition.

Definition 1.2.6. [Partitions] $\mathscr{A}$ is a partition of a set A iff

1. For all $X \in \mathscr{A}$ it holds that $X \neq \emptyset$
2. For all $X, Y \in \mathscr{A}$ it holds that if $X \neq Y$, then $X \cap Y = \emptyset$
3. $\cup\mathscr{A} = A$

The blocks in a partition are thus equivalence classes; for a block X in a partition $\mathscr{A}$ of a set A, any member x of X is related to every other member y in X by an equivalence relation R, for which the following properties hold:

Definition 1.2.7 (Equivalence Relation)**.** For $x, y, z \in X$ a block in a partition $\mathscr{A}$ on the set A, a relation R is an equivalence relation iff

1. R is reflexive: $\forall x : Rxx$
2. R is transitive: $\forall x, y, z : Rxy \wedge Ryz \Rightarrow Rxz$
3. R is symmetric: $\forall x, y : Rxy \Rightarrow Ryx$

Issues raised by interrogatives correspond to alternative blocks in a partition induced by the interrogative. Answer to questions can be either address or resolve issues. If an answer ψ addresses an issue, it is *partial*; if it resolves an issue, it is *complete*. Complete answers are maximally informative with respect to the issues raised by an interrogative in that they eliminate all but the one block of worlds consistent with the ψ. Partial answers on the other hand provide incomplete information with respect to the query by eliminating some partitions or worlds, but not necessarily along the lines delimited by the question. Definition 1.2.8 formalizes the distinction between types of answers (from Aloni (2001)):

Definition 1.2.8. [Answers] Let ψ and $?\overrightarrow{x}\phi$ be closed sentences in our language.

1. ψ is a *partial answer* to $?\overrightarrow{x}\phi$ in a model M iff
$$\exists X \subset [[?\overrightarrow{x}\phi]]_M : [[\psi]]_M = \cup\{\alpha \mid \alpha \in X\} \neq \emptyset$$
2. ψ is a *complete answer* to $?\overrightarrow{x}\phi$ in a model M iff
$$[[\psi]]_M \in [[?\overrightarrow{x}\phi]]_M$$

Having provided the absolute minimum background required for the chapters that follow, we now turn to the characterization of the central topic of the thesis: concealed questions.

2

Concealed Questions

2.1 Descriptive Background

The term 'concealed question' describes a nominal which can be paraphrased as a *covert indirect question*. For instance, (4a) and (5a) might be paraphrased as (4b) and (5b), respectively. Notice that the complex nominal, *e.g., the mayor of the town*, receives an identity question paraphrase, *e.g., who the mayor of the town is*.

(4) a. John guessed the mayor of the town
    b. John guessed who the mayor of the town is

(5) a. John predicted the cost of the reform
    b. John predicted what the cost of the reform is

Currently, there is no general consensus among semantic theories as to the denotation of concealed questions (CQs), save what they *are not*. Namely, it is generally agreed that CQs are not, despite their name, questions of any sort.[1] Further, most current proposals claim that CQs do not denote individuals, at least individuals as typically denoted by other nouns (see Frana (2007) for a lucid dissenting view). Three current proposals which attempt to derive CQ denotations from the key type of the noun are briefly reviewed in §2.2; the insights and attendant difficulties with the proposals are considered.

[1] See Grimshaw (1979) for the view that concealed questions are semantically questions, but are syntactically DPs. This division, cashed out in terms of s-selection and c-selection has the advantage of predicting the distribution of admissible CQ contexts, but also faces several empirical problems, as detailed in Nathan (2006). For reasons of space, Grimshaw's proposal will not be discussed below.

First, however, we turn to the classic arguments that deny the denotation of concealed questions membership among that of their fellow nouns or question cousins. The argument against CQs as individual denoting DPs is based on five 'core properties' of CQs. To briefly glance ahead, note that not all of these properties is above contention as a defining characteristic of a CQ. In fact, I argue that the fifth tenant of CQs – that *only an individual* can be said to answer the identity question associated with the paraphrase – is misguided. In Chapter 3, I offer a more flexible account of CQ meanings, with this consideration as a conceptual core of the theory. But, we get ahead of ourselves.

2.1.1 The Core Properties of CQs

The major accounts of CQ in the literature have spawned a stock of observations about the linguistic behavior of those DPs paraphrasable as covert questions. These core descriptive properties are summarized in (6); any account of CQs must provide a proposal that adequately addresses each of these properties.

(6) Core Descriptive Properties of CQs
  i. CQs are paraphrasable as indirect questions
  ii. CQs fail entailments under propositional attitude reports
  iii. CQs don't support gendered anaphora cross-linguistically
  iv. CQs don't conjoin with extensional verbs taking individuals as arguments
  v. Acceptable answers to question paraphrases are restricted to individuals

2.1.1.1 Question Paraphrase

First and foremost, CQs are, by definition, paraphrasable by covert questions, which typically involve a sort of identification of the description mentioned in the noun phrase. Naturally, the fact that CQs are defined as such begs the question of whether and to what extent the correspondence to a paraphrase reveals any kind of semantic kinship. Viewed in such a light, theories of CQs can be categorized according to what kind of relationship they assume to hold between the DP and the question paraphrase. Approaches can be divided into two very coarse categories, with two finer grained distinctions apiece.

First, there are ontological *reductionist* approaches, the first subtype of which relates the interpretation of the DP more closely to its question paraphrase;[2] the second attempts to

[2]Grimshaw (1979) and Nathan (2006) fall into this category, although through markedly different proposals. For Grimshaw, CQs semantically denote questions, but are syntactically DPs. Assuming that questions are sets of true propositional answers after Hamblin (1973) and Karttunen (1977), Nathan proposes that CQs are in fact propositions (see §2.2.2). To some extent, Romero (2004) is also in this group, since CQs are turned into propositions through a covert answer operator.

associate the DP with more standard accounts of DP denotations.[3] These theories implicitly claim that CQs are best thought of in terms of other, more primitive notions. Reductionist approaches have far been the most popular, trending towards taking the paraphrase as indicative of semantic class.

Second, there are *non-reductionist* approaches, the first subtype of which takes the CQ as an ontological primitive and relates its meaning neither to the question or the noun phrase. Such a theory would have to suppose that the paraphrase relation is accidental, and explain away the fact that the CQ appears syntactically as a DP. To my knowledge, such an approach has gained no adherents. The second subtype of non-reductionist analysis endeavors to establish connections between the question paraphrase and the noun denotation without reducing its semantics to one or the other. The last type of proposal is one that the present thesis argues for, in that the paraphrase is thought to reflect not ontological reduction, but rather *functional congruence* between the CQ and the corresponding question. That is, CQs and their paraphrases are argued to have the same effect on the context, without completely reducing the semantics of one to the other. Whatever the details, a maximally informative account must not only uncover the correct denotation for an NP with a CQ interpretation, but must also explain its Janus-like nature.

2.1.1.2 Entailment Patterns

Secondly, as noted by Heim (1979), verbs that support CQ readings typically fail to support valid entailment patterns with CQs:

(7) I. John knows the capital of Italy
II. The capital of Italy is the largest town in Italy
____
III. John knows the largest town in Italy

From the CQ reading of the first premise, *John knows what the capital of Italy is*, the extensional equivalence in the second premise does not guarantee the validity of the conclusion from the substitution of identical terms (see §3.1.1 for more discussion). Simply put, John's knowledge in the first premise is limited to knowing what town the description *the capital of Italy* denotes – as such no other equivalences, such as the size of its population, are relevant. As a consequence, the equivalence in the second premise does not yield a valid substitution instance under the attitude verb in the conclusion.

As with other puzzles of identity and opaque contexts (among the extensive literature, see Aloni (2001); Bonomi (1995); Crimmins and Perry (1989); Kaplan (1969); Kripke (1979); Quine (1956)), whether the substitution holds depends on the perspective of the content of the attitude. For instance, supposing that we are not speaking about a *de dicto* belief that John holds, but rather a *de re* attribution possibly evaluated outside of John's beliefs, (8a)

[3] As Frana (2007) argues that CQs are individuals, her account could be classified as a member of this category. See also Grimshaw (1979), for whom CQs are syntactically DPs.

would not necessarily be contradictory, although this shift is marginal with a fully explicit question paraphrase (8b).

(8) a. John knows the largest town in Italy, but he doesn't know he knows it

b. ? John knows *what the largest town in Italy is*, but he doesn't know he knows it

Example (8a) lends itself more easily than example (8b) to an interpretation in which it is *our* way of identifying the *the largest town in Italy* which is important. We take this contrast to mean that the way of identifying an object, and not merely the object simpliciter, is a crucial component of the CQs truth conditions. However, complete discussion of this issue would lead us too far astray from the task at hand, and is thus postponed until the next chapter.

2.1.1.3 Anaphora

Thirdly, CQs do not support gendered anaphora, in that CQs must be anaphorically related to the neuter pronoun (examples from Romero (2004)).

(9) a. The winner of the Oscar for best actress walked in. **She/*he/*it** was wearing a red dress.

b. John guessed the winner of the Oscar for best actress before I guessed **it/*her**.

The examples in (9) are typically taken to show that CQs do not denote individuals. The DP *the winner of the Oscar for best actress* denotes an individual in (9a), and accordingly, the gender marking is reflected in the referring pronoun (*she*). Example (9b) does not denote an individual; in English, a gendered pronoun is inappropriate in such a context. As Romero (2004) reports, this pattern is replicated in other languages, including Finnish below:

(10)

Naispääosa-Oscarin	voittaja	astui	sisään
Female-lead-Oscar-GEN	winner-NOM	stepped/walked	in

Hän	oli	pikeutunut	punaiseen	pukuun
She/he-NOM	was	dressed	red-ILLATIVE	evening-dress-ILLATIVE

'The winner of the Oscar for best actress walked in. She was wearing a red dress.'

(11)

Tytttö	joka	aiheutti	tämän	ongelman	ei	ollut	Mari
Girl-NOM	who	caused	this-ACC	problem	neg	was	Mari-NOM

Se	/ # **hä**	oli	Liina
It-Nom	/ # **she/he**	was	Liina-NOM

'The girl who caused the problem was not Mary. **It /#she** was Liina.

However, neuter pronoun reference in English is a wide and still ill-understood phenomena (see Huang (2001) for a detailed survey). The neuter pronoun *it* may refer to (inanimate) individuals, to intensionalized individuals, to propositions, or even to events. Hence, although these examples may indeed show that CQs are not interpreted as individuals, the question of what they *do* denote is not resolved by consideration of the present anaphora facts. Distinguishing between these ontological possibilities for CQ denotations, *e.g.,* propositions or other sorts of intensional objects, is simply not possible by means of the anaphora facts alone.

2.1.1.4 Coordination

Fourth, CQ verbs, such as *predict* or *determine*, do not conjoin with verbs that select for individual entities only, such as *notice* or *kiss*. In general, verbs conjoin just so long as they match on the relevant dimension;[4] in this case, verbs must take the same sort of argument to successfully coordinate, as in example (12). In (12a), verbs selecting for a CQ argument conjoin; in (12b), verbs exclusively selecting for individual arguments conjoin as well (examples modified from Romero (2004)). But when the two respective sorts of verbs combine, as in (13a), the CQ interpretation of the DP cannot be maintained, as shown by the preference for gendered anaphora (13b) over the neuter form (13c).

(12) Matching verbs
  a. John predicted and later determined the winner of the contest
  b. John noticed and later kissed the winner of the contest

(13) Conflicting verbs
  a. # John predicted and later kissed the winner of the contest
  b. She was wearing a read dress
  c. # Before I guessed it

Along with the third point on anaphora, the matching verb effect shows only what CQs are *not*, namely regular noun denotations. To further establish what, in fact, CQs are requires additional empirical and theoretical insight.

2.1.1.5 Greenberg's Observation

The fifth and last of the classical CQ characteristics, is based on an observation by Greenberg (1977). It has often been observed that CQs lack an ambiguity present in their question paraphrase (see Greenberg (1977); Grimshaw (1979); Heim (1979)). Given the CQ in (14a) and its question paraphrase (14b), the CQ is said to lack the reading expressed in (14.ii), which the full embedded question enjoys. Both (14a–b) can be interpreted as (14.i).

[4] See Partee and Rooth (1983) for an influential formal treatment of conjunction.

(14) a. Officer Hopkins found out the murderer of Smith
b. Officer Hopkins found out who the murderer of Smith was

i. Officer Hopkins resolved the question of who murdered Smith, by identifying the individual.

ii. Officer Hopkins resolved the question of who murdered Smith, by finding out some essential fact about individual denoted by *the murderer of Smith*.

Greenberg's observation has long since been cashed out as a strict condition on answerhood; the question which paraphrases the CQ is about *the identity of an individual*, and any proper answer to that question must involve picking out an individual under identity. Although this position has been codified in various ways, the moral is always the same: there is no situation in which a property or fact of equivalence could satisfy a CQ interpretation. In other words, situations or facts about identity are assumed never to provide a felicitous answer to the corresponding question paraphrase.[5]

For the great majority of cases, this restriction on CQ interpretations is perfectly adequate. However, as argued below, there are numerous situations in which a property can serve as an answer to the identity question under discussion, namely when resolving the question of identity is satisfied by specifying a concept essential to that individual. In this situation, readings (14.i) and (14.ii) converge.

Such cases, along with some initial limitations, are presented in §2.1.2 immediately below and expanded upon further in §3.3.1. I argue that these cases are systematic, no matter how restricted, and that a proper theory of CQs must address the conditions under which such a reading is possible.

2.1.2 On the Interpretation and Structure of CQs: Two novel observations

This section presents two novel observations about concealed questions, which, either in part or in whole, prove to be problematic for previous analyses of CQs. The main topics are listed in (15) below, and comprise a test for critically examining the adequacy of the existing views in §2.2.

(15) i. Concepts and Properties as Answers
ii. CQs in ACD and Sluicing Environments

These observations are discussed here because the account I advocate in the next chapter is designed explicitly around these facts. In §3.6.2, I present additional facts that distinguish different types of relational nouns that, to date, no account of CQs can account for.

[5]Recently, Romero (2005) and Frana (2007) discuss alternative contexts of use for CQs in passing. However, a detailed examination of the contexts which support deviant interpretations is more or less absent.

Before delving into the details of these observations, the central gist of each might first be established. Observation *(i)* constitutes extended evidence against the generality of Greenberg's observation, that is, the constraint limiting the answers to the CQ paraphrase to individuals. Observation *(ii)* explores CQ interpretations in markedly different syntactic environments, contexts that do not straightforwardly support any previous analysis. Taken together, the central goal of this section is to illustrate that CQs encompass a much wider range of interpretations in remarkably subtle ways, and to convince the reader that a full account of CQ interpretations must, at the absolute minimum, address the data presented below.

Lastly, a minor qualification is in order. In the immediately following sections, there is little by way of concrete analysis. The aim is to present the data without prematurely encumbering it with a particular interpretation. Those anxious for a formal account of the data will find one in the next chapter. The hope is that the data are sufficiently interesting in their own right, quite apart from specifying exactly what classes of theory they support.

2.1.2.1 Concepts and Properties as CQ Answers

Since Greenberg (1977) observed that CQ interpretations are typically not ambiguous in the way that their question paraphrases often are, every subsequent analysis of CQs has ignored or failed to elaborate on exceptions, although the contexts that support such exceptions are systematic.

In this section, we mean to establish that *individual concepts* and *identifying properties* may sufficiently answer the question paraphrase associated with the CQ. For instance, suppose that one of Pontius Pilate's many duties is to set the price of milk, the availability of which varies according to season, violence in the region, and so on. Pilate, however, has no interest in attaching a numerical value to milk, only to change it in times of crisis. So, clever Pilate declares that the milk should always cost the same as another product: orange juice, say, which comes from the same regions from which milk is imported. Thus, a dairy manufacturer could utter (16a), knowing full well that Pilate has not actually set the price of milk to a numerical value of any sort, but rather to an individual concept, as in (16b–c). Thus, what Pilate does when he sets the price of milk in this indirect fashion, is to fix the intension of the price to another intension, in our case the price of orange juice.

(16) a. Pilate set/determined the price of milk
  b. Pilate set/determined the price of milk is, namely, the price of orange juice, whatever that may be
  c. It's the same as orange juice, which his subordinate fixes

The basic pattern extends to other CQ verbs as well. Upon hearing (16), I can assert (17) to the reader, although neither one of us knows anything else about how much one might have paid for milk (or orange juice, for that matter).

(17) You now know the price of milk in Pilate's era. It cost whatever orange juice did.

The appearance of a wider meaning in CQ interpretations is not limited to *set* and *determine*. It enjoys the full range of CQ selecting verbs, including *know, predict, guess*, and so on.

(18) a. John predicted/guessed/recalled the president of the LSA

b. He predicted/guessed/recalled it is the linguist with the most articles published in Language

Imagine that John is a statistician researching management trends in professional and academic organizations. He discovers that the person elected president of the LSA is always the person who has published the most articles in Language the year before. Perhaps this fact is merely an odd anomaly. Perhaps it is stated as such in the by-laws of the organization. In any event, interpreting (18a) as (18b) is perfectly felicitous.

It is furthermore instructive to examine cases in which a property answer does *not* satisfy a CQ interpretation. Returning to Smith's murder case, above, suppose that from the way in which Smith was murdered, Officer Hopkins concludes that Smith's murderer, whoever he is, must be insane. On the basis of (19) and/or (20), neither (21a) nor (21b) is felicitous.

(19) Officer Hopkins knows that Smith was murdered

(20) Officer Hopkins predicts that Smith's murderer is insane

(21) a. # Officer Hopkins knows Smith's murderer, *i.e.*, that he is insane

b. # Officer Hopkins predicted Smith's murderer, *i.e.*, that he is insane

Under normal circumstances, we understand that the property of *being insane* is not unique enough to pick out the denotation of the CQ *Smith's murderer*. Suppose, however, that Officer Hopkins remembers an unsolved case from 1973 in which a murderer mutilated his victims in an identical manner. Reports of the killings were never published and so it could not have been a copycat. Officer Hopkins has reason to believe that the identity of the murderer in these crimes past and present is one and the same, despite the fact that beyond a collection of facts collected at the crime scene, no further identification can be made.

Thus, although the true identity of the killer has not been discovered, both (22) and (23) is felicitous and might be truthfully uttered about Officer Hopkins.

(22) Officer Hopkins has guessed Smith's murderer; he thinks it's the same monster who evaded capture in 1973.

(23) Officer Hopkins has guessed who Smith's murderer is; he thinks it's the same monster who evaded capture in 1973.

To recapitulate, a property as answer must uniquely identify the individual in question, either as a sufficiently unique coextensive property (Smith's murderer) or as an individual concept (the price of milk). The appropriateness of such a identifying relation is sensitive to the context in which it is uttered, and it explored further in the next chapter.

2.1.2.2 CQs and Ellipsis

The second novel observation centers around an observation made by Heim (1979), whose insights will be discussed in detail in the next section. Here, I present the so-called Heim's ambiguity: CQs containing a relative clause are sometimes ambiguous with respect to whether the semantic contribution of the relative clause is evaluated with respect to the subjet's belief worlds or not. I argue that these different readings in fact have very different syntactic structures as well. Evidence for this claim is brought out by particular syntactic environments involving ellipsis: in particular, Antecedent Contained Deletion (ACD) and Sluicing contexts.[6]

In her classic paper on concealed questions, Heim (1979) observed that there is an ambiguity in sentence (24) between two types of knowledge that John might have. Reading A supports an interpretation in which John and Fred know the same price. Reading B supports an interpretation in which John knows the answer to the question "What price does Fred know?" Accordingly, Romero (2005) has categorized Reading B as the *meta-question* interpretation.

(24) John knows the price that Fred knows

  A. John knows the same price that Fred knows

  B. John knows the answer to the question "What price does Fred know?"

The arguments below crucially depend on the two following assumptions: (i) Reading A affords a concealed question interpretation, and (ii) that the complement of Reading B, the meta-question interpretation, is semantically a proposition. These assumptions are commonly shared in most of the previous literature, and trivially so for those theories which equate the semantic type of reading A and B as propositions (Nathan (2006), for instance, argues that all CQs are propositions and derives reading B from reading A from another type shifting rule). Whether these semantic types are truly one and the same is the central topic of this section.

That at least Reading B should be analyzed as a proposition makes intuitive good sense, at least when the meta-question is recast as a typical propositional attitude. In particular, the meta-question regarding John in (24B) is a garden-variety statement about John's knowledge state.

(24′) B′. John knows which price it is that Fred knows

However, the common ground between some prior proposals and the one presented in the next chapter stops with these assumptions, for I will argue that although Reading A is a

[6]This section greatly benefited from Kyle Johnson's instruction and good will. In addition, many of these judgments and arguments were discussed with Paul Elbourne, Lyn Frazier, Angelika Kratzer, Chris Potts, Craige Roberts, and Maribel Romero.

CQ, and Reading B is a proposition, CQs are not propositions, as a minor extension of Heim's ambiguity (24) to ACD contexts (25) illustrates.

It is a surprisingly robust fact that when example (24) is modified to elide the most embedded verb *know* that Reading B of Heim's ambiguity is altogether absent.

(25) John knows the price that Fred does
  A. John knows the same price as Fred
  B. # John knows the answer to the question "What price does Fred know?"

Not only is this contrast surprising under the reductionist view, it also goes unexplained under the traditional Quantifier Raising (QR) account of ACD (see Berman and Hestvik (1991) for a review of the various conceptions of QR and ACD). However, before showing precisely how the absence of Reading B is problematic, it must be first established that example (25) concerns ACD itself and that it is not a general fact about ellipsis.

As evidence that the contrast in (25) concerns ACD in particular and not ellipsis in general, observe the paradigm in (26), which illustrates the fact that cross-sentential ellipsis licenses either reading (26A–B), but not both at once (26C).

(26) John knows the price that Fred knows. Sally does, too.
  A. John, Fred, and Sally know the same price
  B. John knows which price Fred knows. Sally knows which price Fred knows, too.
  C. # John and Fred know the same price. Sally knows which price John/Fred knows.

As expected, the cross-sentential VP ellipsis cannot be reinterpreted as Reading B, as a propositional attitude, when the preceding context does not supply this reading.

(27) John knows the price that Fred does. Sally does, too.
  A. John, Fred, and Sally know the same price.
  B. # John knows which price Fred knows. Sally knows which price Fred knows, too.

Together, examples (26) and (27) illustrate cross-sentential ellipsis licenses both Readings A and B, but that the available reading is dependent on whether the reading is likewise supported by the preceding phrase. This suggests that the lack of Reading B cannot be attributed to VP ellipsis itself, and that the antecedent for ellipsis is sensitive to the semantic representation made available to it. We conclude that the failure of ACD to license reading B is due to reasons which are independent from VP ellipsis in general, but are instead specific to conditions licensing ACD in particular.

Antecedent contained deletion (ACD) is so-called because the syntactic constituency that would serve as an antecedent to the elided VP contains the elided VP itself, demarcated by $\langle \cdot \rangle$ as in (28).

(28) John [$_{VP}$ ridiculed every syntactician Fred did ⟨ ridicule *t*⟩]

Although ellipsis has been much explored, the precise conditions that license it are not firmly established. For instance, although some degree of semantic parallelism between the elided and overt VP is required, as expected in VP ellipsis, it is not clear what semantic dimensions parallelism respects.

Nonetheless, I will assume, following Fox (2002) and Heim (1997) and works cited therein, some version of parallelism which requires that the deleted VP and its antecedent share logically equivalent semantic representations. This rough characterization is formalized in (29).[7] Only when it is satisfied does the VP have license to elide.

(29)
i. PARALLELISM: An elided VP must be identical to an antecedent VP at LF modulo alphabetic variance and indexation.
ii. ELIDE! Do not pronounce a string which satisfies PARALLELISM.

One popular explanation, then, is that the covert operation of Quantifier Raising (QR) applies to the complement DP (see Sag (1976) and Larson and May (1990), as well as Fox (2002, §1) for discussion), so that the elided VP satisfies PARALLELISM. For example, our example (28) would be analyzed as (30), in which the QR applies to the DP before PARALLELISM is checked.

(30) [every syntactician that Fred did ⟨ ridicules *t*⟩]
John ridicules *t*

We see below that while this analysis can derive Reading A, its success cannot account for the lack of Reading B, at least not without alteration.

Applying the standard QR analysis to the CQ complement in (31) yields a representation in which (i) QR has adjoined the DP to a position dominating the matrix clause,[8] and (ii) the trace left behind is indistinguishable from the trace in the elided VP (see Heim and Kratzer (1998)). The syntactic representation and its corresponding Logical Form (LF) are shown in (31).

[7] The constraints listed below are formalized in the style of Optimality Theory (Prince and Smolensky (1993); McCarthy (2002)), in order to give only the most general idea of what such constraints could look like.

[8] Although the exact cite of adjunction is often left unspecified, it is often assumed to adjoin to the CP dominating the matrix clause.

(31) John knows the price that Fred does

i. Syntactic Representation of ACD under CP adjoining QR:

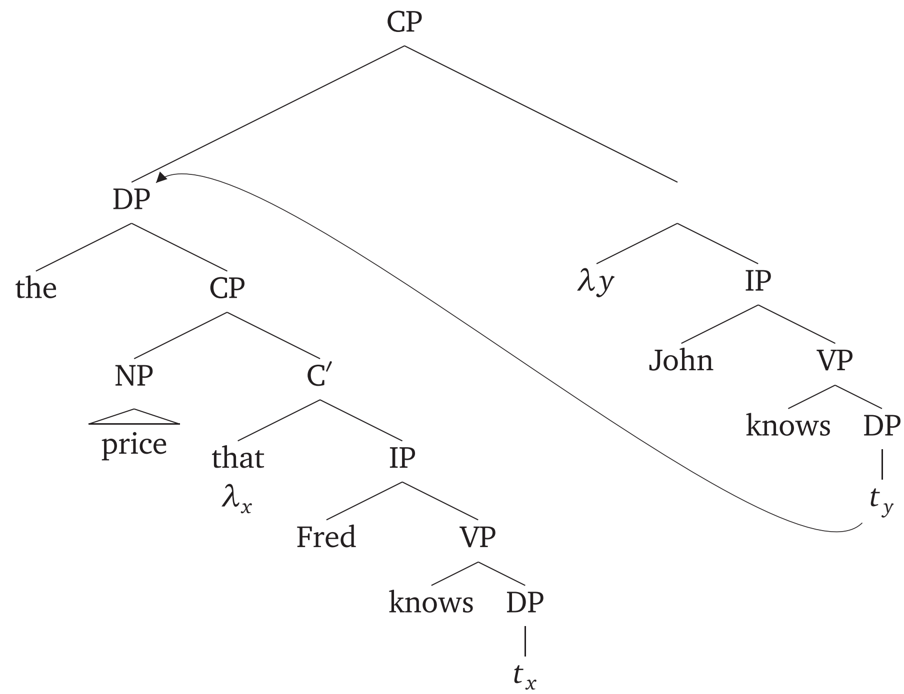

ii. Input to Logical Form
[[the price λx Fred knows t_x]
λy John knows t_y]

iii. PARALLELISM is satisfied, ELIDE
[the price λx Fred *does* ⟨ knows t_x⟩]
λy John knows t_y

While the uniformity of the trace representation is able to explain why the embedded VP elides, it does not distinguish between Readings A and B, as their representations are identical on this analysis. A simple, and accurate, amendment could differentiate the two merely by stipulating that the types of trace left by a nominal element and a propositional element are qualitatively distinct, and further that PARALLELISM is sensitive to this difference. Thus, (32B) fails to be licensed because it disobeys PARALLELISM.

(32) John knows the price that Fred does
B [the price that Fred does ⟨ knows t_e ⟩]
John knows t_{st}

Another equally plausible solution defines QR as an operation which applies to nominal (DP) element only. Thus, a propositional complement can never be QR'd in such a way that will allow it to satisfy the licensing conditions for ACD. Under this conception, QR would be sensitive the semantic denotation of an expression, rather than (only) its syntactic category. On many current models of the syntax-semantics interface, this would strongly suggest that QR applies exclusively to semantic representations. It remains to be seen what other evidence could be marshalled in favor of such a limitation.

However, Merchant (2000) and Fox (2002) note that there are independent reasons to revise the above analysis of ACD, as illustrated by the disjoint reference effect in (33a), which goes unexplained under the simple trace account representation (33b).

(33) a. ?? Guess which friend of John's$_i$ he$_i$ visited *t*

b. [which friend of John's$_i$] he$_i$ visited *t*

The problem is that, according to the Binding Theory,[9] the proper name *John* and the pronoun *he* in (33a) are thought not to be interpreted as co-referential, as doing so would violate Principle C. However, if the Binding Theory holds over semantic representations, the representation (33b) affords no straightforward explanation for why Principle C should still hold. Given the representation above, we should expect coreference to be possible as *John* is no longer in the binding domain of *he*, a coreferential pronoun.

Fox (2002) develops a solution to the bleeding of Condition C that pins the blame squarely on three properties of the above account: firstly, the 'impoverished status' of traces, secondly, the treatment of QR, and three, the ways in which an adjuncts are merged. Fox (2002) adopts the "Copy Theory of Movement" (Chomsky, 1995), in which *copies*, not *traces*, of a constituent remain in every position where it is merged, and are subsequently deleted at the interface to phonology. By so doing, he can account for the failure of QR to bleed Condition C of the Binding Theory, since the R-expression *John*$_i$ occurs free in its binding domain, whereas the bound pronoun *he*$_1$ has *Peter* as its antecedent. However, PARALLELISM is violated, as the copy is not itself a trace, as in the complex nominal raised by QR.

(34) [every guy Peter$_1$ wanted me to ⟨ introduce him$_1$ to t ⟩]

I introduced him$_1$ to [every guy Peter wanted me to *t*]

[9] The binding of anaphors has received much attention from linguists and philosophers, in general, and generative grammarians, in particular. The generative conception of anaphora binding is encoded in the binding theory; three conditions which determined when a type of anaphor is licensed.

(1) The Binding Theory

A. An anaphor is bound in a local domain

B. A pronominal is free in a local domain

C. A R(eferring)-expression is free

Generally, an anaphoric element is bound when it depends on an antecedent for interpretation, and is free otherwise. The precise notion of 'domain' has been the subject of much debate, and will not be reviewed here (but see Reuland and Everaert (2003) and references cited therein).

Fox (2002) introduces a rule of *trace conversion*, a conjunction of rules which converts "the copy at the tail of the chain to an interpretable object." In effect, trace conversion adds a variable to the representation of a predicate (35.i), and then lowers the complex into a definite determiner (35.ii), so that Parallelism is satisfied.

(35) Trace Conversion
  i. Variable Insertion: (Det) Pred → (Det)[Pred $\lambda y(y = x)$]
  ii. Determiner Replacement: (Det)[Pred $\lambda y(y = x)$] → the [Pred $\lambda y(y = x)$]

In addition, Fox (2002) assumes (i) that QR moves the complement DP rightward as an adjunct to VP, and (ii) that the restrictive relative clause appears in syntax *afterwards* as a consequence of Late (Adjunct) Merger, forming an externally headed relative clause.[10] Crucially, the relative clause is formed by adjoining a CP to an NP. The internal head is thought to move to a position where it and all its copies are deleted under the somewhat mysterious condition Delete under Identity, defined with Late Merger below:

(36) i. Late Merger: An instance of (countercyclic) subtree insertion in which an *adjunct* is merged *after* the head to which it is a sister is merged into the syntax.
  ii. Deletion under Identity: A terminal node which is in some way 'identical' to another string adjacent terminal node must or can be deleted.

After these operations have created a relative clause, Trace Conversion applies, making the LF representations of the elided VP and the matrix VP similar, or similar enough *modulo* the choice of variable inserted, to satisfy Parallelism. The derivation in (37) illustrates how Fox's system handles the case of interest. Note that struck-out text signifies that it is unpronounced.

(37) John knows the price that Fred does
  i. The DP is merged sister to the verb creating a VP. The subject *John* is later merged. For convenience, the resulting phrase is presented as a VP.

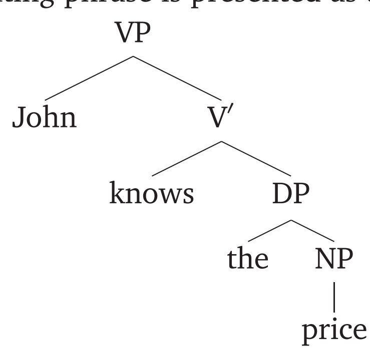

[10]See Carlson (1977) for an early discussion of the ambiguity of relative clauses, and more recent developments from Bhatt (2002), Hulsey and Sauerland (2006), and Sauerland (1998). Late Merger originated with Lebeaux (1988) and Lebeaux (1990), with extensions and developments discussed in Hulsey and Sauerland (2006) and Takahashi (2006). See Bianchi (2002a, 2002b) for review.

ii. QR right adjoins a copy of the DP to VP.

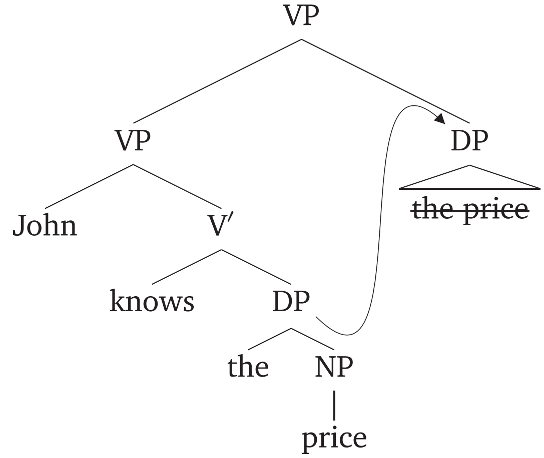

iii. The relative clause late merges to the adjunct, forming an externally and internally headed relative clause in which internal head raises to the specifier of CP and is deleted under identity.

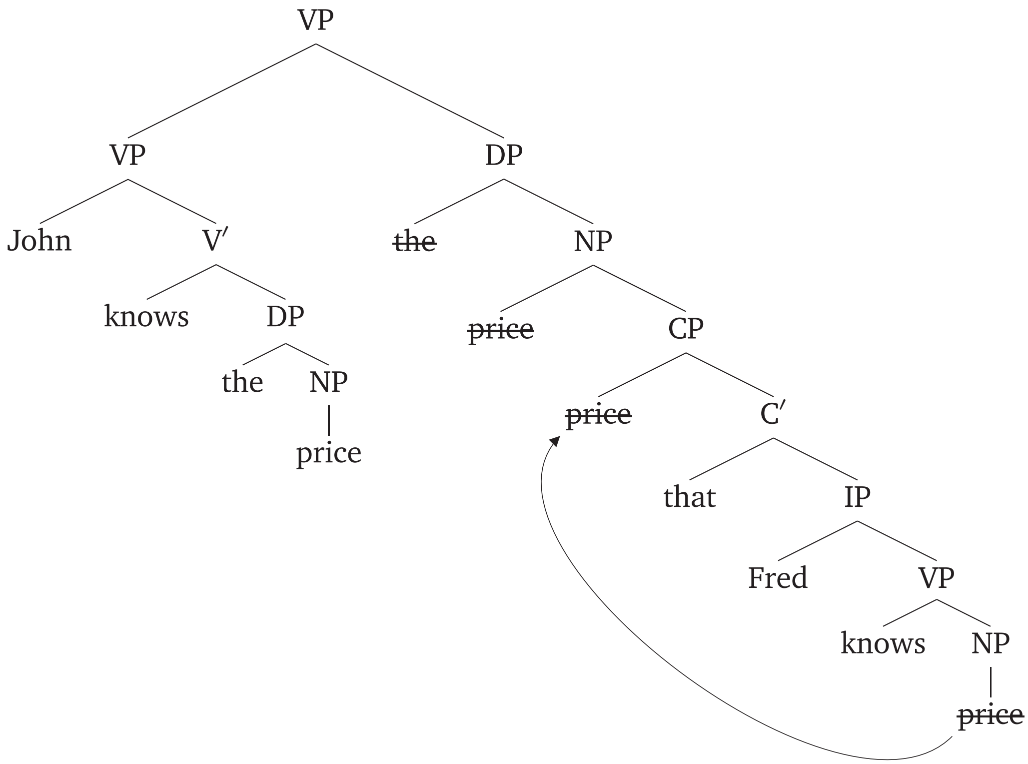

This derivation is schematized in (38), for future reference. The structure (38a.iii) is the input to LF where different copies are pronounced. At LF, (38b) undergoes Trace Conversion (38c), and then is subject to ELIDE! in (38d).

(38) a. Copy Theory

i. [$_{VP}$ John knows the price] $\longrightarrow^{QR}$

ii. [[$_{VP}$ John knows the price] ~~the price~~] $\longrightarrow^{Late\ Merger}$

iii. [[$_{VP}$ John knows the price] ~~the price~~ that Fred knows ~~price~~]

b. Input to LF:

[the price [$_{VP}$ Fred knows price]
[$_{VP}$ John knows the price]

c. Trace Conversion

[the price λx. Fred knows the price x
λy. John knows the price y]

d. ELIDE!

the price λx. Fred does ⟨ knows the price x⟩
λy. John $\underbrace{\text{knows the price } y}$

Note that this derivation is not the only one available to Fox (2002). An alternate route involves merging the relative clause in-situ in the matrix clause. However, if QR right adjoins the full structure as before, PARALLELISM will not be satisfied at LF, as shown in (39) below. In contrast to the earlier analysis of ACD, the representation of the relative clause is essential in that only relative clauses formed via Late Merger are available for ACD.

(39) a. Copy Theory

i. [$_{VP}$ John knows the price ~~price~~ that Fred knows ~~price~~] $\longrightarrow^{QR}$

ii. [[$_{VP}$ John knows the price ~~price~~ that Fred knows ~~price~~]
~~the price that Fred knows price~~]

b. Input to LF:

[[the price that Fred knows price]
John knows the price that Fred knows price]

c. Trace Conversion:

[[the price λx Fred knows the price x]
λy John knows the price that Fred knows the price y]

The central claim of this subsection is that the fact that Reading B is absent in ACD contexts should be united with Fox's idea that the parallelism required for ellipsis in such contexts depends on the Late Merger of the relative clause. In other words, a late merged relative clause is responsible for Reading A, whereas an early merger of a relative clause to an NP is uniquely linked to Reading B. We may now organize the data we've seen into the following table:

Merge	Reading	ACD Licensed?	Relative Clause Head	Relative Clause Type
Late	A	✓	external	Matching
Early	B	*	internal	Raising

Table 2.1: Relative Clause Structure and Reading Correspondence

The specific hypothesis is codified into the Headedness Conjecture below, in which the two readings employ two different, and exclusive, representations.

Hypothesis 2.1.1 (The Headedness Conjecture (HC)).

1. *The semantics of Reading A require an externally headed relative clause (the Matching analysis)*

2. *The semantics of Reading B require an internally headed relative clause (the Raising analysis)*

The HC can potentially explain the difference in available readings between (24A) and (24B), without further stipulation to Fox's approach to ACD. If (i) Reading B manifests only as a CP and (ii) QR (ii.a) applies only to nominal categories, and (ii.b) is required to license ACD, then it follows that Reading B is illicit in ACD structures. That is, the type of structure needed to produce Reading B is simply not available in ACD environments.

Although much more is required to flesh out such a hypothesis, some observations initially appear to be in its favor.[11] First of all, we expect that only Raising structures allow phrasal idioms. Thus, if the HC is correct, we expect only reading B is possible with a phrasal idiom, such as *make headway*:[12]

(40) #/✓ John knows the headway that Fred made

Given the assumption that Matching structures don't support phrasal idioms, Hulsey and Sauerland (2006) argue on the basis of (41) that Raising structures don't permit extraposition past temporal adverbs.

(41) a. Mary praised the headway that John made
b. *Mary praised the headway last year that John made

If Reading B requires a Raising analysis, then we expect that extraposition past a temporal adverb should be impossible with such an interpretation. Despite some difficult judgments, I believe this prediction is borne out in (42). While the Reading B paraphrase cannot be

[11] See Henderson (2007) for a critical overview of many of these diagnostics.

[12] Unless the immediate context is contrasting something specific, the acceptability marks #, ?, ∗ on the left of the '/' symbol represent judgments for the reading A interpretation, while those to the right of the '/' symbol represent judgments for reading B.

obtained either as a full relative clause or in the ACD context, Reading A is allowed in either case, as expected.[13] If on the left side of the slash symbol '/', the acceptability judgment corresponds to the the full verb (*knows*), and if on the right, it corresponds to the elided verb (*does*).

(42) a. John knew the price yesterday that Fred knew/did

A. ✓/✓ John and Fred both knew the same price yesterday

B. ??/#John knew which price Fred knew yesterday

Secondly, Reading A and B behave differently with respect to superlative adjectives, such as *highest*. Obscuring the details here for reasons of space, Hulsey and Sauerland (2006) suggest that the available readings depend on possible positions created by movement to interpret the world variable associated with the adjective. It seems possible to attribute the difference of interpretation in (44) to different relative clause types. At the very least, it seems that the most salient interpretation of the ACD variant is one that bans the low reading of the adjective. This idea forges a close connection between the fact that Reading B is unavailable in ACD contexts and the absence of a low reading in the very same cases. This connection is naturally expected under the HC above.

(43) John knows the highest price that Fred knows

i. *High reading: Matching Structure*
the highest price about which John knew that Fred knew

ii. *Low reading: Raising Structure*
the price x such that John knows the highest price Fred knows is x

(44) John knows the highest price that Fred knows/does

A. ✓/✓ John and Fred know a price, which happens to be the highest price that Fred knows (*high*)

B. ✓/# Of all the prices (that Fred knows), John at least knows which price is the highest one that Fred knows (*low*)

Clearly, additional evidence and argumentation must be offered in order to validate this pattern. Yet, as an initial examination, the judgments pattern with the predictions of the HC.

[13] The judgments become significantly less clear with multiple temporal adverbs. It is unclear to me whether (1a) can have a Reading A interpretation, although it seems clear that Reading B in the ACD environment is still impossible. Interestingly, I think that (1a) and (1b) have the same readings, although only in (1a) is the relative clause head thought to be extraposed.

(1) a. John knows the price today that Fred knew/did yesterday

A. ✓/? John knows the same price today that Fred knew yesterday

B. ?#/# Today, John knows which price it is that Fred knew yesterday

b. John knows the price that Fred knew/did yesterday today

Lastly, sluiced variants of Heim's ambiguity admit interpretations which are the mirror image of the ACD pattern. Identified by Ross (1969), the term sluicing describes sentences in which an interrogative clause (46) reduces to a *wh*-phrase, as in (45).

(45) a. John bought something, but I don't know what
b. Sally made a cheesecake – you'll never guess for who!

(46) a. John bought something, but I don't know what he bought
b. Sally made a cheesecake – you'll never guess for who she made the cake

We follow the extensive work in Merchant (2001), and assume that sluicing involves a CP that dominates an IP which is deleted for semantic reasons. In (47), the general form of a sluiced clause is depicted, in which the struck-out text is to be interpreted as phonologically null, but syntactically present (see Merchant (2001) for details).

(47) General schema for sluiced clauses

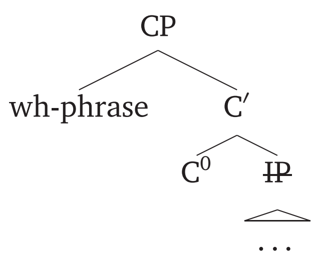

Thus, if the HC is correct, then Reading A should be banned from sluicing contexts, as Reading A requires a relative clause which of the wrong syntactic type. Indeed, the pattern in (48) supports the HC in that only Reading B is available:[14]

(48) John knows the price that Fred knows, but I can't remember which
A. # John knows the same price that Fred knows, but I can't remember which price they both know
B. John knows the answer to the question "What price does Fred know?", but I can't remember which price that Fred knows, which John knows

[14]It has been reported to me that some speakers strongly prefer an indefinite in these environments as in (1). The preference for an indefinite is likely due to independent restrictions on the sluicing contexts, and is essentially inconsequential to the analysis.

(1) John knows a price that Fred knows, but I can't remember which
A. # John and Fred both know a price, but I can't remember which price they both know
B. John knows the answer to the question "What is a price that Fred knows?", but I can't remember which price that Fred knows, which John knows

When the two above facts are taken together, they predict that a structure containing both ACD and a sluice should be uninterpretable. I believe that (49) shows that this prediction obtains.

(49) John knows the price that Fred does, but I can't remember which

A. # John knows the same price that Fred knows, but I can't remember which price they both know

B. # John knows the answer to the question "What price does Fred know?", but I can't remember which price it is that John knows Fred knows

The pattern displayed in (49) is consistent with the Headedness Conjecture because if ACD and Sluicing support different readings, then a sentence whose syntax involves both ACD and Sluicing is predicted to be semantically ill-formed. It is interesting, though tangential to the discussion at hand, that the (49) sounds grammatically well-formed, although its meaning, should it have one, is entirely unclear.

To conclude this subsection, we have examined a hitherto unexplored extension of Heim's ambiguity which showed that a propositional attitude interpretation is banned from ACD environments. Two analyses of ACD were compared, and it was argued that the approach which made use of Late Merger could account for the difference in interpretations (Fox, 2002). Lastly, the Headedness Conjecture, which associated each of the two readings in Heim's ambiguity to either the Matching or Raising analysis of relative clauses, was presented and briefly defended with additional evidence. In short, it was proposed that Reading A depends on the availability of an internally headed relative clause (the Raising analysis), whereas Reading B required an externally headed relative clause (the Matching analysis). The respective structures are provided below.

(50) John knows the price that Fred does

A. John knows the same price as Fred

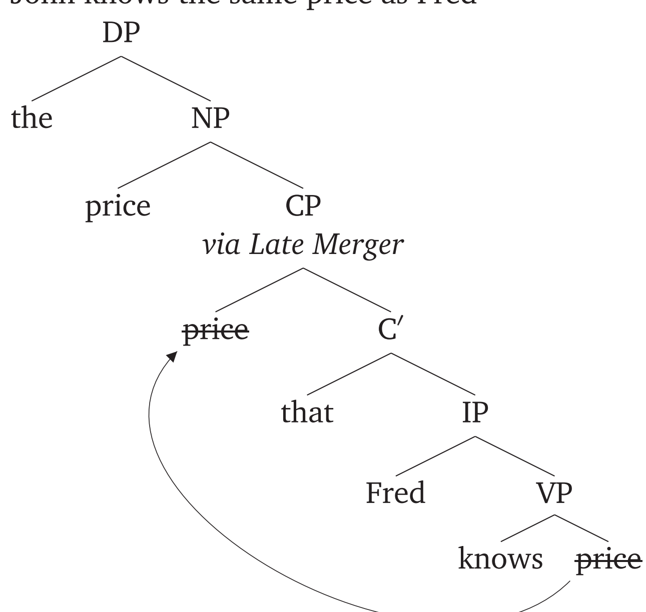

B. John knows which price it is that Fred knows

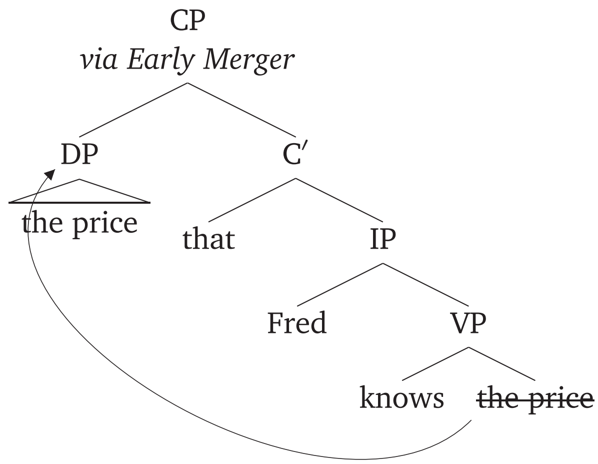

2.1.3 Summary of the CQ data

To review the terrain covered so far, the key characteristics of concealed question interpretations were presented and critically examined. We found evidence supporting the idea that CQs do not behave as individual denoting nouns. In addition to positing that Heim's ambiguity actually reflects two distinct syntactic environments, it was also determined that although CQs often express an identity question that is satisfied by an individual as an answer, there are other coherent readings in which the answer to the identity question can be satisfied in other ways. In particular, we argued that *individual concepts* and *properties of individuals*, can answer the CQ paraphrase just so long as they are unique enough to establish an identity relation. This observation is the guiding insight in much of the next chapter, where we present a novel account of CQs. Before that, however, a survey of previous approaches awaits.

2.2 A Budget of Proposals

In this section, we review three major classes of CQ accounts. First, we examine the view that CQs are individual concepts, the reasons that Heim (1979) rejected this view, and a series of insightful papers by Romero (2004, 2005, 2007a, 2007b) which greatly improves upon this idea (see also a much earlier proposal by Janssen (1984)). Our discussion of Romero's work leads naturally to the second type of CQ account: the view that CQs denote propositions, as recently proposed by Nathan (2006). From here, we return to Heim (1979) for a review of her sketch of CQs as context sensitive relations, recapitulating some of the basic criticisms against this view (Nathan (2006); Romero (2005)). As I hope will be evident, the specific account I advance owes something to each of the works reviewed below.

2.2.1 CQs as individual concepts

In her now classic paper, Heim (1979) develops three possible analyses for the interpretation of concealed questions: CQs are analyzed in turn as *(i)* definite DPs that have been 'quantified-in' over the argument in complement position, *(ii)* individual concepts, and *(iii)* context dependent relations. Heim (1979) ultimately rejects the first two and casts considerable doubt on the third proposal. Yet, I will argue that elements from both of the second and third views contain important insights that form the basis for a sound theory of CQs, and that their dismissal was far too hasty. We will not belabor the presentation with a recapitulation of the first option, as it does not weigh heavily on the present discussion. As Romero has developed an account of CQs as individual concepts which has wide empirical coverage, both Heim's individual concept analysis and her concerns will be presented in some detail.

2.2.1.1 Heim (1979)

In her first analysis, Heim (1979) takes CQs to denote individual concepts, functions from worlds (or more neutrally, indices) to individuals, as discussed in the Introduction. An intensionalized system containing variables ranging over individual concepts is supposed throughout.

Heim (1979) observed that there are two verbs of knowledge in German, *kennen*, whose argument is interpreted as an individual, and *wissen*, whose argument is always a concealed question or proposition.[15] She argues for an ambiguity of English *know* that corresponds to the two readings in other languages. The first homonym variant, $know_1$ is an extensional verb, in the familiar sense of knowing someone. Heim posits a second homophonic variant, $know_{CQ}$, which is not extensional, but not fully intensional either.[16] Instead of selecting for

[15]A similar typology holds in numerous other languages, Germanic and others, including Dutch and Spanish.

[16]Heim's $know_{CQ}$ obeys the following meaning postulate which does not hold of other intensional verbs:

(1) $\exists T \forall x \forall \mathscr{P} \Box(\delta(x,\mathscr{P}) \longleftrightarrow {}^{\vee}\mathscr{P}({}^{\wedge}\lambda y {}^{\vee}T({}^{\vee}x, y)))$

individuals, $know_{CQ}$ composes with individual concepts as complements, and represents a relation between an individual (the bearer of knowledge) and the CQ complement.

(51) Homophonous *know*.

i. $[[know_1]]: \lambda x_e.\lambda y_e.know'(x)(y)$ type $\langle e, \langle e, t\rangle\rangle$

ii. $[[know_{CQ}]]: \lambda x_{\langle s,e\rangle}\lambda y_e.know'(x)(y)$ type $\langle\langle s, e\rangle, \langle e, t\rangle\rangle$

This analysis solves the entailment puzzles in practically the same way as Montague (1970) solves his "temperature paradox."[17]. That is, the relevant statement of identity in the second premise is between variables ranging over individuals not the requisite individual concepts. Using a Russellian definite description, translations are provided beneath the premisses of the invalid syllogism (52).

(52) John knows the capital of Italy
$\exists y[\forall x[\text{capital-of-Italy}'(x) \longleftrightarrow x = y] \wedge K(j, y)]$
The capital of Italy is the largest town in Italy
$\exists y_1[\forall x_1[\text{capital-of-Italy}'(x_1) \longleftrightarrow x_1 = y_1] \wedge$
$\exists y_2[\forall x_2[\text{largest-town-in-Italy}'(x_2) \longleftrightarrow {}^{\vee}y_1 = {}^{\vee}y_2]]$
John knows the largest town in Italy
$\exists y[\forall x[\text{largest-town-in-Italy}'(x) \longleftrightarrow x = y] \wedge K(j, y)]$

Despite its good standing with the entailment puzzle, Heim (1979) locates numerous objectionable features of the individual concept analysis. First, she worries that the intensionalization will overgenerate in the grammar; that is, few common nouns will survive as extensional entities. Second, Heim (1979) believes that individual concepts cannot represent the meta-question reading of (53) discussed in § 2.1.2, without introducing higher order homonyms of price (for example, $\langle s, \langle s, e\rangle\rangle$). Since this process is, in principle, infinitely iterable (54), such a solution would introduce an infinite hierarchy of intensionalized individual concepts.[18]

(53) John knows the price that Fred knows

A. John knows the answer to the question about a particular price, namely that milk costs, *e.g.*, \$1.29, which is the same answer to the same question that Fred knows

B. John knows what the question is to which Fred knows the answer
(John needn't know the answer to this question)

Example (54) has at least four readings to the meagre two in (53). The important reading is the last one, in which the CQ denotation would have three distinct types: $\langle s, e\rangle$, $\langle s, \langle s, e\rangle\rangle$,

[17]See Dowty et al. (1981) and Gamut (1991) for discussion.
[18]The discussion here is based on Nathan (2006).

and $\langle s, \langle s, \langle s, e \rangle \rangle \rangle$. Under this interpretation, Bill knows the meta-question to an answer that Fred knows, and John knows a meta-meta-question to the question Bill knows about Fred.

(54) John knows the price known to Fred that Bill knows

I. John, Bill, and Fred each know the same price, *e.g.*, $1.29

II. John knows the same price as Bill, Bill knows what price Fred knows, and Fred knows some price

III. John knows what price Bill knows, Bill knows the same price as Fred, and Fred knows some price

IV. John knows what price Bill knows, Bill knows what price Fred knows, and Fred knows some price

As noted by Nathan (2006), each higher typed individual concept, under this account, requires a corresponding variant in the denotation of *know*, which selects for the higher typed DP as its complement. For example, in addition to $know_{CQ}$ as presented in (51.ii) above, the semantics now needs the cross-categorical variant in (55.ii) to take a complement of type $\langle s, \langle s, e \rangle \rangle$, the individual concept concept, and the lexical entry (55.iii) to capture the meta-meta-question reading, which is of type $\langle s, \langle s, \langle s, e \rangle \rangle \rangle$:

(55) i. $[[know_{CQ_1}]]$: $\lambda x_{\langle s.e \rangle} \lambda y_e . know'(x)(y)$

ii. $[[know_{CQ_2}]] = \lambda x_{\langle s, \langle s,e \rangle \rangle} . \lambda z_e . \forall w_1 \in Dox_j(w)[x(w_1) = x(w)]$

iii. $[[know_{CQ_3}]] = \lambda x_{\langle s, \langle s, \langle s,e \rangle \rangle \rangle} . \lambda z_e . \forall w_1 \in Dox_j(w)[x(w)(w_1) = x(w)(w_1)]$

from Nathan (2006)

Aside from introducing an aesthetic blemish into the semantics, the third variant of $know_{CQ}$ suggests an unattested reading of (53), namely that Fred knows a meta-question to a question that John knows (56). This reading is simply not available.

(56) John knows$_{CQ_3}$ [the price]$_{meta-question}$ that Fred knows$_{CQ_2}$

Before reviewing previous criticisms of CQs as individual concepts, Romero's (2004, 2005) analysis is introduced as it presents an elegant solution to several of the concerns above, while maintaining CQs as individual concepts.

2.2.1.2 Romero (2004, 2005, 2007a, 2007b)

Romero defends the analysis of CQs as individual concepts while avoiding the pitfalls associated with adding $know_{CQ_3}$ to the lexicon. Following Heim (1979), she notes that a purely extensional denotation is an inadequate CQ interpretation, and that the intensional type of a CQ can originate either:

(57) i. from the *intension* of the NP complement, or
ii. from the *extension* of a higher type NP.

The A and B readings of (53) are the result of forming the appropriate intensional argument in different ways. Summing her insight up nicely, she writes

> the reading A/reading B ambiguity is nothing more than the possibility of drawing an intensional object from the extension or from the intension of the NP. Reading A results when this intensional object corresponds to the *extension* of the NP. Reading B obtains when the intensional object arises from the *intension* of the NP. Romero (2004, 152–153)

A schematic of the division displays the formal and conceptual simplicity of the analysis:

(58) John knows the price that Fred knows
A. $[[know]] + \text{EXT}([_{NP}\text{the price that Fred knows}])$
B. $[[know]] + \text{INT}([_{NP}\text{the price that Fred knows}])$

There are a few ways to implement this analysis formally, but a complete discussion is outside the scope of our current aims. Suffice to say, Romero (2004) considers two such approaches which yield equivalent representations. The first assumes that the CQ combines directly with the intensional verb (59I). The second posits a more traditionally uniform semantics for *know*, which takes a propositional concept as a complement (59II.i), and an intervening ANS operator which lifts the CQ to the correct type (59II.ii).[19]

(59) Romero (2004, 2005)
I. $[[know_{CQ}]] = \lambda y_{\langle s,e\rangle}.\lambda x.\lambda w.\forall w' \in Dox_x(w)[y(w') = y(w)]$
II. i. $[[know_{CQ}]] = \lambda p_{\langle s,\langle s,t\rangle\rangle}.\lambda x_e.\lambda w.\forall w' \in Dox_x(w)[p(w)(w') = 1]$
ii. $\text{ANS}(y_{\langle s,e\rangle}) = \lambda w.\lambda w'.y_{\langle s,e\rangle}(w') = y_{\langle s,e\rangle}(w)$

When applied to an intensional verb, the CQ complement is representationally identical under either approach. Romero (2004) argues that the second approach can account for both the anaphora and coordination facts discussed in §2.1.1. That the CQ does not agree with

[19] See Beck and Rullman (1999) and Beck and Sharvit (2002) for analyses of questions that propose such an operator directly into the semantics.

the gender of its referent follows from the fact that pronominalization of the CQ argument is just like pronominalization of a proposition. Propositions are not individuals and as such do not require agreement in grammatical gender in English. That no CQ reading is available for coordinated structures is a testament to the ontological disparity of the arguments; that is, CQs are of a higher type than the type with which individual selecting verbs combine. Given the option, the argument must be interpreted at an appropriate level, *i.e.*, an individual.

To illustrate the combinatorics of Romero's second implementation, derivations of the A and B readings are provided below.

(60) Reading A from Romero (2004)

a. $[[NP]]_g$(w):
$\iota x_{\langle s,e\rangle}[price(x,w) \wedge \forall w''' \in Dox_{fred}(w)[x(w''') = x(w)]]$

b. ANS($[[NP]]_g$(w)):
$\lambda w^*.\lambda w'.\iota x_{\langle s,e\rangle}[price(x,w) \wedge \forall w''' \in Dox_{fred}(w)[x(w''') = x(w)]](w') =$
$\iota x_{\langle s,e\rangle}[price(x,w) \wedge \forall w''' \in Dox_{fred}(w)[x(w''') = x(w)]](w^*)$

c. know(ANS($[[NP]]_g$(w))):
$\lambda w.\forall w' \in Dox_j(w)\iota x_{\langle s,e\rangle}[price(x,w) \wedge \forall w''' \in Dox_{fred}(w)[x(w''') = x(w)]](w') =$
$\iota x_{\langle s,e\rangle}[price(x,w) \wedge \forall w''' \in Dox_{fred}(w)[x(w''') = x(w)]](w)$

(61) Reading B from Romero (2004)

a. $[[NP]]_g$:
$\lambda w''.\iota x_{\langle s,e\rangle}[price(x,w'') \wedge \forall w''' \in Dox_{fred}(w)[x(w''') = x(w)]]$

b. ANS($[[NP]]_g$):
$\lambda w.\lambda w'.[\lambda w''.\iota x_{\langle s,e\rangle}[price(x,w'') \wedge \forall w''' \in Dox_{fred}(w)[x(w''') = x(w)]]](w') =$
$[\lambda w''.\iota x_{\langle s,e\rangle}[price(x,w'') \wedge \forall w''' \in Dox_{fred}(w)[x(w''') = x(w)]]](w)$

c. λ-conversion:
$\lambda w.\lambda w'.[\iota x_{\langle s,e\rangle}[price(x,w') \wedge \forall w''' \in Dox_{fred}(w)[x(w''') = x(w')]] =$
$\iota x_{\langle s,e\rangle}[price(x,w) \wedge \forall w''' \in Dox_{fred}(w)[x(w''') = x(w)]]$

d. know(ANS($[[NP]]_g$(w))):
$\lambda w.\forall w' \in Dox_j(w)[\iota x_{\langle s,e\rangle}[price(x,w') \wedge \forall w''' \in Dox_{fred}(w)[x(w''') = x(w')]] =$
$\iota x_{\langle s,e\rangle}[price(x,w) \wedge \forall w''' \in Dox_{fred}(w)[x(w''') = x(w)]]]$

The advantages of this analysis are clear: the interpretations are completely compositional, and the different readings are derived from whether or not the ANS operator selects for an argument of type $\langle s,e\rangle$ or $\langle s,\langle s,e\rangle\rangle$, that is, the *extension* (reading A) or the *intension* (reading B) of an individual concept denoting DP, respectively.

Despite its advantages, the above analysis does not address several key concerns regarding the representational and distributional nature of individual concepts. First, the account does not address Heim's original concern that about the growing hierarchy of types. Second, the two readings are achieved by assuming that the ANS operator selects for two (cross-categorical) types of argument; presumably, the ambiguity between individual, individual concept and individual concept concept readings is reduced to lexical selection. Thus, to achieve a CQ interpretation, the verb must select for an ANS operator, which in turn is a propositional concept, of type $\langle s, \langle s, t \rangle \rangle$. Lastly, Nathan (2006) presents three main arguments against the individual concept account in general, the last of which constitutes a major complication in ther analysis of CQs for all accounts.

Nathan's first criticism of the individual concept approach rests on problematic coordination facts between individual concept selecting verbs (*rise, fall*) and CQ selecting verbs (*know, forget*).

(62) a. [The price of milk]$_{IC}$ fell last week and is rising this week
b. [The price of milk]$_{CQ}$ is known to John and has been forgotten by Fred
c. * [The price of milk] fell last week and is known to John

From (62c), Nathan (2006) argues that individual concepts and CQs diverge in certain environments, and that, as a result, CQs do not denote individual concepts. However, it is unclear whether these examples are ungrammatical for the sole reason of argument selection mismatch. Slight modification permits a perfectly acceptable reading with neuter anaphora:

(63) John thinks he knows the price of milk. But it fell last week when he was on vacation.

In the first conjunct, *the price of milk* carries an indisputable concealed question interpretation. As the second conjunct, according to Nathan's argument, supports only an individual concept, it is surprising that *it*, which is anaphorically related to the denotation of the CQ (see §2.1.1 for discussion), is felicitous where only individual concepts are permitted. At the very least, that examples such as (63) are generated so easily should cast some suspicion on the insight (62) purports to afford.

The second concern in Nathan (2006) centers around the compositional construction of more complex CQs, such as *the price of milk*, from individual concept denoting nouns, *(the) price*. In particular, if both *the price* and *the price of milk* are of type $\langle s, e \rangle$, then the modifier does not change the type, but only the semantic content.

Yet, it is clear that this concern is, at best, very minor. In fact, constructions with modifiers tend *not* to change the category of object they modify. Indeed, returning an identical type is a common property of modification and is seen in other constructions involving the addition of an *of*-phrase, in addition to prepositional and adjectival modifiers.

(64) a. the leg of the table
b. the cat on the table
c. the fluffy cat on the table

Each noun in (64) has maintained its ontological type, despite modification, which is all to say that a fluffy cat on a table is still a cat, no matter the extent of its mane or where it perches. Interestingly, (64b–c) immediately bring to mind plausible CQ interpretations, whereas (64a) does not. Without additional context, the most natural question paraphrase is one which selects a candidate table leg or cat from a set of relevant possibilities. Presumably, (64a) cannot not *automatically* be interpreted as index dependent.

(65) a. # John knew/guessed/predicted the leg of the table

b. John knew/guessed/predicted the cat on the table

c. John knew/guessed/predicted the fluffy cat on the table

Third among Nathan's major criticisms is that certain nouns which are plausibly individual concepts do not straightforwardly pattern as CQs.

(66) a. John knows the picture on Jordan's wall

b. $\neq$ John knows what the picture on Jordan's wall is

c. $\approx$ John knows which picture of Jordan's is on her wall

Under most accounts equating CQs with individual concepts, this distribution goes unexplained. In the next chapter, it is argued that while CQs are indeed individual concepts, they are subject to a very stringent pragmatic constraint which requires that only sets of concepts which uniquely cover the domain of individuals can be answers to CQs. While we do not attempt to solve this puzzle in the framework developed below, it is suggested that CQ NPs must be (conventionally) unique enough to maintain a bijection between the CQ NPs and answers to their question paraphrases. We now discuss Nathan's proposal which connects CQ denotations to a propositional theory of questions.

2.2.2 CQs as identity propositions

As discussed above, Nathan (2006) enumerates several difficulties for the CQ as individual concept analysis. The most damaging of which was the lack of explanation for the distribution of CQ interpretations: individual concepts appear where CQs do not, and *visa versa*. As it stands, the individual concept account does not clearly indicate which nouns make good CQs and under what conditions. The problem is accentuated by a discrepancy between the great many nouns can be individual concepts, and the few that might be interpreted as CQs.

Nathan (2006) argues that individual denotations cannot adequately explain the selection of CQ interpretations alone. Nathan (2006) opts to incorporate individual concept meanings into the construction of an identity proposition via a series of type shifting rules.

The idea that CQs express propositions has some intuitive currency, as most verbs that permit a CQ reading also allow a proposition. This correlation is summarized below:

(67) Propositional/Concealed Question Correlation (PCQC): A CQ can fill a predicate's argument position iff

i. the Case requirements of the proposition are met, and

ii. a (clausal) question can fill the position, and

iii. a (clausal) proposition can fill the position.

In Nathan (2006), CQ interpretations are ultimately represented as propositions, as in (68). Our canonical examples are shown in (69), in pseudo-logical form.

(68) $[\![\text{the NP}]\!] = \iota p_{\langle s,t\rangle}.[[\exists x_e.p = \lambda w'.\ [\![\text{NP}]\!](x)(w')] \wedge C(p)]$

(69) a. John guessed the mayor of the town

b. John guessed $\exists x.p = \lambda w'.x$ is a mayor of the town in $w' \wedge C(p)$

(70) a. John predicted the cost of the reform

b. John predicted $\exists x.p = \lambda w'.x$ is a cost of the reform in $w' \wedge C(p)$

The interpretation of CQ propositions involves two parts; (69a), for example, is interpreted as (i) *John guessed for some individual x, the proposition which is the set of worlds, such that* x is a mayor of the town *in those worlds*, (ii) *such that this proposition meets a contextually specified restriction*.[20]

There is a systematic difference between relational nouns and non-relational nouns in the distribution of CQ interpretations (see discussion on relational nouns in the Introduction). Given the appropriate verb, relational nouns almost always permit a CQ reading, whereas non-relational nouns require a particular type of context, one that arguably transforms them into objects which are more like relational nouns in crucial respects.

Nathan (2006) captures the dichotomy by means of stipulating the paths that different nouns must travel in order to achieve the CQ interpretation. Among the battery of type-shifting rules he proposes, the main rule shifts the intension of a relational noun, $\langle s, \langle e, \langle e, t\rangle\rangle\rangle$,[21] directly to a relation between individuals and sets of propositions $\langle e, \langle\langle s, t\rangle, t\rangle\rangle$, as in ($\mathscr{S}$).

($\mathscr{S}$) $\lambda P_{\langle s,\langle e,\langle e,t\rangle\rangle\rangle}.\lambda y_e.\lambda p_{\langle s,t\rangle}.[\exists x_e.p = \lambda w'.P(w')(y)(x)]$

At this point, a relational NP (*mayor*) can either combine with an individual denoting argument, such as *the town* in (71) or have its second argument closed off via existential closure, as in (72).

[20] I assume that (69) should be read as (1), in which the existential quantifier binds the variable x not within the proposition p, but rather into the λ-term. Nonetheless, it unclear whether the binding relationship is opaque.

(1) $[\![\text{the NP}]\!] = \iota p_{\langle s,t\rangle}.[[\exists x_e.[p = \lambda w'.\ [\![\text{NP}]\!](x)(w')] \wedge C(p)]]$

[21] We will assume a different type of analysis for relational NPs in the next chapter, following Dekker (1993a)

(71) RELATIONAL NOUNS: *Combination with individual argument*

a. [[mayor]]
$\lambda y_e.\lambda x_e$[x is mayor of y in w]

b. lift([[mayor]])
$\lambda w_1.\lambda y_e.\lambda x_e$[x is mayor of y in w_1]

c. $\mathscr{S}$(lift([[mayor]]))
$\lambda y_e.\lambda p_{\langle s,t\rangle}.[\exists x_e.p = \lambda w_1[\lambda w_1.\lambda y_e.\lambda x_e[\text{ x is mayor of y in } w_1](w_1)(y)(x)]] =$
$\lambda y_e.\lambda p_{\langle s,t\rangle}.[\exists x_e.p = [\lambda w_1.\text{ x is mayor of y in } w_1]]$

d. ($\mathscr{S}$(lift([[mayor]])))([[the town]])
$\lambda p_{\langle s,t\rangle}[\exists x_e.p = [\lambda w_1 \text{x is mayor of [the town] in } w_1]]$

e. [[the]]((($\mathscr{S}$(lift([[mayor]])))([[the town]])))
$\iota p_{\langle s,t\rangle}.[\exists x_e.p = [\lambda w_1 \text{x is mayor of [the town] in } w_1] \wedge C(p)]$

(72) RELATIONAL NOUNS: *Existential closure of argument*

d'. ∃-closure
$\exists y_e \lambda p_{\langle s,t\rangle}[\exists x_e.p = [\lambda w_1 \text{x is mayor of y in } w_1]]$

e'. [[the]]($\mathscr{S}$(lift([[mayor]])))
$\iota p_{\langle s,t\rangle}.\exists y_e[\exists x_e.p = [\lambda w_1 \text{x is mayor of y in } w_1] \wedge C(p)]$

Non-relational nouns, however, follow a markedly different route; Nathan (2006) argues that their type denotation is first lifted to the type of a relational noun. Nathan (2006) observes that the sentences in (73), require different degrees of contextual embellishment in order to license a CQ interpretation.

(73) a. # John knows UvA's philosopher

b. John knows UvA's most popular philosopher

c. John knows the philosopher who adores Sellars

As Nathan (2006) observes, (73a) is infelicitous when uttered without context, whereas both (73b – 73c) require relatively little context and are easily interpreted as CQs. In particular, (73a) requires a context that satisfies the presupposition that UvA has more than one philosopher, whereas (73b–c) appear to be sufficiently specific. If the discourse participants share the belief that UvA has only one philosopher, thereby satisfying the uniqueness presuppositions of (73a) becomes as felicitous as the others. Nathan (2006) attributes the considerable improvement in such a context to the notion that the NP *UvA's philosopher* acts here as a relational noun. However, he does not specify whether there is an operation from non-relational nouns of type $\langle e,t\rangle$, or if he intends the analogy to relational nouns to be somehow pragmatically determined. Either option is problematic for his theory, as there would be no need for the operation $\mathscr{T}$ on the one hand, and a possible counterargument for a wholly semantic treatment of CQs on the other.

In any event, Nathan (2006) posits a 'cost-free' type-shifting operation, which instead of the usual predicate modification operation, assigns a propositional meaning to the conjunction of the noun and the modifier. As expected, the type shifter takes (the intension of) the modifier and returns a function from (the intension of) non-relational noun type to a set of propositions:

($\mathscr{T}$) $\lambda Q_{\langle s,\langle e,t\rangle\rangle}\lambda P_{\langle s,\langle e,t\rangle\rangle}.\lambda p_{\langle s,t\rangle}.[\exists x_e.[p = \lambda w_1.P(w_1)(x) \wedge Q(w_1)(x)]]$

Once the operation $\mathscr{T}$ composes with a non-relational head noun, *e.g., the philosopher*, the propositional meaning emerges and is input for $\mathscr{S}$. A sample derivation, ignoring intensions for readability, is given below:

(74) NON-RELATIONAL NOUNS

a. [[who adores Sellars]]
λx_e.x adores Sellars

b. $\mathscr{T}$([[who adores Sellars]])
$\lambda P_{\langle e,t\rangle}.\lambda p_{\langle s,t\rangle}.[\exists x.[p = \lambda w_1.P(w_1)(x) \wedge \text{x adores Sellars}]]$

c. ($\mathscr{T}$ [[who adores Sellars]])([[philosopher]])
$\lambda p_{\langle s,t\rangle}.[\exists x.[p = \text{philosopher(x)} \wedge \text{x adores Sellars}]]$

d. [[the]](($\mathscr{T}$ [[who adores Sellars]])([[philosopher]]))
$\iota p_{\langle s,t\rangle}.[\exists x.[p = \text{philosopher(x)} \wedge \text{x adores Sellars}](p)] \wedge C_{\langle\langle s,t\rangle,t\rangle}(p)])$

At this point, both relational and non-relational nouns have been encoded as 'propositional nouns', denoting objects of type $\langle s,t\rangle$, and as such may compose with predicates taking propositions. We take the output from the type-shifter $\mathscr{S}$ with the relational noun in (71) above for illustration:

(75) John knows the mayor of the town (m.o.t.)

a. [[know]]
$\lambda p_{\langle s,t\rangle}.\lambda z_e.\forall w_2 \in Dox_z(w)[p(w_2) = 1]$

b. [[know]](71)
$\lambda z_e.\forall w_2 \in Dox_z(w).[\iota p_{\langle s,t\rangle}[\exists x_e.[p = [\lambda w_1.\text{x is m.o.t. in } w_1] \wedge C(p)]](w_2) = 1]$

c. ([[know]](71))([[john]])
$\forall w_2 \in Dox_{john}(w).[\iota p_{\langle s,t\rangle}[\exists x_e.[p = [\lambda w_1.\text{x is m.o.t. in } w_1] \wedge C(p)]](w_2) = 1]$

Sparing the reader a complete derivation, reading A follows the computation detailed above, with the caveat that the argument to *price* has been existentially closed, as in (72). The resulting representation is displayed in (76):

(76) Reading A

a. John knows the (same) price that Fred knows

b. $\forall w_3 \in Dox_{john}.[\iota p.[\forall w_1 \in Dox_{fred}(w)[p(w_1)=1] \wedge$
$\exists y.\exists x.p = [\lambda w_2.y \text{ costs x at } w_2] \wedge C(p)(w_3)=1]]$

In the calculation of reading A, Nathan's analysis has followed Romero (2004) closely, although perhaps more in spirit than technical detail. However, in order to relate reading B to reading A, Nathan (2006) proposes a type-shifting rule that "changes the set of propositions derived by $[\mathscr{S}]$ into a new set of propositions that has an extra propositional layer added to it," as in (AB).

(AB) $\lambda P_{\langle s,\langle\langle s,t\rangle,t\rangle\rangle}.\lambda p_{\langle s,t\rangle}.[\exists q_{\langle s,t\rangle}.p = \lambda w_1.P(w_1)(q)]$

Applying *the price that Fred knows* from (76b) to the type shifter AB, yields (77B), forming the argument of *know*, as in (77c).

(77) a. Intension of *price Fred knows*
$\lambda w'''.\lambda r_{\langle s,t\rangle}.$
$[\forall w' \in Dox_{fred}(w''')[r(w')=1] \wedge \exists y.\exists x.r = [\lambda w''. \text{ y costs x in } w'']]$

b. Applied to AB, after application of [[the]]
$\lambda p_{\langle s,t\rangle}.\exists q_{\langle s,t\rangle}.p = \lambda w''''.$
$[\forall w' \in Dox_{fred}(w'''')[q(w')=1] \wedge \exists y.\exists x.q = [\lambda w''. \text{ y costs x in } w'']]$

c. John knows the answer to question "What price does Fred know?"
$\forall w''' \in Dox_{john}(w).[\iota p_{\langle s,t\rangle}[C(p) \wedge \exists q_{\langle s,t\rangle}.p = \lambda w''''.$
$[\forall w' \in Dox_{fred}(w'''')[q(w')=1] \wedge \exists y.\exists x.q = [\lambda w''. \text{ y costs x at } w'']]](w''')=1]$

The representation (77c) simply states that all of John's doxastic alternatives contain the proposition that Fred knows some price, or, in other words, that John knows *the proposition* that Fred knows a price. It is argued in the next chapter that Nathan's intuition about the propositional nature of knowledge expressed in reading B is essentially correct. For it is clear that reading B does not require that John know anything more than the fact *that* Fred knows some price. However, we advocate a markedly different route to obtaining this reading, without a specialized type shifting rule fashioned to suit the purpose.

To briefly summarize Nathan's (2006) proposal, CQs are thought to be propositional entities on the basis of a correspondence between verbs that select for propositions and verbs that support CQ interpretations from DPs. Two classes of relational noun, abstract and concrete, were differentiated by means of the naturalness with which they yield question paraphrases and the compositional mechanics posited to derive these readings. As discussed above, the difference between concrete and relational nouns appears to have grammatical import, in that concrete relational nouns are prohibited in environments where abstract relational nouns are not. Lastly, Nathan (2006) provides a compositional semantics that relates the two readings of Heim's ambiguity via a type shifting operation (AB) which adds an additional propositional layer to the default CQ denotation to yield Reading B.

While a contextual variable C is said to select the most salient proposition from the context, it is unclear how it does so, and what ramifications the role of this contextual dependence plays in the interpretation of CQs. In our next section, we turn to a theory which makes crucial use of context in determining the relation between the CQs and the salient properties which define them.

2.2.3 CQs as context dependent relations

The last section of our review of the literature on concealed questions, brings us back around to the final, and somewhat tentative, analysis presented in Heim (1979). Here, Heim modifies her individual concept account so that $know_{CQ}$ is no longer a relation between individuals and *individual concepts*, but instead a relation between individuals and *second order properties of individual concepts* which depends on some contextually salient property relating the DP of the CQ to the context in a variety of ways. The leading intuition behind this approach is that when *John knows* α, John's knowledge consists of *knowing* α *as* α; Heim (1979, 58) lists a few possible paraphrases for (78) below.

(78) John knows Bill's phone number

i. John knows Bill's phone number *as* Bill's phone number

ii. John knows Bill's phone number *with respect to* its being Bill's phone number

Intuitively, then, the relation between the content of John's knowledge, and what it represents under John's perspective is of crucial importance. In particular, a CQ reading of a DP is obtained, in part, when the denotation of the DP in the actual world is related to an object of knowledge, belief, inquiry, *etc.* under an appropriate perspective. In Heim (1979), this relation is contextually supplied, as are the assignment functions valuing any free variables contained in the intensional complement.

More formally, $know_{CQ}$ is still homophonous with its extensional cousin, but it now is taken to be an IL-constant KNOW(P), of type $\langle\langle s,\langle\langle s,e\rangle,t\rangle\rangle,f(TV)\rangle$, where $f(TV)$ is the type of a transitive verb $\langle e,\langle e,t\rangle\rangle$. Heim (1979) proposes that this constant, KNOW(P), allows for a variety of translations, just so long as the variable P: *(i)* is of the correct type, $\langle s,\langle\langle s,e\rangle,t\rangle\rangle$, and *(ii)* is made up of only variables. Variable assignment functions, g_C, are provided by the context C. That is, within P is a free variable which is subject to specification from the particular context at hand. The value of the variable in P will be determined by an assignment function g_C made available in context C. We return to the way the assignment function works shortly.

In her translations of the constant KNOW(P), Heim (1979) requires that the following meaning postulate apply to KNOW(P):

(MP1) $\exists S \forall x \forall \mathscr{P} \Box[\delta(x, \mathscr{P}) \leftrightarrow {}^{\vee}\mathscr{P}({}^{\wedge}\lambda y {}^{\vee}S({}^{\wedge}x, y))]$, for

i. x, y of type e,

ii. S of type $\langle s, \langle\langle s, e\rangle, \langle e, t\rangle\rangle\rangle$, and

iii. $\mathscr{P}$ of type $\langle s, \langle\langle s, \langle e, t\rangle\rangle, t\rangle\rangle$.

Gamut (1991) discuss a nearly identical meaning postulate MP2. The only difference is that the first argument of S is an *individual concept* in Heim's approach (MP1), while S takes an individual simpliciter in (MP2).

(MP2) $\exists S \forall x \forall \mathscr{P} \Box[\delta(x, \mathscr{P}) \leftrightarrow {}^{\vee}\mathscr{P}({}^{\wedge}\lambda y {}^{\vee}S(x, y))]$, for

i. x, y of type e,

ii. S of type $\langle s, \langle e, \langle e, t\rangle\rangle\rangle$, and

iii. $\mathscr{P}$ of type $\langle s, \langle\langle s, \langle e, t\rangle\rangle, t\rangle\rangle$.

On MP2, Gamut (1991) write, modulo names of variable constants:

> MP2 expresses that for each δ for which it is defined, there is a relation S between individuals such that $\delta(x, \mathscr{P})$ is true iff $\lambda y {}^{\vee}S(x, y)$, or in other words, the property of standing in the relation S to x belongs to the set of properties ${}^{\vee}\mathscr{P}$, *viz.* ${}^{\vee}\mathscr{P}(\lambda y {}^{\vee}S(x, y))$ ... [F]or every δ for which MP2 is defined, there is exactly one S that fulfills this condition.

By analogy to the explanation above, we see that MP1 requires that for a relation δ between individuals x and second order properties of individual concepts $\mathscr{P}$ holds just when $\mathscr{P}$ contains the property of individuals y standing in a particular relation S to x.

Heim assumes that the following convention applies to KNOW(P), reducing it to an extensional relation, $[\text{KNOW(P)}]_*$. We generalize this convention in (79i):

(79) CONVENTION:
If δ is an expression of type $\langle\langle s, \langle\langle s, e\rangle, t\rangle\rangle, \langle e, \langle e, t\rangle\rangle\rangle$, we may write δ_* instead of δ.

In essence, convention (79) allows us to replace a relation of knowing an individual by a unique second-order property with a relation of knowing that individual simpliciter. Note that not all predicate relations of the form $\delta(x, \mathscr{P})$ can be extensionalized to a δ_* equivalent, *e.g.* if the the property $\mathscr{P}$ does not refer to a second order relation related to a specific individual. Thus, δ_* substitutes for δ when the relation δ holds between an individual concept ${}^{\wedge}x$ and a second order property $\mathscr{P}$ holding of a specific individual y. As Gamut (1991, 176) write, this $\mathscr{P}$ can be conceived of as "the second order property that refers in every world to the set of properties of y (in that world)."[22]

To illustrate Heim's third solution, we turn to its treatment of the entailment puzzle in (80), repeated below for convenience:

[22] In addition, we may need a meaning postulate that will ensure *rigidity* for names, *i.e.* that a name m referring to some individual **m** will refer to the same individual in all possible worlds, since we don't want

(80) I. John knows the capital of Italy
$\exists v[\forall u[\text{capital-of-Italy}_*(u) \longleftrightarrow u = v] \wedge [KNOW(P)_*(j,v)]$

II. The capital of Italy is the largest town in Italy
$\underline{\exists v[\forall u[\text{capital-of-Italy}_*(u) \longleftrightarrow u = v] \wedge \forall t[\text{largest-town-in-Italy}_*(t) \longleftrightarrow t = v]]}$

III. John knows the largest town in Italy
$\exists v[\forall u[\text{largest-town-in-Italy}_*(u) \longleftrightarrow u = v] \wedge [KNOW(P)_*(j,v)]$

As Heim (1979) notes, the conclusion in (80) holds whenever the referential index (world, information state, *etc.*) and variable assignment function g_{C^I} remains constant in the context C^I of the first premise. To account for the (more common) perception that (80) is, in fact, invalid, Heim proposes that uttering the second premise (80II) creates a strong pragmatic bias to create a new context C^{II}, which contains another variable assignment $g_{C^{II}}$ which differs just on the property P which is assigned to the object denoted by the DP.

If this bias is operative, then the second premise creates a context C^{II} which renders the syllogism invalid, as the property P is assigned two separate properties in C^I and C^{II}. In C^I, it is the property of being *being the capital of Italy*, while in C^{II} it is *being the largest town in Italy*.

Given that the contexts may differ between premise and conclusion, Heim (1979) correctly anticipates that there is a potential confound for genuinely valid entailments, as in (81). Specifically, there is no guarantee that the contexts should remain the same from the general premise, that John knows every phone number, to the specific conclusion, that John knows one individual's number. Yet, our intuitions demand that (81II) follow from (81I), at least on the most readily available reading.

(81) I. John knows every phone number
II. John knows Bill's phone number

To amend this problem, Heim (1979) proposes that KNOW(P) alternates between translations, as discussed above. In the family of translations, one possibility is the relation-in-intension $\lambda w.\lambda y.R(y,x)$. Here, the relation R (of type $\langle\langle s,e\rangle,\langle\langle s,\langle s,e\rangle,t\rangle\rangle\rangle$) is a variable to be specified by the context according to its salience. In the case of (81), it expresses *being-the-phone-number-of* in both premise and conclusion, necessitating a valid entailment:

$seek'_*(j,m) \equiv seek'(j,{}^\wedge\lambda X^\vee X(m)))$ to mean that John seeks some individual that has all the properties of Mary but which m does not refer to Mary in the world in which he finds $\mathbf{m}'$, the referent of m in that world.

(1) $\exists x\Box(x = \alpha)$, where α may be any name in our language.

(82) I. John knows every phone number
$\underline{\forall x[\text{phone-number(x)} \rightarrow [KNOW(\lambda w.\lambda y.R(y,x))_*(j,x(w))]]}$

II. John knows Bill's phone number
$\exists y[\forall x[of'(\lambda w.b^*(w))(\lambda w'.\text{phone-number}(x)(w') \longleftrightarrow x = y)] \wedge$
$[KNOW(\lambda w''.\lambda z.R(z,y)))]_*(j,y(w))]$

Heim (1979) herself expressed "... uneasiness about the extent to which this theory relies on pragmatic explanations wherever its semantic machinery fails to pick out the intuitively good readings of an example among a host of intuitively bad ones." Perhaps reservation about the formal parsimony of the approach is well-founded; indeed, the pragmatic mechanisms which are intended to "cut down on a vast overgeneration of semantically admissible readings" prove unconstrained themselves. However, I do not think that Heim's doubts should completly dissuade us from pursuing at least a partially pragmatic treatment of concealed questions. In the next chapter, it is argued that a full treatment of CQs requires a formalization of the ways in which CQs are context dependent, and that the pragmatic aspects which govern the interpretation of CQs in various contexts are natural and are definable along intuitive and sufficiently systematic lines.

However, problems with this approach are not limited to the merely conceptual. In her interpretation of Heim (1979), Romero (2005) notes that the context dependence analysis fails to meet two main empirical difficulties. Firstly, the account does not answer why only *some* NPs are candidates for CQ interpretations. In particular, why the individual denoting name in (83) fails to be interpreted as *John knows Rome as P*, where *P* is a contextually salient property.

(83) #John knows Rome

The second criticism is far more challenging; under the context dependence approach, it is unclear how to capture reading B in Heim's ambiguity puzzle. The translation of (84a) is given as (84b), which, for simplicity of representation, employs the iota operator as the definite description over the longer Russellian statement. Letting P be the property corresponding to knowing *the answer to the question 'what price does Fred know?'* and Q be the property of *being the price of milk*, John's knowledge consists of a price which is extensionally equivalent to the price that Fred knows. Even if P and Q pick out the same object, say \$1.79, the problem is that John needn't know that Fred knows \$1.79 *as the price of milk*. Romero claims that this does not adequately represent reading B.

(84) a. John knows the price that Fred knows
b. $\lambda w.\text{know}(j, \iota x_e[\text{price}(x,w) \wedge \text{know}(f,x,Q,w)], P, w)$

In the next chapter, much space is devoted to presenting an analysis which incorporates both the context sensitivity of Heim (1979) and the individual concept approach of Romero. While I agree with Romero that Heim's context dependence account does not provide an explanation for why proper names cannot be CQs, and that the second reading of Heim's ambiguity puzzle goes unaccounted for in this approach, it is far from clear that these deficiencies are essential failings of Heim's analysis. The account of the next chapter argues that (i) names, such as Rome, cannot be interpreted as a CQ for pragmatic reasons and (ii) that Reading B is a proposition, not a CQ complement. In theory, then, we must re-interpret the extent to which Romero's criticisms of Heim's approach hold, despite advocating a markedly different approach.

2.2.4 Summary of Literature Review

We have reviewed three major approaches to concealed questions: Heim (1979) who argued that CQs are context dependent relations involving a family of IL-variables; Romero's string of articles (2004, 2005, 2007a, 2007b) suggesting that CQs are individual concepts in identity relations, and Nathan (2006) who suggests that CQs are in fact propositions to exploit the fact that most CQ-selecting verbs also select for propositions. In the theory developed in the next chapter, I take something from each one of these proposals, but synthesized in a completely novel fashion. Concealed questions are argued to be highly context dependent, but also very tightly constrained, relations between individual concepts and identity questions.

3 Towards a Unified Analysis of CQs

The present chapter presents a novel analysis of Concealed Questions (CQs) that explicitly addresses several of the puzzles presented in the previous chapter. The basic conception of CQs adopted here is that CQs express relations between individual concepts and identity questions, such that the individual concept functions as a constituent *answer* to the question of the identity of the NP which constitutes the CQ. Thus, CQs are argued to be semantic objects of type $\langle\langle s,e\rangle,\langle s,\langle s,t\rangle\rangle\rangle$. A specific type shifting rule $\mathscr{Q}$ is posited which lifts the nominal complement of type $\langle s,e\rangle$ to the CQ type $\langle\langle s,e\rangle,\langle s,\langle s,t\rangle\rangle\rangle$. In effect, the individual concept in the argument of the result type of the CQ *specifies* the extension of the DP at the relevant world of evaluation. In an attempt to clarify this notion, two different – but inter-definable – notions of $\mathscr{Q}$ are explained in detail in §3.3.1.

This idea borrows from several sources, most crucially from Aloni's (2001) theory of Conceptual Covers, which places pragmatic constraints on how individuals are identified in both identity questions and belief reports. This account is useful to the discussion of CQs because it allows us to characterize what sort of answer can satisfy the concealed question in a natural and intuitive way. A brief introduction to Conceptual Covers is provided in §3.2.

Additionally, this account of CQs incorporates a partially dynamic notion of meaning. Crucial to the analysis of Heim's ambiguity (Heim, 1979), Dekker's (1993a, 1993b) dynamic semantics for relational nouns is adopted wholesale. It is argued that the formal system presented there allows us to account for the subtle differences in interpretation which may have been overlooked in the existing literature on CQs, including *why* and *how* CQs are context-sensitive to the extent that they are.

After reviewing the necessary background and presenting an intuitive and formal sketch of the account, the proposal is evaluated in light of how well it performs on the major puzzles.

I show that it not only conforms to the defining characteristics of CQs, but that it gives some additional insight into novel data as well. For ease of reference, the central claims of the proposal are listed in (85).

(85) CENTRAL TENETS OF A PRAGMATIC CQ ACCOUNT:

I. Syntactically, CQs are DPs.

II. Semantically, CQs express relations between individuating concepts from sets of conceptual covers and questions involving the identity of the DP from which the CQ interpretation stems.

III. CQs are thus of type $\langle\langle s,e\rangle,\langle s,\langle s,t\rangle\rangle\rangle$ and are derived from nouns denoting individual concepts via a type shifting operation $\mathscr{Q}$.

Before moving onto an explicitly formal account, I discuss some intuitive motivations for incorporating Conceptual Covers into the analysis directly below. Here, I hope to convince the reader that CQs display some of the same properties as both descriptions under belief reports and identity questions.

3.1 The Pragmatics of CQs

Given that canonical CQ typically involve *definite descriptions*, and are paraphrasable as an *identity question*, it is somewhat unsurprising that CQs are subject to exactly the same sort of philosophical puzzles and semantic confounds as those that litter the expansive literature on definite descriptions and identity questions. Yet, the connection is not as firmly established in existing analyses as one might imagine. The topics of definite descriptions and informative identity have both generated an impressive literature that constitutes some of the most familiar and important work in the last century of philosophy. A substantive discussion of the debate, let alone any sort of balanced treatment, is well-beyond the scope of the present study. Instead, a few examples are offered as evidence towards a connection between the long standing philosophical debates and some of the varieties of CQ interpretation discussed above. To this limited end, the next section reconsiders Heim's entailment problem in light of familiar transparency paradoxes. The subsequent section examines how identity in the CQ is subject to similar puzzles of identity in its question paraphrase.

3.1.1 The Entailment Problem

Although Heim's entailment problem will be recalled by the reader from §2.1.1, the original example is reprinted below, omitting any specific logical rendering.

(86) I. John knows the capital of Italy
II. The capital of Italy is the largest town in Italy
III. John knows the largest town in Italy

Much has been made about the correspondence to Montague's (1970) solution to the "temperature paradox", which was treated by identity of individual concepts (for a cogent discussion see Romero (2007b)). Importantly, Heim (1979) observed that although the syllogism (86) is not *valid*, per se, it does hold just in case the second premise which establishes identity between the two terms is known to be true by John in all of his accessible states of knowledge.

Strikingly similar cases have been thoroughly discussed by Quine (1956), and many others, in so-called *referentially opaque contexts*. These contexts violate *Leibniz's law of the indiscernibility of identicals*,[1] in which co-referring terms and expressions can be exchanged without altering the truth value of the expression in which those terms are contained. There are several types of contexts, sometimes known as transparency paradoxes, which appear to be referentially opaque. Two such contexts are (i) quotation and (ii) belief contexts. For illustration, consider the invalid syllogism in (87). The equivalence in premise II cannot, as it were, ensure that the terms *Cicero* and *Tully* are interchangeable in quotation contexts (demarcated by the corner quotes).

(87) I. ⌜Cicero⌝ contains six letters
II. Cicero is Tully
III. ⌜Tully⌝ contains six letters

The second case of referential opacity is in belief contexts, as in the invalid pattern (88). Here, it is clear that the *de dicto* – of what is said – belief John has about Cicero is not equivalent to a belief about Tully, despite the fact that the two might be extensionally equivalent. That is, the extensional equivalence of the singular terms in the second premise does not survive under a propositional attitude.

(88) I. John believes Cicero denounced Cataline
II. Cicero is Tully
III. John believes Tully denounced Cataline

Just as substitution of co-referential terms is not always licit, we cannot always existentially quantify over an object *inside* a belief context, as in (89).

[1] See Gamut (1991, 5) who formalize the principle as $s = t \models \phi \longleftrightarrow [t/s]\phi$, for terms s and t.

(89) $\exists x$(John believes that x denounced Catiline)

The problem with exporting the existential quantifier out of belief contexts is that we again cannot identify the individual about whom Ralph holds this particular belief of denouncement – given that Ralph believes of Cicero and not Tully that he denounced Catiline, the co-reference of Cicero and Tully notwithstanding.

Assuming again the validity of Leibniz's law, Quine saw the failure of referentially opaque contexts to obey this principle as motivation to ban quantification in all kinds of intensional contexts, including in particular propositional attitudes. Another influential example illustrating Quine's motivation for banning quantification in belief contexts turns on the idea that to do so could charge an individual with contradictory beliefs. Quine's original (1956) example is quoted below at length:

> There is a certain man in a brown hat whom Ralph has glimpsed several times under questionable circumstances on which we need not enter here; suffice it to say that Ralph suspects he is a spy. Also there is a grey-haired man, vaguely known to Ralph as rather a pillar of the community, whom Ralph is not aware of having seen except once at the beach. Now Ralph does not know it, but the men are one and the same. Can we say of this *man* (Bernard J. Ortcutt, to give him a name) that Ralph believes him to be a spy?

From the quotation above, we might say the following about Ralph, as he appears to belive that the man in the brown hat, namely Ortcutt, is a spy:

(90) Ralph believes that Ortcutt is a spy
$\exists x(x =$ Ortcutt $\wedge$ Ralph believes that x is a spy)

Again, we can say the following about Ralph, for he does not believe that the grey-haired man, a pillar of the community, is a spy:

(91) Ralph believes that Ortcutt is not a spy
$\exists x(x =$ Ortcutt $\wedge$ Ralph believes that x is not a spy)

Given that Ralph holds very different beliefs about Ortcutt, Quine asked whether we can charge Ralph with inconsistent beliefs, as in (92).

(92) Ralph believes that there is someone who is both a spy and not a spy
$\exists x(x =$ Ortcutt $\wedge$ Ralph believes that (x is a spy $\wedge$ x is not a spy))

Note that we have utterly disregarded the *description* under which Ralph himself would identify the object of his beliefs, namely Ortcutt. Surely, Ralph would not admit the contradictory statements (90–91), stated about *Ortcutt*, the man. Rather, Ralph might assent to the statements provided they were identified by a description associated with Ralph's beliefs, in this case *the man in the brown hat* and *the grey-haired pillar of the community*. The error, then, originates in ignoring Ralph's perspective; by naming the object *Ortcutt* according to our beliefs, we have invalidated any claim about the internal consistency of Ralph's beliefs.

Discussion of these cases has persisted through the last 50 years of analytic philosophy (cf., Kaplan (1969), Crimmins and Perry (1989), Recanati (2000), and Aloni (2001), to cite but a few). Although we cannot review all of what has been observed about such cases, let us briefly consider Aloni's remarks on descriptions in similar environments. Let *Portcutt* name the man who Ralph believes is a spy. According to the story above, in all of Ralph's belief worlds, the description *the man in the brown hat* refers uniquely to Portcutt, whereas the description *the grey-haired pillar of the community* refers to Ortcutt. Ralph, of course, does not know that the descriptions both pick out Ortcutt in the actual world.

Although a more complete discussion of Aloni's solution to the Ortcutt puzzle is postponed until the next section, the charge of inconsistency is avoided by properly relativizing the individuals picked out by the descriptions to the appropriate context of evaluation. That is, Ralph's beliefs may be wrong (in that his belief worlds may be inaccurate), but they are not internally inconsistent, in that the relevant descriptions pick out different objects in his belief worlds.

A similar complication of belief attribution holds in concealed question environments. To avoid the veridicality of *know*, let us modify the example slightly to *thinks he knows*. Suppose again that Ralph believes that the man in the brown hat, call him Portcutt, is a spy, and that the grey-haired pillar of the community, Ortcutt, is not. Under the CQ reading (93b), example (93a) states that Ralph believes that he can identify the individual denoted by the term *the man in the brown hat*, and similarly for the following case (94):

(93) a. Ralph thinks he knows the man in the brown hat
b. Ralph thinks he knows who the man in the brown hat is

(94) a. Ralph thinks he knows the grey-haired pillar of the community
b. Ralph thinks he knows who the grey-haired pillar of the community is

Yet, if asked whether the man in the brown hat is in fact Ortcutt, Ralph would have to disagree. Further, suppose that the man in the brown hat could be individuated by another unique description, say the man in the trenchcoat. While Ralph might in fact know who the man in the brown hat is, he need not think he knows the man in the trenchcoat. In particular, the following syllogism fails, despite the extensional equivalence of the terms in the second premise.

(95) I. Ralph thinks he knows the man in the brown hat
II. <u>The man in the brown hat is the man in the trenchcoat</u>
III. Ralph think he knows the man in the trenchcoat

If Ralph does not know that the man in the brown hat and the man in the trenchcoat refer to one and the same individual, the conclusion III does not hold. Yet, if he does acknowledge the equivalence, the conclusion is indeed warranted. This is precisely the same pattern observed by Heim and illustrated by Quine's dilemma. In the next section, we present the basic notion of conceptual covers and address its application to transparency paradoxes. In subsequent sections, we discuss how Aloni's solution to Quine's dilemma resolves Heim's entailment problem as well.

3.2 Conceptual Covers

In order to solve a number of long standing problems of identity ascription in philosophy, Aloni (2001) presents a formal pragmatics to account for the various ways in which identifying individuals varies according to different discourse contexts. In particular, she addresses exactly the sort of puzzles which we have argued apply to CQs, namely identity puzzles in transparency paradoxes.

The insight of Aloni (2001) is that methods for identifying individuals vary according to perspective and discourse environment.[2] As such, the *ways we specify* individuals are relativized to *conceptual covers*, sets of individual concepts, which "represent different ways of looking at one domain."

Conceptual covers are subject to a stringent constraint; for a set of individual concepts, CC, to qualify as a conceptual cover, each individual of D_e must be identified by one and only one individual concept $c_i \in CC$. Aloni (2001, 64) sums up her position:

> A conceptual cover is a set of individual concepts that satisfies the following condition: in a conceptual cover, in each world, each individual constitutes the instantiation of one and only one concept.

Formally, a conceptual cover must satisfy condition 3.2.1 below.

Definition 3.2.1. [Conceptual Covers] Let W be a set of worlds and D the domain of individuals. A CC based on $\langle W, D\rangle$ is a set of functions $W \mapsto D$, such that:

$$\forall w \in W : \forall d \in D : \exists! c \in CC : c(w) = d$$

Adherence to Definition 3.2.1 requires that each individual is uniquely selected by one and only individual concept. Conceptual covers thus enforce the two conditions of *existence* and *uniqueness* on the domain of individuating concepts. An informal definition of these conditions is worth repeating:

> Conceptual covers are sets of [individual] concepts which exhaustively and exclusively cover the domain of individuals. In a conceptual cover, each individual d is identified by at least one concept in each world (existence), but in no world is an individual counted more than once (uniqueness). Aloni (2001, 64)

[2] "Different methods of identification are operative in different conversational circumstances and the evaluation of fragments of discourses can vary relative to these methods."Aloni (2001, ix)

There are many ways to identify an individual, *e.g.,* naming, ostension, unique definite description, *etc.* A conceptual cover collects various ways of individuating each individual in the domain, just as long as each individual is uniquely identified once and only once.

To slightly alter an example from Aloni (2001, 64), suppose that the domain consists of just two individuals, the brothers *Aaron* and *Baron*. You see that Aaron is standing to your left, and Baron to your right. You now have a two concepts which suffice to select these two individuals uniquely. Collecting them into a set constitutes the creation of a conceptual cover.

Suppose further that you ascertain a list of facts about the characters in front of you, summarized in Table 3.1.

	Individuating Concepts		CC?
1.	Aaron	Baron	✓
2.	to the left	to the right	✓
3.	linguist	economist	✓
4.	curly hair	straight hair	✓
5.	brown hair	brown hair	*
6.	Howard's grandson	Howard's grandson	*

Table 3.1: Aaron and Baron

Clearly, there are several ways to individuate Aaron from Baron, and Baron from Aaron: by name (1), by relative position (2), by profession (3), even by hair shape (4). These qualities can be collected into conceptual covers, as each 1 – 4 picks out each individual in the domain uniquely and exhaustively (96). However, Aaron and Baron share a few things in common, such as hair color (5) and lineage (6). These qualities do not individuate Aaron and Baron uniquely, and thus cannot be included in any conceptual cover considered here. Furthermore, there are qualities which fail to pick out any individual in this domain, for instance, being female or Dutch. Such concepts are likewise excluded from any cover on this domain.

(96)
i. $\mathrm{CC}_1 = \{\text{Aaron, Baron}\}$
ii. $\mathrm{CC}_2 = \{\text{left, right}\}$
iii. $\mathrm{CC}_3 = \{\text{linguist, economist}\}$
iv. $\mathrm{CC}_4 = \{\text{curly, straight}\}$

Let the concepts c_1 and c_2 represent the names 'Aaron' and 'Baron', respectively. Let the concepts c'_1 and c'_2 represent the relative positions 'left' and 'right'. Taking covers CC_1 and CC_2 for illustration, Figure 3.1 depicts the mapping between individuating concepts c_i in various conceptual covers CC_i in the set of all covers C and individuals **a** in the domain of individuals. In our example, CC_1 the set of naming concepts exhaustively specifies D_e by mapping each name to the individual it denotes. The cover CC_2 similarly maps the positions 'left' and 'right' to Aaron and Baron respectively. Both covers individuate the domain completely and exhaustively, albeit with rather different concepts. Note that the mapping between elements in CC_i and D_e must be a bijection according to our definition of Conceptual covers.

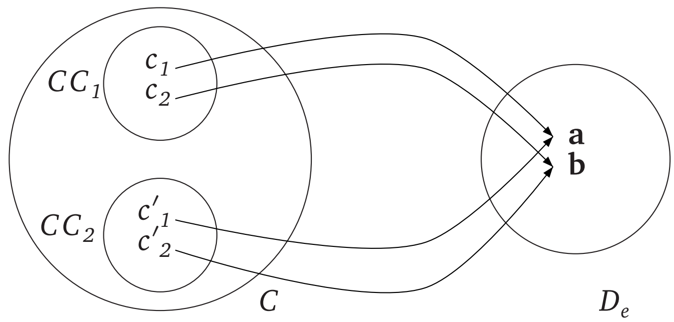

Figure 3.1: Mapping from conceptual covers in C to D_e

Note, however, that conceptual covers are not required to select individuals by consistent methods, it only matters that the individual concepts conform to existence and uniqueness. For example, we could individuate Aaron and Baron by a cross between name and relative position (97). Such a conceptual cover would be entirely appropriate if, for instance, you know Aaron's name but not Baron's.

(97) $\text{CC}_{1\times 2}$ = {Aaron, right}

Including redundant information in a CC, such as a name and a relative position denoting the same individual, would violate uniqueness (98a), while omitting a concept that picks out an individual would violate existence (98b).

(98) IMPROPER COVERS
  a. *$\text{CC}_{1\times 2}$ = {Aaron, left, Baron}
  b. *CC_1 = {Aaron}

Suppose that you have discovered a great many facts about the two brothers, and yet have not met them in person. You walk into a room and see them standing side by side. For all the information you have collected, you can't tell which one is older.

For simplicity, let us consider only two CCs, {Aaron, Baron} and {left, right}. Each relevant difference is modeled by a world below (the * marks the older brother):

w_1	$\mapsto$	A	B^*
w_2	$\mapsto$	B^*	A
w_3	$\mapsto$	A^*	B
w_4	$\mapsto$	B	A^*

In w_1 and w_2, Baron is older; in w_3 and w_4 Aaron is older. In w_2 and w_3 the brother on the left is older, and in w_1 and w_4 the one on the right is older.

The worlds in the model above can be split up according to the method of identification used. In Aloni (2001) the division is expressed as a *conceptual perspective*, a mapping from indices to conceptual covers. Thus perspectives group CCs together that make the relevant distinctions between individuals.

Definition 3.2.2 (Conceptual Perspectives)**.** Let $M = \langle W, D, C \rangle$ be a model and N be the set of indices in our language, and C the set of conceptual covers based on $\langle W, D \rangle$. A *conceptual perspective* $\wp$ in M is a function from N to C.

The perspectives below carve the worlds above according to different criteria: $\wp$ assigns the CC that identifies the brother by means of his name to the index of the variable x, whereas the later individuates the brothers by their relative position. Making intensions explicit, our example perspectives yield the following conceptual covers:

(99) Conceptual Perspectives
   i. $\wp(x) = \{\lambda w.\text{Aaron}(w), \lambda w.\text{Baron}(w)\}$
   ii. $\wp'(x) = \{\lambda w.\text{left}(w), \lambda w.\text{right}(w)\}$

Perspectives prove to be an important feature in structuring worlds according to relevant dimensions. Consider the following question (100).

(100) Which brother is older?

As discussed in the Introduction, a partition semantics for questions structures logical space by collecting equivalent worlds that agree on possible answers to the question in that world. For any given world v, the denotation of an interrogative then corresponds to the complete true answer ϕ of that question in v. Interrogatives group worlds in which the individual or individuals satisfying the proposition ϕ is the same.

The essential difference in the theory of conceptual covers is that the individuals satisfying the propositional answer ϕ are relativized to concepts referring to those individuals in the relevant worlds of evaluation.

Definition 3.2.3 (Interrogatives under Cover)**.**

$$[?\overrightarrow{x}\phi]^{\wp}_{w,g} = \{v \in W | \forall \overrightarrow{c} \in \prod_{i \in n}(\wp(x_i)) : [\phi]_{w,g[\overrightarrow{x}/\overrightarrow{c}(w)]} = [\phi]_{v,g[\overrightarrow{x}/\overrightarrow{c}(v)]}\}$$

$\lambda w.$ no $c_i(w)$ is older in w
$\lambda w.$ only $c_1(w)$ is older in w
$\lambda w.$ only $c_2(w)$ is older in w
$\vdots$
$\lambda w.$ only $c_1(w)$ is older in w $\wedge$ $c_2(w)$ is older in w
$\vdots$
$\lambda w.$ all $c_i(w)$ is older in w

Figure 3.2: Partitions induced by (100) under $\wp$

For a perspective $\wp(x_i) = \{c_1, c_2, \ldots\}$, the interrogative in (100) induces a set of partitions, depicted in Figure 3.2, in which c_i denotes an individual d in w, for which the proposition *d is older* holds true.

Taken together, these partitions correspond to the set of possible complete answers to the question. The perspective $\wp$ groups worlds together in which the eldest brother is identified by name, while $\wp'$ collects worlds by positions. For example, under $\wp$ the block $\{w_1, w_2\}$ collects the worlds in which (a) Baron is the oldest, and (b) Baron is identified by name. In contrast, under $\wp'$ the block $\{w_1, w_2\}$ is formed by gathering worlds in which (a) Baron is the oldest, and (b) he is identified not by name, but by position.

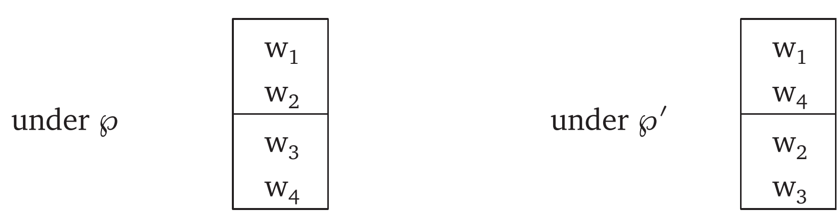

Figure 3.3: Partitions of Logical Space under $\wp$ and $\wp'$

As Aloni (2001) observes, conceptual covers provide a natural way to preserve the informativity of identification questions. For example, knowing that Baron is the older brother may not help you identify which person is which.

(101) Which brother is which?

Thus, question (101) structures logical space into two partitions, one in which the Aaron is on the *left*, and Baron on the *right*, $\{w_2, w_4\}$, and another in which Baron is the *left*, and Aaron on the *right*, $\{w_1, w_3\}$.

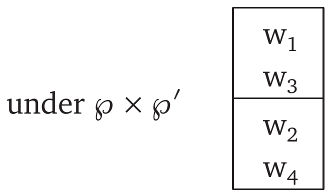

Figure 3.4: Partition induced by (101) under $\wp \times \wp'$

By introducing conceptual covers, multiple ways of identifying individuals are allowed to be active in questions, and some ways may be more useful or appropriate in different contexts. The theory of conceptual covers also avoids trivializing identity questions (102), under partition semantics.

(102) Who is Ortcutt?

$?x.x = o$

Aloni argues that if we assume, after Kripke (1972), that names are rigid designators, and a complete answer to a constituent identity question like (102) must be a rigid term, then the answer to the above question is trivial; the answer is simply that Ortcutt is Ortcutt, which by itself is completely uninformative. Yet, such questions have natural and intuitive uses, for instance if you join a conversation about Ortcutt and do not know who that individual is.

The solution that Aloni (2001, 23) offers relativizes the notion of answerhood to conceptual perspectives, so that whether an identity question remains trivial depends on the perspective $\wp$ employed (*cf.,* Definition 1.2.8).

Definition 3.2.4 (Answers under Cover)**.** Let ψ and $?\overrightarrow{x}\phi$ be closed sentences.

1. ψ is a (partial) answer to $?\overrightarrow{x}\phi$ in M *under* $\wp, \psi \rhd_{M,\wp} ?\overrightarrow{x}\phi$ iff

$$\exists X \subset [[?\overrightarrow{x}\phi]]_{M,\wp} : [[\psi]]_{M,\wp} = \cup\{\alpha | \alpha \in X\} \neq \emptyset$$

2. ψ is a complete answer to $?\overrightarrow{x}\phi$ in M *under* $\wp, \psi \rhd_{M,\wp} ?\overrightarrow{x}\phi$, iff

$$[[\psi]]_{M,wp} \in [[?\overrightarrow{x}\phi]]_{M,\wp}$$

In particular, identity questions are trivial only if the perspective $\wp$ contains the rigid designator queried already; that is, an identity question is trivial only if one asks about the identity of an individual who one can already identify.

As the theory of conceptual covers also provides an interesting solution to Quine's dilemma, we briefly discuss the proposal in the next section. This particular solution will be able to provide some important insights into Heim's entailment property in the sections that follow.

3.2.1 Conceptual covers and Quine's dilemma

Recall from §3.1.1 that Ralph held very different beliefs about an individual he identified in different ways. Specifically, Ralph believes that the man with the brown hat is a spy, and that the grey-haired pillar of the community is not, despite the fact that they pick out one and the same individual in the actual world. Let the domain D_e consist of just two individuals Ortcutt and Portcutt, and w_1 be the sole belief state in Ralph's belief alternatives $Bel_R(w_0)$. Thus in Ralph's belief world w_1, Ortcutt is the pillar of the community, and Portcutt is the man with the brown hat, who Ralph believes is a spy. In the actual world, w_0, both terms refer to Ortcutt. In all worlds Portcutt is a spy, and Ortuctt is not. In this situation, there are four possible concepts, **a, b, c,** and **d**, and 2 possible covers.

(103) CC_1

		a	**b**
w_0	$\mapsto$	o	p
w_1	$\mapsto$	o	p

(104) CC_2

		c	**d**
w_0	$\mapsto$	o	p
w_1	$\mapsto$	p	o

Thus, whether x is a spy depends on what concept is assigned to n below (in which $\Box$ represents 'believes that'). If n is mapped to **a**, (105) is false, and if n is mapped to **c**, then it is true. For the concept **a** is the interpretation of 'the pillar of the community', whereas **c** is 'the man in the brown hat.'

(105) $\Box S(x_n)$

So, Ralph's beliefs about Ortcutt can be modeled by mapping the index on the variable x to different covers:

(106) Ralph believes that Ortcutt is a spy
$\exists x_n(x_n = o \wedge \Box S(x_n))$

(107) Ralph believes that Ortcutt is not a spy
$\exists x_m(x_m = o \wedge \Box S(x_m))$

Aloni (2001) thereby accounts for the two ways of viewing Ortcutt by differentiating the conceptual covers that Ralph uses to individuate Ortcutt. Thus, we are able to distinguish Ralph's *de re* beliefs from our own method of identifying the objects in those beliefs. Notably, (109) is satisfiable, but the contradictory (108) is not, for $i \in \{n, m\}$, in which $n \neq m$:

(108) $\exists x_i(x_i = o \wedge \Box(S(x_i) \wedge \neg S(x_i)))$

(109) $\exists x_n(x_n = o \wedge \exists y_m(y_m = o \wedge \Box(S(x_n) \wedge \neg S(y_m))$

Note that the mutual exclusivity of the covers $\{\mathbf{a}, \mathbf{b}\}$ and $\{\mathbf{c}, \mathbf{d}\}$ is not arbitrary. We simply cannot ascribe Ralph inconsistent beliefs under a single cover as the set of concepts required to do so $\{\mathbf{a}, \mathbf{c}\}$ does not qualify as a conceptual cover, as it both fails to pick out Portcutt in w_0 (violating existence), and picks out Ortcutt more than once in w_0 (violating uniqueness).

To summarize in brief, conceptual covers are sets of individual concepts which uniquely pick out each member of the domain of individuals. The sets of concepts used varies according to perspectives $\wp(x_i)$, functions from indices to conceptual covers. Questions about the identity of an individual, then, quantify directly over individual concepts, rather than the individuals they refer to. Intuitively, this grants a certain degree of flexibility in answering the question of *who* or *what* something is. We systematize this flexibility in our analysis of CQs below, in which CQs are taken as constituent answers to identity questions, namely as objects of type $\langle\langle s, e\rangle, \langle s, \langle s, t\rangle\rangle\rangle$. We will see that CQs crucially involve the pragmatic notion of conceptual cover in order to account for the flexibility of CQ interpretations and Heim's entailment problems.

3.3 Concealed Questions and Conceptual Covers

The primary goal of this section is to introduce a theory of CQs which *(i)* incorporates conceptual covers, and *(ii)* explicitly addresses the flexibility of interpretations discussed above. The theory advocated here considers CQ meanings as answers to identity questions. Specifically, it is argued that concealed questions denote a relation between an individual concept from a conceptual cover and questions in which that concept serves to individuate the noun in the CQ. Informally, this amounts to answering the question of how the NP should be identified. We argue for this account both on philosophical and linguistic grounds.

Secondary aims of this section include illustrating a novel way of deriving Heim's ambiguity – one that gives semantic content to the syntactic notion of Late Merger as discussed in §2.1.2, and respects the intuition that the relative clause modifies very different linguistic objects across the two readings. In particular, it is argued that, in Reading A, the relative clause restricts the price concept with semantic content evaluated *outside* of the matrix subject's belief worlds, whereas the relative clause in Reading B modifies an object evaluated *with respect to* the subject's belief worlds.

In addition, another aim of following section involves representing the proper source of context dependency in CQs. As illustrated above, the primary role of contextual information is quite limited. It is assumed that variance due to context sensitivity does not stem from different relations being picked out in context (as in Heim (1979)) or from choosing between a collection of salient propositions (as in Nathan (2006)). Rather, it is argued that the major source of context sensitivity is due to the following two facts: (a) different conceptual covers divide up the domain in different ways, such that these methods are informative in different ways, and (b) Conceptual Covers requires that the CQ NP and the concept that individuates it be unique enough so as to specify exactly one individual. The present theory has the advantage of tightly constraining the role of context, while offering a principled explanation of apparent exceptions. With additional evidence from Heim's ambiguity, it is argued that the specificity condition (b) constrains how implicit arguments in relational nouns are interpreted.

Before illustrating the relationship between CCs and CQs, it will be useful to define the notion of *meta-perspective of an agent a*. A meta-perspective is simply the set of Conceptual Covers from the set C of all Conceptual Covers based on $\langle M, D \rangle$ with which the agent can, given the individuating concepts she has at her disposal, uniquely and exhaustively specify the members of the domain of individuals.

Definition 3.3.1 (Meta-perspective.). In a model $\mathscr{M} = \langle W, D, R, I, C \rangle$, a *meta-perspective*, $\mathfrak{M}\wp(a) \subseteq C$, is the set of conceptual covers available to agent a, namely the set $\mathfrak{M}\wp(a) = \{\wp_i(a) | \forall i \in N\}$.

It is worth noting that a meta-perspective is not the set of all logically possible perspectives, but is restricted to those that an agent could logically generate using the individuating concepts she has. In other words, the meta-perspective is simply all the ways an agent could uniquely and exhaustively specify the domain D. For example, if an agent has but one way, c, to individuate D_e then the meta-perspective will consist of a single set, $\{c\}$. Further note that we stipulate that $\mathfrak{M}\wp$ will not typically contain every possible way to individate D_e, given that there may be individuating concepts which are to remain unknown to the agent.[3] Table 3.2 displays the hierarchical and conceptual relationship between meta-perspectives, perspectives, covers, and concepts.

[3]In particular I want to keep *logical* and *acknowledged* meta-perspectives separate, as the number of logically possible covers that can be formed from the various perspectives grows much faster than set $\mathfrak{M}\wp$ as I have defined it above. Note that $\mathfrak{M}\wp$ is a subset of the logically possible meta-perspectives.

Notation	Name	Description	Type
$\mathfrak{M}\wp$	meta-perspective	set of all CCs available to the agent	$\langle\langle\langle s,e\rangle,t\rangle,t\rangle$
$\wp_i$	perspective	function from indices N to CCs	
CC	conceptual cover	set of concepts that uniquely and exhaustively specify D_e	$\langle\langle s,e\rangle,t\rangle$
c	concept	individual concept that individuates some $d \in D_e$ uniquely	$\langle s,e\rangle$

Table 3.2: Meta-perspective to concept hierarchy

It is important to note that meta-perspectives are not formally crucial to the account that follows and could possibly be dispensed with. However, their intuitive use will become clear when accounting for Greenberg's generalization below. Armed with the additional notion of meta-perspective, we are prepared to present a novel semantics for concealed questions.

3.3.1 The Denotation of CQs

We have labored above to show that CQs represent a shift of perspective from the object the NP denotes to a characteristic way of picking that object out, according to a given perspective. A formal analysis that accords with this intuition should *(i)* allow the NP to continue to denote an individual concept (evaluated in the appropriate world), *(ii)* incorporate a relation of identity between the CQ NP and another individual concept, *(iii)* constrain the types of individual concepts that can pick out the object in the correct circumstances, *(iv)* include an element of context dependence to account for changes in perspective which appeared in the paradoxes of equivalence above. Two conceptions of CQs are presented below; the first attempts to elucidate CQs as relations between *constituent answers* and *identity questions*. The second conception exploits Schönfinkelization (Schönfinkel, 1924) and explains CQs as mapping between pairs of worlds $\langle w, w'\rangle$ and sets of concept that individuate the CQ NP in both w and w' and are contained in CCs in a meta-perspective $\mathfrak{M}\wp$ of some agent a.

3.3.1.1 CQs as relations between answers and questions

Our first conception of concealed questions relates individuating concepts to identity questions in which those concepts serve to individuate the CQ NP in a question about its identity. The semantics are formalized as follows:

Definition 3.3.2. [Concealed Question Denotation. Conception 1 of 2.]

In a model $M = \langle W, D, R, I, C\rangle$, a *concealed question interpretation of* $\alpha_{\langle s,e\rangle}$ expresses a function $CC \mapsto (W \mapsto W')$, such that a concept $c \in CC \in \mathfrak{M}\wp(a)$ evaluated at w specifies the individual $d \in D_e$ denoted by α at w and c at w' specifies the individual denoted by α at w'.

Semantically, a concealed question has the representation in λ-notation (110.i), in which α represents the DP from which the CQ originates. A set-theoretic counterpart is shown in (110.ii)

(110) CONCEALED QUESTION REPRESENTATION:

i. $\lambda c_{\langle s,e\rangle}.\lambda w.\lambda w'.c \in \wp_i(a)[c(w) = \alpha_{\langle s,e\rangle}(w) \wedge c(w') = \alpha_{\langle s,e\rangle}(w')]$

ii. $\{c \in CC \in \wp_i \mid \{\langle w, w'\rangle \in W \times W \mid c(w) = \alpha(w) \wedge c(w') = \alpha(w')\}\}$

As (110) is intended to show, it is clear that under Definition 3.3.2, a CQ is defined as the set of concepts which for a pair of worlds, w, w' specifies the CQ NP α in the same way in w and w'. In effect, given a concept c and a world w, a CQ induces equivalence classes on logical space precisely as a matrix question might on the partition semantics view of questions (Groenendijk and Stokhof, 1984). The schematic effect a CQ has on its context is show in Figure 3.8. For ease of reference, I'll adopt the following convention:

(111) CQ DENOTATION CONVENTION:

Write $\lambda c.\lambda w.\lambda w'.c \in \wp_i : c(w)(w') \doteqdot \alpha(w)(w')$ for

$\lambda c_{\langle s,e\rangle}.\lambda w.\lambda w'.c \in \wp_i(a)[c(w) = \alpha(w) \wedge c(w') = \alpha(w')]$.

For a given meta-perspective of agent a, $\mathfrak{M}\wp(a_i)$,[4] suppose that there are $n \in \mathbb{N}$ conceptual covers, such that $\mathfrak{M}\wp = \{CC_1, CC_2, \ldots, CC_n\}$. Given the definition of conceptual covers, one and only one individuating concept c_i in each cover $CC_{m\leq n}$ that maps onto each individual of the domain of individuals $D_e(w)$ at world w.[5]

Note that a concept c_i need not be uniquely associated with one cover, although it could be in principle. In fact, there could be only one concept c_1 which appears in every $CC_{m\leq n}$, just so long as there is no other concept c_2 in any $CC_i \in \mathfrak{M}\wp$ such that $\exists w.\exists d : c_1(w) = c_2(w) = d \in D_e(w)$. In the case of a single $c(w) = d$ in all $CC \in \mathfrak{M}\wp$, the agent would only have one way at her disposal to successfully individuate that particular d.

How does the denotation in (110) relate to questions? Feeding the CQ representation a concept, say c_1, from $\mathfrak{M}\wp$ yields a propositional concept, a function from worlds to propositions, as shown in (112a). Further, when evaluated at a particular world, say w, the result is a propsition which picks out the set of worlds u for which the same concept c_1 specifies α in u and in the reference world w, as in (112).

(112) a. $\lambda w.\lambda w'.[c_1(w)(w') \doteqdot \alpha(w)(w')]$

b. $\lambda w'.[c_1(w)(w') \doteqdot \alpha(w)(w')]$

Under a partition semantics for questions, we may represent effect that the proposition has on logical space as a bipartition, depicted in Figure 3.5, divides the worlds into equivalence class based on whether or not they satisfy (112b).

[4]The relativization to agent a and index $i \in N$ will sometimes be omitted, sacrificing the more precise notation for the sake of readability.

[5]I remain agnostic about whether the domain of individuals remain constant or vary across diffrent possible worlds. For our purposes here, a constant domain is assumed to simplify the discussion.

$\lambda w'.c_1 \in \wp_i : c_1(w)(w') \doteqdot \alpha(x)(w)(w')$	Worlds where concept c_1 specifies α
$\lambda w'.\neg[c_1 \in \wp_i : c_1(w)(w') \doteqdot \alpha(x)(w)(w')]$	Worlds where concept $c_i \neq c_1$ specifies α

Figure 3.5: Bipartition of Logical Space

Taking another concept c_2 in the relevant set of covers, $\wp$, and applying it to the CQ representation in (110), we again get a bipartition of logical space as in Figure 3.6.

$\lambda w'.c_2 \in \wp_i : c_1(w)(w') \doteqdot \alpha(w)(w')$	Worlds where concept c_2 specifies α
$\lambda w'.\neg[c_2 \in \wp_i : c_2(w)(w') \doteqdot \alpha(w)(w')]$	Worlds where concept $c_i \neq c_2$ specifies α

Figure 3.6: Bipartition of Logical Space

Taking the union of the bipartitions in Figures 3.5 and 3.6 yields a third division of logical space as shown in 3.7.

If the procedure is repeated pointwise through the set of covers $\wp$, the result is exactly that of a identity question in which each cell in the partition contains worlds which identify the individual in question by one unique concept, as in Figure 3.8.[6]

Given that the CQ, when supplied an individuating concept, partitions logical space in the same way that an identity question does, the relationship between CQs and identity question is made clear. Concealed questions represent functions from ways of identifying the CQ NP in all the different ways of individuating the domain, to questions regarding the identity of the CQ NP.

The intuition that this analysis exploits is that CQs and identity questions carve up logical space in very similar ways. Worlds are eliminated if they do not agree in the concept that specifies the CQ NP, through the elimination of entire cells of the partition which do not properly pick out the individual in question with an adequate concept.

[6]The same result is achievable through somewhat different means. We might instead collect the propositions formed by combining the CQ denotation and each concept c and a reference world w_0 and take the generalized union. Given that every conceptual cover CC must individuate the domain exhaustively, the result will be a complete partition of logical space as in Figure 3.8.

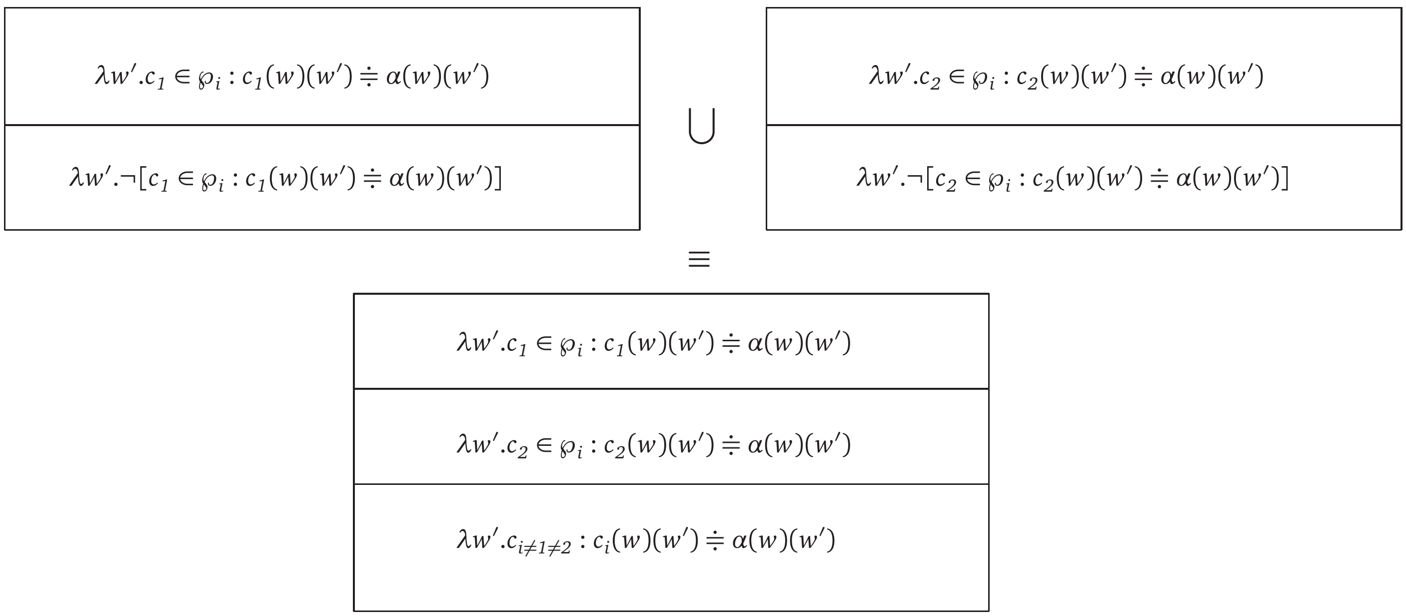

Figure 3.7: Union of Bipartitions

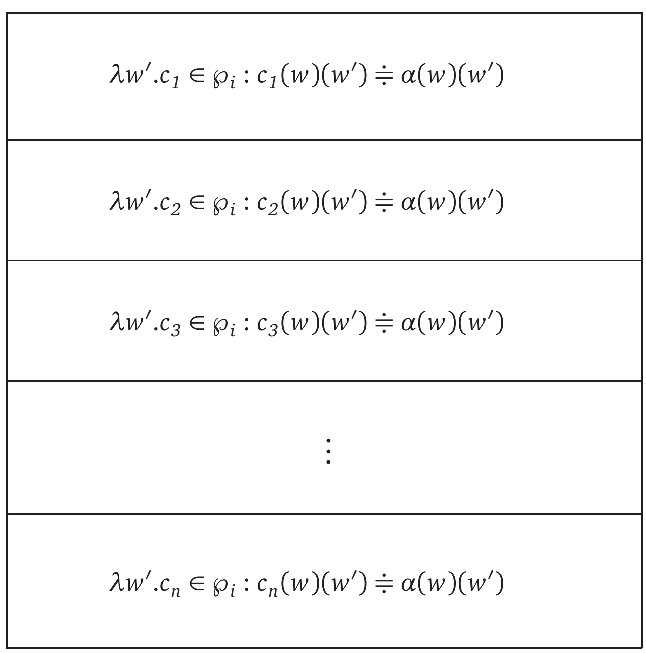

Figure 3.8: CQ division of logical space

Note that an agent may have multiple ways of identifying an individual, that is, have multiple concepts for specifying that individual. This accords entirely with the intuition that an agent may qualify how to conceive of the individual in question, as in (113). These qualifications are not entirely redundant, for the fact that the *Aaron Aaronson* is also the *newest member on the faculty* is non-trivial information. It potentially relates two concepts by noting that their extensions are equivalent in the actual world, and that they might be distinct, but equally adequate methods for specifying this individual. This means that any set of concepts which contains both these concepts fails to be a proper cover.

(113) A. Can you guess UvA's most popular philosopher?
B. I know UvA's most popular philosopher; it's *Aaron Aaronson*: the newest member on the faculty.

Further, B can continue to add information about Aaron Aaronson without it necessarily applying uniquely to that individual. For instance, mentioning *he's also a squash fanatic*, does not indicate that this information suffices to identify that individual under discussion, only that the property applies to Prof. Aaronson. The point to keep in mind is that CQs, at least under this analysis, are simply about possible ways of identifying individuals denoted by the CQ NP. Other properties and concepts still apply of these individuals, but just not as proper answers. The denotation presented in Definition 3.3.2 allows for a distinction between various kinds of concepts that can serve as an answer to the question of the CQ NPs identity. This distinction, in turn, allows us to gain significant insight not only into cases where the specification of the NP is flexible, but into the ways in which it effects the context in these different circumstances. The resolution is illustrated more fully in sections that follow.

3.3.1.2 Defining the type shifting operation $\mathscr{Q}$

As it stands, it remains mysterious how CQ NPs ever come to denote these higher-typed intensional objects. To ensure the CQ is of the correct type, we posit the type-shifting rule $\mathscr{Q}$:

($\mathscr{Q}$) CONCEALED QUESTION TYPE SHIFTER

$$\langle s,e\rangle \mapsto \langle\langle s,e\rangle,\langle s,\langle s,t\rangle\rangle\rangle$$
$$\alpha_{\langle s,e\rangle} \mapsto \lambda c_{\langle s,e\rangle}.\lambda w.\lambda w'.c \in \wp_i(a)[c(w)=\alpha(w) \wedge c(w')=\alpha(w')]$$

A more extensive discussion on type-shifting and compositionality follows in Chapter 4, preceding a presentation of an experimental study showing that there is a processing cost associated with interpreting concealed questions in sentential contexts. Let it suffice it to say here that this rule is not *merely ontological*, but that it represents something *natural*, in the sense of Partee (1986, 1995). I argue that the type-logical shift is accompanied by a cognitive shift in perspective which *maximally perserves information across domains*. In essence, the idea is that $\mathscr{Q}$ can be defended on several levels of explanation in addition to merely satisfying the mechanistic requirements of semantic compositionality.

In her classic article, Partee (1986) proposes the idea that there is no single uniform and universal set of type-shifting principles. Type-shifting operations claim a variety of origins: (i) they may be derived directly from type theory, (ii) they may be general, but dependent on the domain, (iii) they may impose additional structure on the domain, (iv) they may reveal language particular rules. The little structure that operation $\mathscr{Q}$ imposes on its argument is closely related to how that argument is viewed in a slightly different context.

Much like an highly-intensionalized variant of Partee's *ident* operation, which lifts entities to singleton sets containing that entity, $\mathscr{Q}$ alters the way we conceive of the domain. Instead of picking out individuals directly, $\mathscr{Q}$ asks us to consider the various ways in which individuals could be identified. As such, I believe that $\mathscr{Q}$ is a 'cognitively natural' operation. As the second, Schöfinkelized variant of the CQ denotation better reveals the cognitive naturalness of $\mathscr{Q}$, it is to this conception that we now attend.

3.3.1.3 CQs as an intensionalized set of individual concepts

This section is an exposition of an alternative view of CQs, which derives from the first via a correspondence to another type domain secured by Schönfinkelization. On this latter conception, CQs are merely sets of individual concepts, relativized to a pair of worlds, that are uniquely paired to the CQ NP when evaluated at these indices. A full discussion of the limitations that Schönfinkelization imposes on a logic is beyond the reach of this section.[7] Let it suffice to merely state the technique in general, and to gesture towards its application in the current context.

Definition 3.3.3. [The Schönfinkel correspondence (Informal).] Let σ, τ, ρ be (possibly non-distinct) types in the set of well-formed types TT. There is an isomorphism between $D_{\langle\sigma,\langle\rho,\tau\rangle\rangle}$ and $D_{\langle\sigma\times\rho,t\rangle}$.

For our purposes, Schönfinkelization can be exploited in the following way: since any object in $D_{\langle\langle s,e\rangle,\langle s,\langle s,t\rangle\rangle\rangle}$ has a corresponding object in $D_{\langle s\times s,\langle\langle s,e\rangle,t\rangle\rangle}$,[8] we redefine the concealed question denotation of Definition 3.3.2:

[7] For excellent discussions, see Muskens (1995) and Tichý (1982) who show that the technique is a redundancy limited only to total models. Partial models lose the result in general, though it remains to be seen what classes of type-theoretic object retain the correspondence after partialization.

[8] Let $\delta = \langle s, e\rangle$, $\rho = s$, and $\tau = \langle s, t\rangle$. There is a isomorphism between objects in $D_{\langle\langle s,e\rangle,\langle s,\langle s,t\rangle\rangle\rangle}$ and $D_{\langle\langle s,e\rangle\times s,\langle s,t\rangle\rangle}$, by Schönfinkelization. The pairs $\delta\times\rho$ and $\rho\times\delta$ are equivalent. So there is an isomorphism between $D_{\langle\langle s,e\rangle,\langle s,\langle s,t\rangle\rangle\rangle}$ and $D_{\langle s\times\langle s,e\rangle,\langle s,t\rangle\rangle}$. There exists an isomorphism between $D_{\langle s\times\langle s,e\rangle,\langle s,t\rangle\rangle}$ and $D_{\langle s\times s\times\langle s,e\rangle,t\rangle}$, again by Schönfinkelization, which guarantees an isomorphism between $D_{\langle s\times s\times\langle s,e\rangle,t\rangle}$ and $D_{\langle s\times s\,\langle\langle s,e\rangle,t\rangle\rangle}$. Therefore, there exists an isomorphism between $D_{\langle\langle s,e\rangle,\langle s,\langle s,t\rangle\rangle\rangle}$ and $D_{\langle s\times s,\langle\langle s,e\rangle,t\rangle\rangle}$.

Definition 3.3.4. [Concealed Question Denotation. Conception 2 of 2.] In a model $\mathscr{M} = \langle W, D, R, I, C \rangle$, a *concealed question interpretation of* $\alpha_{\langle s,e \rangle}$ expresses a function $g : W \times W \mapsto \cup CC$, such that any pair of worlds $\langle w, w' \rangle \in W \times W$ returns the set of concepts $(\!|c|\!)_{\alpha,w,w'}$ that agree on the specification of α in w and w'. We call the set $(\!|c|\!)_{\alpha,w,w'}$ the set of *concepts generated by* α *in w and w'.*

Intuitively, $(\!|c|\!)_{\alpha,w,w'}$ is just the set of ways to uniquely specify α in w or w'. The following result is immediate:

Note 3.3.1 (Uniqueness). For any $c_i, c_j \in (\!|c|\!)_{\alpha,w,w'}$, where $i \neq j$, there is no cover CC_k such that $c_i, c_j \in CC_k$.

Simply put, no concept in $(\!|c|\!)$ can occur with any other concept in $(\!|c|\!)$ in any conceptual cover. To do so would violate uniqueness. This does not mean that a concept in $(\!|c|\!)$ cannot occur in *multiple* covers. Suppose for example, that $(\!|c|\!)_{\alpha,w,w'}$ is just the singleton set $\{\lambda w.Aaron(w)\}$, meaning that this agent has one and only one way of specifying Aaron uniquely against the other members in the domain at the relevant world of evaluation. Even in this limited case, there may be any number of other covers in the metaperspective $\mathfrak{M}\wp$ that pick out other members according to different methods of identification. For instance if $CC_1 = \{\lambda w.Aaron(w), \lambda w.Baron(w)\}$, there could exist another cover $CC_2 = \{\lambda w.Aaron(w), \lambda w.\iota x.economist(x)(w)\}$, for which the only common concept between them is the one that names *Aaron*. From the viewpoint of *covers*, the following fact means that $(\!|c|\!)_{\alpha,w,w'}$ collects all the concepts *exhaustively*.

Note 3.3.2 (Exhaustivity). There is no $CC_k \in \mathfrak{M}\wp$ such that $c \in CC_k$ and $c(w) = \alpha(w)$ and $c \notin (\!|c|\!)_{\alpha,w}$.

In other words, each cover in $\mathfrak{M}\wp$ must contain one and only one member of $(\!|c|\!)_{\alpha,w,w'}$ after the context is updated with the CQ.

The intuitive relationship between Conceptual Covers, the domain of individuals, and the set $(\!|c|\!)$ is depicted in Figure 3.9. Like Figure 3.1, concepts c_i from covers CC_i in a metaperspective $\mathfrak{M}\wp$ map onto elements **a,b** in D_e. In general, the set $(\!|c|\!)_{\mathbf{n}}$ collects all the ways to individuate object **n** according to $\mathfrak{M}\wp$.

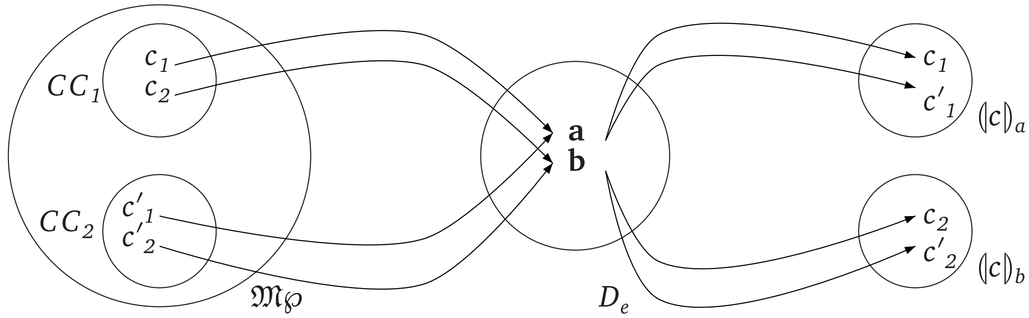

Figure 3.9: Generation of $(\!|c|\!)_a$ and $(\!|c|\!)_b$

Having gained some intuition about how CQs function, it is appropriate to return to the question of whether $\mathcal{Q}$ can be defended on conceptual grounds to be 'cognitively natural.' At an intuitive level, the result of applying $\mathcal{Q}$ to a noun α consists only in a shift of perspective: from the *denotation* of α to alternative conceptions of α. The different ways of viewing α are collected in the set $(\!|c|\!)_\alpha$. A central reason for limiting the domain of conceptual covers to $\mathfrak{M}\wp$ is that it might seem unreasonable to allow *any logically possible* individuation of α into $(\!|c|\!)_\alpha$. Rather, we want to consider individuating concepts available to a particular agent to accord with the description of $\mathcal{Q}$ as '*cognitively* natural', rather than, say, '*logically* natural.'

Furthermore, $\mathcal{Q}$ is 'natural' in the sense that it maximally preserves information across domains (see Partee (1995, 348) for a discussion of a family of natural functions). The identity of each object in D_e has a unique match in each conceptual cover. In a sense, applying Q simply does what an identity question asks of the discourse particpants – the question of what something or other is is just a request for determining how to view that object or thing. Which concept is among those in $(\!|c|\!)$ is best for the occasion, on the other hand, will be determined by factors external to semantics, such as politeness, informativity and relevance (*e.g.,* a name may not be useful in all contexts). Choosing the most appropriate concept given a particular context is expected to be subject to pragmatic constraints, as discussed in Aloni (2001).

3.3.1.4 Concept variance: Explaining Greenberg's generalization

One intuitive result of this framework is that different methods of individuating an individual vary in how and when they are informative. The gradience displayed between sorts of concepts in $(\!|c|\!)$ allows us to explain why Greenberg's generalization, namely that answers to CQs are only individuals, holds in most cases, but also *why* it has the exceptions it does. Specifically, *names* are typically more informative concepts for individuating objects in D_e in CQ contexts since they denote the same object rigidly in every world, including the actual one. Other concepts, such as descriptions, can have different values at different worlds. As such, the relationship, though informative, that descriptions forge with the CQ NP is less strict than with names, as it is subject to contextual variance.

For example, again consider (114a), under the concealed question interpretation (114b).

(114) a. John knows the president of the LSA
b. John knows who the president of the LSA is

As claimed in the previous chapter, either one of the following continuations is acceptable.

(115) i. It's Aaron.
ii. It's the linguist who published the most articles in Language the year before.

To examine the description *the linguist who published the most articles in Language the year before* first, suppose that our agent John has multiple ways of possibly specifying the CQ NP *the president of the LSA*. For simplicity, suppose that there are only two linguists in D_e, Aaron and Baron who are named by the concepts c'' and c''', respectively. Let α be the concept denoted by the NP *the president of the LSA* and c be the description *the linguist who published the most articles in Language the year before*, and c' it's complement.

Let the extensions of the concepts be determined as in Table 3.3; assuming that names are constant functions, let c'' denote Aaron and c''' denote Baron in all possible worlds.

Worlds	α	**c**	c'	c''	c'''
$\mathbf{w}_o$	**a**	**a**	b	a	b
w	**a**	**a**	b	a	b
w$'$	**b**	**b**	a	a	b
w''	b	a	b	a	b

Table 3.3: Mappings from concepts to individuals

Further observe that the concepts constitute distinct conceptual covers, $\wp$ and $\wp'$:

(116) i. $\wp \ = \{c, c'\}$
ii. $\wp' = \{c'', c'''\}$

The effect that the CQ has on the context can be shown by how the two perspectives divide up logical space, as depicted in Figure 3.10. Here, different perspectives give different concepts that could individuate the denotation of α. Worlds are divided according to whether the concept in question and α pick out the same individual in that world as in a reference world w_0. The worlds in our scenario in the partitions are shown to the right of the partition itself.

under $\wp$	$\lambda u.c(w_0)(u) \doteqdot \alpha(v)(u)$	$= \{w_0, w, w'\}$
	$\lambda u.c'(w_0)(u) \doteqdot \alpha(v)(u)$	$= \{w''\}$

under $\wp'$	$\lambda u.c''(w_0)(u) \doteqdot \alpha(v)(u)$	$= \{w_0, w\}$
	$\lambda u.c'''(w_0)(u) \doteqdot \alpha(v)(u)$	$= \{w', w''\}$

Figure 3.10: Representation of CQ under $\wp$ and $\wp'$ at w_0

The crucial difference in the informativity between different sorts of concepts can be observed directly from the scenario above. If all that is known about α, in this case *the president of the LSA*, is that the extension of α is picked out by the description c in those same worlds, it is not guaranteed that the individual that α denotes is known as well, since the denotation c is dependent on the world of evaluation. For all that example (114), when interpreted in light of the continuation (115.ii), ensures is that c and α denote the same individual for the same worlds. It does *not* guarantee that the individual denoted by c in some non-actual world v will be the same individual denoted by α in the actual world w_o. As for our example, this discrepancy between actual and equivalent denotations is shown by the inclusion of world w' in the resulting partition by virtue of the fact that $c(w') = \alpha(w') = \mathbf{b}$, despite the fact that $\alpha(w_0) = \mathbf{a}$. That is, worlds in which α denotes a different individual are still included in the paritition. Thus, using a description to identify a CQ NP does not tell us what the NP is in the actual world, precisely because we can't necessarily tell which world is actual with that information alone.

In contrast, using a name to identify the denotation of α is more restrictive. Assuming that names are simply constant functions denoting the same individual in every accessible world, a name *does* guarantee that the answer will yield the individual that α denotes *in the actual world*. So names are *more informative* in these contexts because they ensure that we know the extension of the CQ NP in the actual world, whereas individuation by description does not. Nonetheless, individuation by description is still informative, and in certain cases, perhaps even more useful (see Aloni (2001) for examples of when a description or demonstrative rather than a name could be preferred.)

Indeed, the fact that the naming concept is more informative than a description in general can be gleaned directly from the example above. In particular, the block of worlds that the naming concept picks out is a proper subset of the block of worlds that the description does, as $\{w_o, w\} \subset \{w_o, w, w'\}$. Although it remains to be shown whether this relationship holds in general, it is clear that in the most restricted case – in which we have settled on what world is actual – and there is both a naming concept and a description in different covers, the name and the description are equally informative. If there are any worlds in which the extension of description and the name diverge and the description specifies the CQ NP in those worlds, then the name is again more restrictive than the description.

Importantly, the theory above explains precisely how individuation differs from identity, and provides general guidelines for when individuation sufficiently answers an identity question. Rescher (2005) suggests that answers to identity questions which merely individuate an individual are in someway inappropriate answer. Rescher (2005, 38) claims that we require "a *direct* answer, one that is presented in the standard way and not just some roundabout indication of the item at issue as satisfying some elaborate albeit pertinent description or other that leaves the matters of an informative identification still unresolved." For Rescher (2005) it appears that the central difference between a direct and an indirect answer is a 'condition of appropriateness.' Specifically, he argues that indirect answers merely *individuate* the relevant individual, without necessarily *identifying* her and are thereby unsatisfactory. The following passage is worth citing at length:

> It is one thing to offer a true response to a question, but something else – and more – actually to answer it directly. Without identification, [an agent] x does not have an informatively *appropriate* answer. The crucial difference here is between the actual identification of an item and its mere individuating specification that is unique to it. Thus "Don Juan, the famous rake" identifies the individual, but "the man – whoever he is – who fathered yon child" does not. For although this response is *individuating*, it does not *identify* the person at issue and so would not count as successfully answering the question "Whom did Pedro the Cruel most notably imprison for licentiousness?" Rescher (2005, 36)

In one sense, Rescher's comment is surely correct; to individuate an individual is not the same as identifying it, and when answering an identity question, a name is typically more appropriate. Yet, there are undoubtably contexts in which specification is sufficient, typically when the name of the individual is unavailable or when the answer to the question co-varies with another property or concept which is more easily discernable (*e.g., the price of orange juice* example). So we ought to be interested in explaining the contrast, not in merely explaining it away. The *reason* why names – direct answers – are *more appropriate* in neutral contexts is that they are typically more informative. Assuming the standard Gricean maxims, we prefer naming concepts as answers over less-direct methods of individuation because we are compelled to be as informative as possible. While much more could be said about how to formalize the pragmatics, we turn to the resolution of another outstanding issue, namely how to treat the entailment properties of concealed questions.

3.3.1.5 Conceptual covers and the entailment problem

Consider again Heim's entailment problem. In §3.1.1 it was argued that the fact that CQs do not validate syllogisms like (118) below is related to Quine's discussion of the failure of Leibniz's law in referentially opaque contexts. Given Aloni's solution to Quine's dilemma, we suggest a related way to dispel the entailment puzzle. As the dynamic quantifier $\mathscr{E}$ is not introduced until the next section, interpret $\mathscr{E}$ as the ordinary static quantifier $\exists$ below.

(117) I. John knows the capital of Italy
$\forall w \in Dox(w_0).\mathscr{E}c \in \wp : c(w)(w_0) \doteqdot \iota x.\text{capital-of-Italy}(x)(w)(w_0)$

II. The capital of Italy is the largest town in Italy

i. $\iota x.\text{capital-of-Italy}(x)(w_0) = \iota x.\text{largest-town-in-Italy}(x)(w_0)$

ii. $\iota x.\text{capital-of-Italy}(x)(w)(w_0) \doteqdot \iota x.\text{largest-town-in-Italy}(x)(w)(w_0)$

III. John knows the largest town in Italy
$\forall w \in Dox(w_0).\mathscr{E}c \in \wp : c(w)(w_0) \doteqdot \iota x.\text{largest-town-in-Italy}(x)(w)(w_0)$

Note that the second premise can be interpreted in one of two ways (118II.i–ii). In the first case, the second premise states that the two relevant concepts are equivalent in the actual world. The syllogism need not go through in this case, as the concepts may assign different values in John's beliefs worlds. In the second case, if $w \in Dox_J(w_0)$, then John knows the two concepts are equivalent in the relevant worlds, and so the syllogism holds.

Note that the first case is derivable from the second. Assuming that one description concept α is lifts to a CQ denotation, type $\langle\langle s,e\rangle,\langle s,\langle s,t\rangle\rangle\rangle$, and the other β is of type $\langle s,e\rangle$, then $\alpha(\beta)$ is a propositional concept: $\lambda w.\lambda w'.[\alpha(w) = \beta(w) \wedge \alpha(w') = \beta(w')]$. If this propositional concept is evaluated at w_0 for both world arguments, the resulting statement is $[\alpha(w_0) = \beta(w_0) \wedge \alpha(w_0) = \beta(w_0)]$, which classically reduces to $[\alpha(w_0) = \beta(w_0)]$. The second interpretation of the second premise, on the other hand, results if the propositional concept is evaluated two different worlds, resulting in $[\alpha(w) = \beta(w) \wedge \alpha(w_0) = \beta(w_0)]$, which reduces to $[\alpha(w)(w_0) \doteqdot \beta(w)(w_0)]$, by convention (111).

This analysis, incomplete though it may be, correctly predicts that a free relative could substitute for one or both concepts in the second premise:

(118) I. John knows the capital of Italy

II. i. What the capital of Italy is is the largest town in Italy

ii. What the capital of Italy is is what the largest town in Italy is

III. John knows the largest town in Italy

Although a complete analysis of this pattern is not attempted here, the relationship between CQ NPs and free relatives in pseudocleft constructions is discussed further in § 3.6.

Thus far, we have given two conceptions of concealed questions and a type shifter $\mathscr{Q}$ which lifts CQ NPs to relations between concepts and identity questions. The naturalness of

$\mathcal{Q}$ was located within the type-shifting framework of Partee (1986, 1995) and defended on conceptual grounds. Lastly, Greenberg's generalization and its exceptions were explained by illustrating the different effects that names and descriptions can have on logical space when they are fed into the CQ denotation. However, all of these results have been presented in the abstract. The next section works through a number of the core cases, concluding with a general solution to Heim's ambiguity that gives semantic content to the syntactic notion of Late Merger. However, since the solutions crucial rely on Dekker's (1993) analysis of relational nouns in the framework of Dynamic Montague Grammar (Groenendijk and Stokhof, 1990a), the requisite background must first be established. The next section presents the basics of Dynamic Montague Grammar (DMG) and Dekker's *existential disclosure* to represent the composition of relational nouns with their extrinsic arguments.

3.4 Dynamic Montague Grammar

Dynamic Montague Grammar (henceforth, DMG) is a wholly compositional dynamic treatment of Montague Grammar. It overcomes the limitations of Dynamic Predicate Logic (Groenendijk and Stokhof (1991)), a language which is restricted to a first-order language, by developing a system of logical dynamics that holds at the subsentential level. Thus, a dynamic treatment of variable binding is available to terms which better correspond to natural language objects than translations to a first-order system can approximate.

Although this section does not claim to provide a full treatment of the depth of DMG, it is hoped that enough basic mechanics and insights of the framework are presented to understand Dekker (1993a), and subsequently, our treatment of Heim's ambiguity. To that aim, the following section reviews the basic language of DMG along with its relation to ordinary, static meanings familiar from Montague Grammar (MG). Key to this section is the relationship between static and dynamic terms in the language, for DMG may be best understood perhaps not as a complete departure from MG, but as an insightful extension of it into the dynamic realm.

The first and foremost point to consider is how MG and DMG treat the basic concept of *meaning*. Montague Grammar, and its static descendants, follow standard modern logic in taking the extension of a sentence to be a *truth-value*, a type-theoretic object of type t, and the intension of a sentence as a proposition, an object of type $\langle s,t\rangle$. Thus, sentences correspond to either the True (1) or the False (0), on their extensional meanings, whose intensions are simply functions from worlds of evaluations to 1 or 0, or equivalently, a set of worlds for which the sentence holds true.

Dynamic systems such as DMG (but also Discourse Representation Theory (Kamp, 1981) and File Change Semantics (Heim, 1983)) take a different approach to meaning. In such systems, the meaning of a sentence is defined as the change its utterance has on the context, encapsulated in the oft-cited slogan "meaning is context change potential." In DMG in particular, it is said that sentence extensions are sets of propositions, objects of type $\langle\langle s,t\rangle,t\rangle$, while intensions are simply functions from states to extensions, $\langle s, \langle\langle s,t\rangle,t\rangle\rangle$. Given a state s, the context change potential (CCP) of a sentence is then just the set of propositions which are true of the state s' that obtains when the sentence updates the original state s. Table 3.4 provides a schematic comparison between the two notions of meaning.

	Static		Dynamic	
extension	truth-value	t	set of propositions	$\langle\langle s,t\rangle,t\rangle$
intension	proposition	$\langle s,t\rangle$	function to set of propositions	$\langle s,\langle\langle s,t\rangle,t\rangle\rangle$

Table 3.4: Static and dynamic views on sentence meanings compared

A word of clarification about the relation between *states* and *intensions* is in order. Groenendijk and Stokhof (1990a) are at pains to distinguish states, which merely provide 'parameters in the notion of interpretation', from intensional possibilia, at least in the traditional sense. Although we will follow Groenendijk and Stokhof (1990a) and speak strictly of states for the purposes of the present exposition, we will nonetheless use the framework developed here to model genuine intensional phenomena in the following sections, painting over any philosophical distinctions with the widest brush possible.

Yet, for all the genuine differences between static and dynamic frameworks, the two notions are in fact compatible. Groenendijk and Stokhof (1990a) show that the static notion of meaning may be derived from a dynamic one, as discussed in detail later in the section. Furthermore, not all terms of the DMG-framework are dynamic. DMG retains ordinary meanings for variables, quantifiers, and λ-terms, arguing that the dynamaticity originates in the passing and binding of special kinds of variables. As such, Groenendijk and Stokhof (1990a) introduce only two dynamic primitives into the syntax: (i) *discourse markers*, $d_1,\ldots,d_n$, which are special kinds of variables, and (ii) *state switchers*, which are special kinds of binding devices that allow the value of discourse markers to be passed to other terms. All other terms inherit a static meaning from MG.

Definition 3.4.1 (DIL Syntax)**.** Given CON_a, the set of constants, VAR_a the set of variables, and DM, the set of discourse markers, and the intensional language IL, for every type $a \in T2$, the set of meaningful expressions ME_a is defined as the smallest set generated by finite applications of 1–3:

1. If $\alpha \in ME_a$ of IL, then $\alpha \in ME_a$ of DIL
2. If $\alpha \in DM$, then $\alpha \in ME_a$
3. If $d \in DM, \alpha \in ME_e, \beta \in ME_a, then \{\alpha/d\}\beta \in ME_a$

The syntax of DIL simply extends that of the familiar IL with two kinds of objects that alone exhibit dynamic behavior. Discourse markers are constants whose valuation is dependent on a state parameter. Accordingly they are determined within the domain ${D_e}^S$, that of individual concepts. Clause 3 introduces the notion of a 'state switcher', $\{\alpha/d\}\beta$, a binding device that switches the state of interpretation of β from the original state s to a state s', which at most differs from s in that d denotes the same object in s' that α does in s. Letting F be the dynamic interpretation function, the following two postulates ensure that there is a unique s' with the intended property:

Postulate 3.4.1 (Distinctness). *If for all $d \in DM : F(d)(s) = F(d)(s')$, then $s = s'$*

Postulate 3.4.2 (Update). *For all $s \in S, d \in DM, \mathbf{d} \in D$ there exists an $s' \in S$ such that:*

1. *$F(d)(s') = \mathbf{d}$; and*
2. *for all $d' \in DM, d' \neq d : F(d')(s) = F(d')(s')$*

The unique state s' differing from s only on the assignment of the discourse marker d is referred to as $\langle d \leftarrow [[\alpha]]_{M,s,g}\rangle s$. The semantics of DIL are spelled out according to the previous informal discussion as follows:

Definition 3.4.2 (Semantics of DIL). For a model $M = \langle F, S, D\rangle$, where g is an assignment function in the set of assignment functions G in F, the semantics of a well-formed expression α is defined as $[[\alpha]]_{M,s,g}$ such that:

1. All terms α in IL are interpreted $[[\alpha]]_{M,s,g}$ as before
2. $[[d]]_{M,s,g} = F(d)(s)$, for every discourse marker d
3. $[[\{\alpha/d\}\beta]]_{M,s,g} = [[\beta]]_{M,\langle d \leftarrow [[\alpha]]_{M,s,g}\rangle s,g}$

As the following facts of state switchers are used in derivations that follow, note that the following equivalences hold:[9]

Fact 3.4.1. *[Properties of state switchers]*

1. *$\{\alpha/d\}\mu \equiv \mu$ for any $\mu \in CON \cup VAR$.*
2. *$\{\alpha/d\}d \equiv \alpha$*
   $\{\alpha/d\}d' \equiv d'$ for all d' distinct from d in DM.
3. *$\{\alpha/d\}(\beta(\gamma)) \equiv \{\alpha/d\}\beta(\{\alpha/d\}\gamma)$*
4. *$\{\alpha/d\}(\phi \oplus \psi) \equiv \{\alpha/d\}\phi \oplus \{\alpha/d\}\psi$ for any connective or the identity symbol $\oplus$*
5. *$\{\alpha/d\}\exists v\phi \equiv \exists v\{\alpha/d\}\phi$, if v is not free in α*

[9]Some facts which were presented in the original article are omitted here for considerations of space. The reader should consult Groenendijk and Stokhof (1990a) or Dekker (1993b) for a more complete exposition.

6. $\{\alpha/d\}\lambda v\beta \equiv \lambda v\{\alpha/d\}\beta$, *if* v *is not free in* α

7. $\{\alpha/d\}^\wedge\beta \equiv {}^\wedge\beta$
   $\{\alpha/d\}^\vee\beta \equiv {}^\vee\beta$ *does* not *hold in the general case*

Briefly, clauses 3 and 4 state that state switchers can be 'pushed' inside an expression to apply past function-argument structure, the logical connectives, and identity, and in limited cases past quantification and λ abstraction, as well. The second part of clause 7 states that $\{\alpha/d\}^\vee\beta$ cannot be reduced any further. Note that a set of propositions $\lambda p.^\vee p$ denotes the set of true propositions in a state s in which it is evaluated. When a state switcher is applied to this expression the result $\{\alpha/d\}\lambda p.^\vee p$ is equivalent with $\lambda p.\{\alpha/d\}^\vee p$ by clause 6, but is irreducible by clause 7. With respect to some state s, the expression is interpreted $[[\lambda p.\{\alpha/d\}^\vee p]]_{M,g,s}$ as the function $h(p) = [[^\vee p]]_{M,\langle d\leftarrow [[\alpha]]_{M,s,g}\rangle s,g}$; the set of all propositions true in the unique state $\langle d \leftarrow [[\alpha]]_{M,s,g}\rangle s$ that at most differs from s in that d has the denotation in s' that α does in s.

Recall that DMG takes the meaning of a sentence to be its CCP. So, given a context for evaluation s, a sentence is just the set of propositions **p** that are true after s is processed with respect to that sentence. If the processing succeeds then **p** is not empty and contains the tautologous proposition $^\wedge\mathtt{true}$, interpreted as 'is a possible state.' Applying this proposition to the state which obtained after processing the sentence, the result is a truth value t. This in turn reflects the ordinary static interpretation of the dynamic meaning.

For example, assume that the sentence *John walks* is translated into DMG as $\lambda p.[walks(j) \wedge {}^\vee p]$, a set of propositions. When the tautologous proposition ($^\wedge\mathtt{true}$) is applied to it, the result is simply $[walk(j) \wedge {}^{\vee\wedge}\mathtt{true}]$, which in turn reduces by Fact 3.4.2 ($^{\vee\wedge}$-elimination) and clause 1 to the classical statement $[walk(j)]$, an object of type t. Thusly is DMG able to retain a static conception of meaning as a derivative of a dynamic one.

Fact 3.4.2. *[$^{\vee\wedge}$-elimination]* $^{\vee\wedge}\alpha$ *is equivalent to* α

Groenendijk and Stokhof (1990a) formalize the correspondence between static and dynamic notions of meaning as operators in the syntax. The first operation $\uparrow$ *lifts* the static interpretation of $\phi \in ME_t$ to a dynamic counterpart. The second $\downarrow$ does just the opposite, by *lowering* a dynamic meaning to a static one. The formal definitions are as follows.

Definition 3.4.3. [Uparrow] $\uparrow \phi = \lambda p[\phi \wedge {}^\vee p]$, where ϕ is an expression of type t, p a variable of type $\langle s,t\rangle$ in which ϕ does not occur free

Definition 3.4.4. [Downarrow] $\downarrow \phi = \phi(^\wedge\mathtt{true})$, where ϕ is an expression of type $\langle\langle s,t\rangle,t\rangle$

Some additional facts regarding the *lift* and *lower* operations will be necessary. When the lower and lift operations apply to a term, in that order, they cancel each other out.[10] As shown in Fact 3.4.3, the lift operator distributes over and out of conjunction.

[10] As suggested, the order of application matters, as $\uparrow\downarrow \phi$ is not, in general, equivalent with ϕ. For although $\downarrow \lambda p[\phi \wedge \{\alpha/d\}^\vee p]$ reduces to ϕ, the expression ϕ when it itself is lifted $\uparrow \phi$ does not yield $\lambda p[\phi \wedge \{\alpha/d\}^\vee p]$, but rather $\lambda p[\phi \wedge {}^\vee p]$ instead.

Fact 3.4.3. *DMG Reduction*[11]

1. *$\uparrow$-export*

   (a) $(\uparrow\beta)(\alpha) \Leftrightarrow \uparrow(\beta(\downarrow\alpha))$

   (b) $\alpha \cong \beta \Leftrightarrow \uparrow(\downarrow\alpha = \downarrow\beta)$

2. *$\downarrow$-import*

   (a) $\downarrow\uparrow\phi = \phi$

   (b) $[\uparrow\phi; \uparrow\psi] = \uparrow[\phi \wedge \psi]$

   (c) $\downarrow\mathscr{E}d\phi \Leftrightarrow \exists d\downarrow\phi$

Although there is much more to DMG than presented in this section, we will conclude the exposition by introducing two more definitions and with a discussion of sentential conjunction. The first of these notions is the dynamic existential quantifier, $\mathscr{E}$, which plays a central role in the sections that follow. Here, the dynamic existential quantifier retains the dynamic character of the system by returning the set of propositions for which hold true when ϕ is evaluated in the state s' in which d has the same denotation as x in s.

Definition 3.4.5. [Dynamic existential quantifier] $\mathscr{E}d\phi = \lambda p\exists x\{x/d\}(\phi(p))$, where x and p have no free occurrences in ϕ

The second notion is just as important, that of dynamic conjunction, defined as intensional function composition. Dynamic conjunction takes the extension of the second sentence ψ (of type $\langle\langle s,t\rangle,t\rangle$) which applies to an arbitrary proposition $\langle s,t\rangle$. The intension of this result is a proposition which in turn composes with the extension of the first conjunct ϕ, and again yields a truth value t. Abstracting over the arbitrary proposition p provides an object of type $\langle\langle s,t\rangle,t\rangle$, an ordinary dynamic meaning. Thus in effect, we get a generalized quantifier over states for which both ϕ and ψ are true in those states.

Definition 3.4.6. [Dynamic conjunction] $\phi;\psi = \lambda p[\phi(^\wedge(\psi(p)))]$, where p has no free occurrences in either ϕ or ψ

As usual, dynamic conjunction is genuinely sequential, as it is associative but not commutative.

[11]These facts are modeled after the presentation in Dekker (1993a), not Groenendijk and Stokhof (1990a), although they hold of course in both systems.

Fact 3.4.4.

1. $[\phi;\psi];\chi$ *is equivalent to* $\phi;[\psi;\chi]$
2. $\phi;\psi$ *is not equivalent to* $\psi;\phi$

Lastly, DMG inherits the familiar dynamic property of existential binding in that the existential quantifier can bind variables outside of its syntactic scope.

Fact 3.4.5. $\mathcal{E}d\phi;\psi$ *is equivalent to* $\mathcal{E}d[\phi;\psi]$

Having provided the necessary notions to work within DMG, a simple example suffices to illustrate the mechanics and formal elegance of DMG. Assuming the sample vocabulary and their corresponding DMG translations list in (119), a full derivation is provided in (120).

(119) VOCABULARY AND TRANSLATIONS:

i. $\texttt{a}_i \rightsquigarrow \lambda P.\lambda Q.\mathcal{E}d_i[P(\uparrow d_i);Q(\uparrow d_i))]$

ii. $\texttt{man} \rightsquigarrow \uparrow man$

iii. $\texttt{walks} \rightsquigarrow \uparrow walks$

(120) A man walks

a. $[[a]]([[man]]) =$
$[\lambda P.\lambda Q.\mathcal{E}d_1[P(d_1);Q(d_1))]](\uparrow man) =$
$\lambda Q.\mathcal{E}d_1[\uparrow man(d_1);Q(\uparrow d_1))]$

b. $[[\text{a man}]]([[walks]]) =$
$[\lambda Q.\mathcal{E}d_1[\uparrow man(d_1);Q(d_1))]](\uparrow walks) =$
$\mathcal{E}d_1[\uparrow man(d_1);\uparrow walks(d_1))]$

i. From DMG translation to CCP:

$\mathcal{E}d_1[\uparrow man(d_1);\uparrow walks(d_1))]$	$\Leftrightarrow$	(by Fact 3.4.3:2.b)
$\mathcal{E}d_1\uparrow[man(d_1)\wedge walks(d_1))]$	$\Leftrightarrow$	(by Def 3.4.5)
$\lambda p\exists x[\{x/d_1\}(\uparrow[man(d_1)\wedge walks(d_1)](p))]$	$\Leftrightarrow$	(by Def 3.4.3 and λ-conversion)
$\lambda p\exists x[\{x/d_1\}[man(d_1)\wedge walks(d_1)\wedge {}^{\vee}p]]$	$\Leftrightarrow$	(by Fact 3.4.1)
$\lambda p\exists x[man(x)\wedge walks(x)\wedge\{x/d_1\}^{\vee}p]$		

ii. From CCP to classical truth: apply $\downarrow$

$\downarrow\lambda p\exists x[man(x)\wedge walks(x)\wedge\{x/d_1\}^{\vee}p]$	$\Leftrightarrow$	(by Def 3.4.4)
$\lambda p\exists x[man(x)\wedge walks(x)\wedge\{x/d_1\}^{\wedge}p](^{\wedge}\texttt{true})$	$\Leftrightarrow$	(by λ-conversion)
$\exists x[man(x)\wedge walks(x)\wedge\{x/d_1\}^{\vee\wedge}\texttt{true})]$	$\Leftrightarrow$	(by Fact 3.4.2, $^{\wedge\vee}$-elimination)
$\exists x[man(x)\wedge walks(x)\wedge\{x/d_1\}\texttt{true}]$	$\Leftrightarrow$	(by Fact 3.4.1:1)
$\exists x[man(x)\wedge walks(x)\wedge\texttt{true}]$	$\Leftrightarrow$	(by classical validity)
$\exists x[man(x)\wedge walks(x)]$		

The derivation in (120) divides the procedure into several steps. Parts (120a–b) only involve the combination of terms via function-application. Part (120.i) shows how to interpret the translation of *A man walks* as a context change potential (CCP). Part (120.ii) further shows how the CCP is transformed to an ordinary classical meaning via lowering the CCP. Since pronouns work in essentially the same way, I will not belabor the discussion with another illustration. The interested reader may find a more complete discussion in Groenendijk and Stokhof (1990a).

To review, we have presented a basic outline of Dynamic Montague Grammer. The relationship between Montague Grammar and competing dynamic systems was briefly discussed. Further, a short derivation served to illustrate the formal mechanisms of DMG. We now investigate an extension of the above system to relational nouns, as developed in Dekker (1993a).

3.4.1 Existential Disclosure

Dekker (1993a, 1993b) develops the notion of 'existential disclosure' in order to treat implicit arguments in a fashion much like the dynamic treatment of existentially bound discourse markers (or variables, if in DPL) and pronouns. Generally speaking, implicit arguments have traditionally been analyzed as syntactic arguments of a variety of syntactic categories which may or may not be overtly present in the syntax (see Bhatt and Pancheva (2006) for an overview). A relevant example is the now familiar category of relational nouns, such as *price* or *president*, the denotation of which we understand as a relation to an extrinsic argument, such as *milk* or *the LSA*, respectively. That is, a *price* is thought to always be a *price of something*, even if the argument is not overtly specified in syntax. Indeed, Dekker's system treats a variety of implicit arguments, including event arguments in adverbial modification and the transfer of tense in discourse, not merely those in relational nouns. However, to pursue the full treatment of such phenomena would lead us far astray from the central purpose at hand.

Implicit arguments are treated as existentially closed arguments, which under a static semantics are semantically unavailable for further specification, such as pronominal reference or various forms of modification.[12] In contrast, the hallmark of dynamic theories is the possibility that indefinites and pronouns can be further specified by constituents appearing later in discourse, as seen above.

Within a modified DMG framework, Dekker (1993a) makes the important distinction between implicit arguments that are *dynamically* closed and those that are *statically closed*. Dynamically closed implicit arguments are able to be further modified; they behave essen-

[12] To avoid confusion, note that static theories are not unable to cope with implicit arguments. Instead of passing semantic objects directly in the semantics for valuation, syntactic rules are responsible for establishing the availability of anaphoric reference. The systems for variable transfer are entirely different between the two types of systems, at least as developed so far, and so a direct comparison on this point is not completely possible, at least not without simultaneously addressing the inherent differences and attendant tradeoffs between the two conceptions.

tially like indefinite objects. Statically closed arguments either are variables of the wrong ontological type, *i.e.,* they are not discourse markers, or have been specified *explicitly*, as in *the price of milk*. As such, statically closed arguments cannot be further modified; they are dynamically inert.

Informally, Dekker (1993b) is able to treat implicit arguments as dynamically closed by introducing the notion of *existential disclosure*, an operation which essentially reverses the effects of *existential closure*, when a free variable is bound off by an existential quantifier. More specifically, existential disclosure frees a discourse marker d in a formula ϕ containing an 'active' occurrence of dynamic existential quantification over d, by conjoining an assertion of the identity of d and a free variable x to ϕ, as in $[\phi;(\uparrow x \cong \uparrow d)]$, the notation for which is written as $\{\uparrow x/d\}\phi$. Note that we now decorate discourse markers d with the lift symbol $\uparrow$, though the meaning is the same as before.

To see that existential disclosure does indeed reverse the effects of the familiar existential binding, I provide an abstract reduction of an existentially disclosed dynamic term to its static representation in (121):

(121) Existential Disclosure.

$\downarrow\{\uparrow x/d\}\mathcal{E}d[\uparrow \zeta(\uparrow d)]$	$\Leftrightarrow$	(by definition)
$\downarrow\mathcal{E}d[\uparrow \zeta(\uparrow d);(\uparrow x \cong\uparrow d)]$	$\Leftrightarrow$	(by $\uparrow$-export; Fact 3.4.3:2.a–b)
$\downarrow\mathcal{E}d[\uparrow \zeta(\downarrow\uparrow d);\uparrow(\downarrow\uparrow x \cong\downarrow\uparrow d)]$	$\Leftrightarrow$	(by $\downarrow$-import; Fact 3.4.3:2.c)
$\exists d[\zeta(d);\uparrow(x=d)]$	$\Leftrightarrow$	(classically valid)
$[\zeta(x)]$		

The discourse marker $\uparrow d$ that was dynamically closed in $\mathcal{E}d[\uparrow \zeta(\uparrow d)]$ is identified with (the lift of) a free variable x, which makes d indirectly available for further modification.

The representation of a relational noun ζ, is simply the DMG translation of ζ with a dynamically closed discourse marker d in place of an explicit argument: $\lambda x_{\langle s,e\rangle}\mathcal{E}d[\uparrow \zeta(\uparrow d)(x)]$. To be concrete, an unsaturated relational noun, *price*, is represented as $\lambda x_{\langle s,e\rangle}\mathcal{E}d[\uparrow \text{price}(\uparrow d)(x)]$, namely as *a price of something*.

An important piece of this analysis is that relational nouns can often be used with or without specification of the extrinsic argument. Therefore, it is important that whatever representation we have for 'unsaturated' relational nouns be compatible with the 'saturated' variants. In order to account for this distribution, Dekker introduces a prepositional phrase, *of*$_2$, which opens the dynamically closed discourse marker in the extrinsic position of the relational noun and discloses it for specification. When combined with an explicit argument, the result is a fully specified relational noun which is *statically* closed. Let ϵ be the type of an individual concept, $\langle s,e\rangle$, and τ the type of a CCP, $\langle s,\langle\langle s,t\rangle,t\rangle\rangle$.

(122) Vocabulary and Translations: Let P, Q be variables of type $\langle\epsilon,\tau\rangle$, and T a variable of type $\langle\langle\epsilon,\tau\rangle,\tau\rangle$.

i. the $\rightsquigarrow \lambda P.\iota x.P(x)$

ii. price $\rightsquigarrow \lambda x_{\langle s,e\rangle}\mathscr{E}d[\uparrow \mathtt{price}(\uparrow d)(x)]$

iii. milk $\rightsquigarrow \lambda Q.Q(\uparrow \mathtt{milk})$

iv. $\mathrm{of}_2 \rightsquigarrow \lambda T.\lambda P.\lambda x.\uparrow\downarrow T(\lambda y\{y/d\}P(x))$

(123) Derivation of *the price of milk*

a. $[[of_2]]([[milk]]) =$
$[\lambda T.\lambda P.\lambda x.\uparrow\downarrow T(\lambda y\{y/d\}P(x))](\lambda Q.Q(\uparrow \mathtt{milk})) =$
$\lambda P.\lambda x.\uparrow\downarrow [\lambda Q.Q(\uparrow \mathtt{milk})](\lambda y\{y/d\}P(x)) =$
$\lambda P.\lambda x.\uparrow\downarrow [\lambda y\{y/d\}P(x)](\uparrow \mathtt{milk}) =$
$\lambda P.\lambda x.\uparrow\downarrow \{\uparrow \mathtt{milk}/d\}P(x)$

b. $[[\mathrm{of}_2 \text{ milk}]]([[\text{price}]]) =$
$[\lambda P.\lambda x.\uparrow\downarrow \{\uparrow \mathtt{milk}/d\}P(x)](\lambda x_{\langle s,e\rangle}\mathscr{E}d[\uparrow \mathtt{price}(\uparrow d)(x)]) =$
$\lambda x.\uparrow\downarrow \{\uparrow \mathtt{milk}/d\}[\lambda x_{\langle s,e\rangle}\mathscr{E}d[\uparrow \mathtt{price}(\uparrow d)(x)]](x)] =$
$\lambda x.\uparrow\downarrow \{\uparrow \mathtt{milk}/d\}\mathscr{E}d[\uparrow \mathtt{price}(\uparrow d)(x)]$

i. Reduction:

$\lambda x.\uparrow\downarrow \{\uparrow \mathtt{milk}/d\}\mathscr{E}d[\uparrow \mathtt{price}(\uparrow d)(x)]$	$\Leftrightarrow$	(by definition)
$\lambda x.\uparrow\downarrow [\mathscr{E}d \uparrow \mathtt{price}(\uparrow d)(x); \uparrow milk \cong \uparrow d]$	$\Leftrightarrow$	(by associativity)
$\lambda x.\uparrow\downarrow \mathscr{E}d[\uparrow \mathtt{price}(\uparrow d)(x); \uparrow milk \cong \uparrow d]$	$\Leftrightarrow$	(by ↑-export; Fact 3.4.3:1.b)
$\lambda x.\uparrow\downarrow \mathscr{E}d[\uparrow \mathtt{price}(d)(\downarrow x); \uparrow(\downarrow\uparrow milk = \downarrow\uparrow d)]$	$\Leftrightarrow$	(by ↑-import; Fact 3.4.3:2.c)
$\lambda x.\uparrow \exists d(\mathtt{price}(d)(\downarrow x) \wedge milk = d)$	$\Leftrightarrow$	(classically valid)
$\lambda x.\uparrow(\mathtt{price}(milk)(\downarrow x))$		

c. $[[\text{the}]]([[\text{price } \mathrm{of}_2 \text{ milk}]]) =$
$[\lambda P.\iota x.P(x)](\lambda x.\uparrow(\mathtt{price}(milk)(\downarrow x))) =$
$\iota x.[\lambda x.\uparrow(\mathtt{price}(milk)(\downarrow x))(x)] =$
$\iota x.\uparrow(\mathtt{price}(milk)(\downarrow x))$

In the above derivation, parts (123a–c) involve only composition of the terms via function application, while subpart (123b.i) shows the saturation of the implicit argument of the relation noun through a series of reductions. It is important to observe again that the preposition *of*$_2$ involves the static closure ↑↓ of its arguments, thereby disallowing any further modification, such as additional applications of the preposition *of*$_2$, or pronominal reference to any of the subparts that exclude the term in its entirety.

Since a composition of *the price* is more transparent than the derivation above, save that there is no application of *of*$_2$, forcing existential disclosure a complete derivation is foregone here. The result, as the reader may confirm, is simply the application of the terms *the* and *price*, resulting in $\iota x_{\langle s,e\rangle}\mathscr{E}d \uparrow \mathtt{price}(\uparrow d)(x)$. Note that the implicit argument is *still* available for further modification and binding. Now that a sufficient grounding of DMG and Existential

Disclosure has been obtained, we now exploit these systems in a completely formal treatment of concealed questions in the next section.

3.5 Deriving CQs: the mechanics

The following sections first present the basic CQ representation for relational nouns with and without overt arguments. We then extend the account to address Heim's ambiguity, and in so doing give a semantics for Late Merger that exploits the unique variable binding system available in DMG.

3.5.1 The standard case

Vocabulary and translations of terms are provided in (124), followed by the compositional derivation of a typical sentence containing a concealed question. Note that the non-CQ interpretation *the price of milk* is given simpliciter, as its derivation was previously illustrated in (123).

(124) Vocabulary and translations:

i. [[the price]] $= \iota x.\mathscr{E}d \uparrow price(\uparrow d)(x)$

ii. [[the price of$_2$ milk]] $= \iota x.\uparrow(price(milk)(\downarrow x(w)))$

iii. $\mathscr{Q}($ [[α]]$) = \lambda c.\lambda w.\lambda w' : c \in \wp[c(w) = \alpha(w) \wedge c(w') = \alpha(w')]$

iv. [[knows$_{CQ}$]] $= \lambda\vartheta_{\langle\langle s,e\rangle,\langle s,\langle s,t\rangle\rangle\rangle}.\lambda x.\forall w \in Dox_x(w_o).\mathscr{E}c.\vartheta(c)(w)(w_o)$

v. [[$John$]] $= j$

(125) John knows the price of milk

a. $\mathscr{Q}($ [[the price of$_2$ milk]]$) =$
$\lambda c.\lambda w.\lambda w' : c \in \wp[\,c(w) = \iota x.\uparrow(price(milk)(\downarrow x))(w) \wedge$
$c(w') = \iota x.\uparrow(price(milk)(\downarrow x))(w')] =$
$\lambda c.\lambda w.\lambda w' : c \in \wp[\,c(w) = \iota x.\uparrow(price(milk)(\downarrow x(w))) \wedge$
$c(w') = \iota x.\uparrow(price(milk)(\downarrow x(w')))] =$
$\lambda c.\lambda w.\lambda w' : c \in \wp[c(w)(w') \doteqdot \iota x.\uparrow(price(milk)(\downarrow x(w)(w')))]$

b. [[know$_{CQ}$]]($\mathscr{Q}($ [[the price of$_2$ milk]])) $=$
$[\lambda\vartheta_{\langle\langle s,e\rangle,\langle s,\langle s,t\rangle\rangle\rangle}.\lambda x.\forall w \in Dox_x(w_o).\mathscr{E}c.\vartheta(c)(w)(w_o)]$
$(\lambda c.\lambda w.\lambda w' : c \in \wp[c(w)(w') \doteqdot \iota x.\uparrow(price(milk)(\downarrow x(w)(w')))]) =$
$\lambda x.\forall w \in Dox_x(w_o).\mathscr{E}c.[\lambda c.\lambda w.\lambda w' : c \in \wp[c(w)(w') \doteqdot \iota x.\uparrow(price(milk)(\downarrow x(w)(w')))]](c)(w)(w_o)] =$
$\lambda x.\forall w \in Dox_x(w_o).\mathscr{E}c : c \in \wp[c(w)(w_o) \doteqdot \iota x.\uparrow(price(milk)(\downarrow x(w)(w_o)))] =$

c. [[know$_{CQ}$ $\mathscr{Q}$(the price of$_2$ milk)]]([[John]]) $=$
$[\lambda x.\forall w \in Dox_x(w_o).\mathscr{E}c : c \in \wp[c(w)(w_o) \doteqdot \iota x.\uparrow(price(milk)(\downarrow x(w)(w_o)))]](j) =$
$\boxed{\forall w \in Dox_j(w_o).\mathscr{E}c : c \in \wp[c(w)(w_o) \doteqdot \iota x.\uparrow(price(milk)(\downarrow x(w)(w_o)))]}$

As witnessed by (125), the computation of the sentence *John knows the price of milk* requires nothing beyond Dekker's representation of relational nouns, the type shifter $\mathscr{Q}$, and standard functional application. Further, the result of the computation, as it is framed here gives a static meaning, that is, a sentence of type t. To further incorporate the sentence into discourse, all that is required is an application of the lift operation $\uparrow$. Lastly, note that the individuating concept is dynamically closed, but that the arguments in the price-concept are statically closed. This allows anaphora to refer back to the individuating concept, but not any argument of the price-concept itself.

This system, then, makes a clear prediction about the sorts of possible anaphoric bindings that are predicted to occur, namely that in CQs with fully specified relational nouns, only the concept that individuates the CQ noun is available for future binding or modification. This prediction appears to be borne out:

(126) John knows the highest price of milk in Albert Heijn.

a. It is $1.29

b. # It's organic.

The proposal also predicts that relational nouns *without* specified arguments *do* make these arguments available for future binding. Although (127b) is somewhat degraded, it is in fact much better than its counterpart with a overtly specified argument (126b). It is arguable whether the distribution in (127) is in accordance with this prediction.

(127) John knows the highest price in Albert Heijn.

a. It is $1.29

b. ? It's milk.

Although much more needs to be said about the above contrast, it provides initial evidence in favor of the current theory. The exact derivation of *John knows the price* is presented below, which assumes the same vocabulary (124) from the above derivation. From (128), it is clear that *the price*, under a CQ interpretation, contains two discourse markers which may be picked up by pronouns later in discourse, as is consistent with (127).

(128) John knows the price

a. $\mathscr{Q}$([[the price]]) =

$\lambda c.\lambda w.\lambda w' : c \in \wp[c(w) = \iota x.\mathscr{E}d \uparrow price(\uparrow d)(x(w)) \wedge c(w') = \iota x.\mathscr{E}d \uparrow price(\uparrow d)(x(w'))] =$

$\lambda c.\lambda w.\lambda w' : c \in \wp[c(w)(w') \doteqdot \iota x.\mathscr{E}d \uparrow price(\uparrow d)(x(w)(w'))$

b. $[[know_{CQ}]](\mathscr{Q}(\text{the price})) =$

$[\lambda\vartheta_{\langle\langle s,e\rangle,\langle s,\langle s,t\rangle\rangle\rangle}.\lambda x.\forall w \in Dox_x(w_o).\mathscr{E}c.\vartheta(c)(w)(w_o)](\lambda c.\lambda w.\lambda w' : c \in \wp[c(w)(w') \doteqdot$
$\iota x.\mathscr{E}d \uparrow price(\uparrow d)(x(w)(w'))) =$

$\lambda x.\forall w \in Dox_x(w_o).\mathscr{E}c.[\lambda c.\lambda w.\lambda w' : c \in \wp[c(w)(w') \doteqdot \iota x.\mathscr{E}d \uparrow price(\uparrow d)(x(w)(w'))](c)(w)(w_o) =$

$\lambda x.\forall w \in Dox_x(w_o).\mathscr{E}c : c \in \wp[c(w)(w_o) \doteqdot \iota x.\mathscr{E}d \uparrow price(\uparrow d)(x(w)(w_o))]$

c. [[know$_{CQ}$ $\mathcal{Q}$(the price)]]([[John]]) =
$[\lambda x.\forall w \in Dox_x(w_o).\mathscr{E}c : c \in \wp[c(w)(w_o) \doteqdot \iota x.\mathscr{E}d \uparrow price(\uparrow d)(x(w)(w_o))]](j) =$
$\boxed{\forall w \in Dox_j(w_o).\mathscr{E}c : c \in \wp[c(w)(w_o) \doteqdot \iota x.\mathscr{E}d \uparrow price(\uparrow d)(x(w)(w_o))]}$

Having laid sufficient groundwork for the formal aspects of this theory, we now turn to Heim's ambiguity. In this section, we articulate a semantic rule to account for the effects of Late Merger. The strategy is as follows: the matrix verb and the CQ NP are combined together before the relative clause can compose. Thus, reading A is formed exactly as (128) was, and the relative clause is argued to compose with the sentence *John knows the price* via sentence level conjunction and anaphoric binding. In contrast, reading B is composed by transparent function application and the matrix verb takes the intension of (128) as its complement. The formal details are developed immediately below.

3.5.2 Deriving Heim's Ambiguity

In the comparison of proposals for CQs, Heim's ambiguity may have been the most influential test of the success of an account. The present account intends to capture structural facts that other accounts have not addressed. Recall that in §2.1.2, a number of arguments for diversifying the syntactic structure of reading A and B were presented. There, it was claimed that reading A contained an externally headed relative clause which was introduced late into the syntactic derivation, *i.e.,* was 'Late Merged' in the sense of Lebeaux (1990). The primary evidence for this analysis was the fact that Reading A survives only ACD contexts, whereby the embedded verb of a relative clause elides (Fox, 2002). In contrast, Reading B was argued to contain an internally headed relative clause which was introduced immediately into the syntactic derivation. The primary basis for this argument was that Reading B was vastly preferred in sluicing contexts.

We now give a semantics for the structural claims made previously. In particular, it is argued that the semantic content of the relative clause in Reading A is made available *only after* the matrix verb and the nominal complement have composed. Late merger, then, is analyzed as *sentential modification* which is licensed only when it can modify a variable made available by the dynamic passing of discourse objects. As such, the content of the relative clause is represented *outside* the beliefs of the matrix subject. This accords with the paraphrase of Reading A. Reading B is treated as a simple propositional attitude. The intension of the phrase *the price that Fred know* is a proposition which is represented as part and parcel of the subject's belief states. Again, this reading accords with the interpretation of Reading B discussed in the previous chapter.

3.5.2.1 Reading A as sentential modification

Turning to Reading A of Heim's ambiguity, we define the semantics of Late Merger as a semantic operation Λ which conjoins two sentences and unites the dynamically bound variables therein.

Definition 3.5.1 (Semantics of Late Merger Λ). Let ϕ and ψ be (dynamic) sentences of type τ, which contain active dynamic quantification over discourse markers d_1 and d_2, respectively. Late merger is the dynamic conjunction of $\{\uparrow x/d_1\}\phi$ and $\{\uparrow x/d_2\}\psi$:

$$\Lambda = \lambda\phi.\lambda\psi[\{\uparrow x/d_1\}\phi\,;\{\uparrow x/d_2\}\psi]$$

(129) Schema for late merger:

a. Reduction:

$$\begin{array}{lll}
\Lambda(\mathscr{E}d_1 \uparrow \phi(\uparrow d_1))(\mathscr{E}d_2 \uparrow \phi(\uparrow d_2)) & \Leftrightarrow & \text{(by definition)} \\
[\{\uparrow x/d_2\}\mathscr{E}d_1 \uparrow \phi(\uparrow d_1);\{\uparrow x/d_2\}\mathscr{E}d_2 \uparrow \psi(\uparrow d_2)] & \Leftrightarrow & \text{(by definition)} \\
[\mathscr{E}d_1 \uparrow \phi(\uparrow d_1);\uparrow x \cong \uparrow d_1;\mathscr{E}d_2 \uparrow \psi(\uparrow d_2);\uparrow x \cong \uparrow d_2] & \Leftrightarrow & \text{(by }\uparrow\text{-export)} \\
[\mathscr{E}d_1 \uparrow (\phi(\uparrow\downarrow d_1));\uparrow(\downarrow\uparrow x = \downarrow\uparrow d_1);\mathscr{E}d_2 \uparrow (\psi(\downarrow\uparrow d_2);\uparrow(\downarrow\uparrow x = \downarrow\uparrow d_2)] & \Leftrightarrow & \text{(by }\downarrow\text{-import)} \\
[\mathscr{E}d_1 \uparrow (\phi(d_1));\uparrow(x = d_1);\mathscr{E}d_2 \uparrow (\psi(d_2);\uparrow(x = d_2)] & \Leftrightarrow & \text{(by Fact 3.4.3)} \\
[\mathscr{E}d_1 \uparrow (\phi(d_1) \wedge (x = d_1));\mathscr{E}d_2 \uparrow (\psi(d_2) \wedge (x = d_2))] & &
\end{array}$$

b. Closure:

$$\begin{array}{lll}
\downarrow[\mathscr{E}d_1 \uparrow (\phi(d_1) \wedge (x = d_1));\mathscr{E}d_2 \uparrow (\psi(d_2) \wedge (x = d_2))] & \Leftrightarrow & \text{(by }\downarrow\text{-import)} \\
[\exists d_1(\phi(d_1) \wedge (x = d_1)) \wedge \exists d_2(\psi(d_2) \wedge (x = d_2))] & \Leftrightarrow & \text{(classically valid)} \\
\phi(x) \wedge \psi(x) & &
\end{array}$$

I assume that the relative clause is externally headed, as in (50). In keeping with this structure, I assume that the type-shifting operation $\mathscr{Q}$ combines with a trace-copy (a) which is later abstracted over (d). Further, I assume that the internal head of the structure, which is phonologically null, combines via functional application with the relative clause (e), so that the relative clause is of a sentential type at the end of the derivation.

(130) Formation of the externally headed relative clause:
~~*the price*~~ *that Fred knows* ~~*the price*~~

a. $\mathscr{Q}(t_1) =$
$\lambda c.\lambda w.\lambda w'.c \in \wp : [c(w)(w') \dotdiv t_1(w)(w')]$

b. $[[\text{know}_{CQ}]](\mathscr{Q}(\text{t}_1)) =$
$[\lambda\vartheta_{\langle\langle s,e\rangle,\langle s,\langle s,t\rangle\rangle\rangle}.\lambda x.\forall w'' \in Dox_x(w_o).\mathscr{E}c.\vartheta(c)(w'')(w_o)](\lambda c.\lambda w.\lambda w'.c \in \wp : [c(w)(w') \dotdiv t_1(w)(w')]) =$
$\lambda x.\forall w'' \in Dox_x(w_o).\mathscr{E}c.[\lambda c.\lambda w.\lambda w'.c \in \wp : [c(w)(w') \dotdiv t_1(w)(w')]](c)(w'')(w_o) =$
$\lambda x.\forall w'' \in Dox_x(w_o).\mathscr{E}c : c \in \wp[c(w'')(w_o) \dotdiv t_1(w'')(w_o)]$

c. $[[\text{know}_{CQ}\ \mathscr{Q}(\text{t}_1)]]([[Fred]]) =$
$[\lambda x.\forall w'' \in Dox_x(w_o).\mathscr{E}c : c \in \wp[c(w)(w') \dotdiv t_1(w)(w')]](F) =$
$\forall w'' \in Dox_F(w_o).\mathscr{E}c : c \in \wp[c(w'')(w_o) \dotdiv t_1(w')(w_o)]$

d. $[[\lambda_1 \cdot \text{ Fred know}_{CQ}\ \mathcal{Q}(\text{t}_1)]] =$
$\lambda z.\forall w'' \in Dox_F(w_o).\mathcal{E}c : [c \in \wp : c(w'')(w_o) \doteqdot z(w'')(w_o)]$

e. $[[\lambda_1 \cdot \text{ Fred know}_{CQ}\ \mathcal{Q}(\text{t})]]([[\text{the price}]]) =$
$[\lambda z.\forall w'' \in Dox_F(w_o).\mathcal{E}c : c \in \wp[c(w'')(w_o) \doteqdot z(w'')(w_o)]](\iota x.\mathcal{E}d \uparrow price(d)(x)) =$
$\boxed{\forall w'' \in Dox_F(w_o).\mathcal{E}c : c \in \wp[c(w'')(w_o) \doteqdot \iota x.\mathcal{E}d \uparrow price(d)(x(w'')(w_o))]\text{w}}$

(131) Reading A: *John knows the price that Fred knows*

a. $\Lambda(\uparrow [[\text{John knows}_{CQ}\ \mathcal{Q}(\text{the price})]])(\uparrow [[\text{Fred knows}_{CQ}\ \mathcal{Q}(\text{the price})]]) =$
$[\{\uparrow x/c\} \uparrow \forall w \in Dox_J(w_o).\mathcal{E}c : c \in \wp[c(w)(w_o) \doteqdot \iota x.\mathcal{E}d_1 \uparrow price(\uparrow d_1)(x(w)(w_o))]\ ;$
$\{\uparrow x/c'\} \uparrow \forall w'' \in Dox_F(w_o).\mathcal{E}c' : c' \in \wp[c'(w'')(w_o) \doteqdot \iota x.\mathcal{E}d_2 \uparrow price(\uparrow d_2)(x(w'')(w_o))]]$

b. By the schema in (129), we get:
$[\uparrow \forall w \in Dox_J(w_o).\mathcal{E}c : c \in \wp[c(w)(w_o) \doteqdot \iota x.\mathcal{E}d_1 \uparrow price(\uparrow d_1)(x(w)(w_o))]\ \wedge (d_1 = x)\ ;$
$\uparrow \forall w'' \in Dox_F(w_o).\mathcal{E}c' : c' \in \wp[c'(w'')(w_o) \doteqdot \iota x.\mathcal{E}d_2 \uparrow price(\uparrow d_2)(x(w'')(w_o))]\ \wedge (d_2 = x)]$

c. In turn, (131b) reduces by associativity (Fact 3.4.4:1) to:
$[\uparrow \forall w \in Dox_J(w_o).\mathcal{E}c : c \in \wp[c(w)(w_o) \doteqdot \iota x.\mathcal{E}d_1 \uparrow (price(\uparrow d_1)(x(w)(w_o))\ \wedge (d_1 = x))]\ ;$
$\uparrow \forall w'' \in Dox_F(w_o).\mathcal{E}c' : c' \in \wp[c'(w'')(w_o) \doteqdot \iota x.\mathcal{E}d_2 \uparrow (price(\uparrow d_2)(x(w'')(w_o))\ \wedge (d_2 = x))]]$

d. Application of the closure operation $\downarrow$ approximately yields:
$[\forall w \in Dox_J(w_o).\exists c : c \in \wp.c(w)(w_o) \doteqdot \iota x.price(y)(x(w)(w_o)) \wedge$
$\forall w'' \in Dox_F(w_o).\exists c' : c' \in \wp.c'(w'')(w_o) \doteqdot \iota x.price(y)(x(w'')(w_o))]$

e. The interpretation of the above is simply that John and Fred have a price concept in all of their accessible doxastic alternatives w and w'', respectively, for the same product y.

Note that in (131d) John and Fred may still have different price-concepts, c and c', for the same product y. I believe that this is one possible interpretation of Reading A. The other possible interpretation is one in which they share the same price concept. This interpretation can be acheived through essential the same means, except that existential disclosure must operate over the concepts, rather than the implicit arguments of the CQ NP. In favor of our theory is again the fact that in (132), John and Fred *must* know the same concept for the price of milk in *both* the fully articulated relative clause (132a) and the ACD context (132b).

(132) a. John knows the price of milk that Fred knows

b. John knows the price of milk that Fred does

We now provide another complete derivation for Reading B of Heim's ambiguity.

3.5.2.2 Reading B as a propositional attitude

In this framework, Reading B can be derived in essentially the same way, except that there is only one noun heading the relative clause. Composing the relative clause with the matrix clause via Late Merger is not an option in this case, since the matrix clause would lack a complement. Thus, I propose that the relative clause forms a proposition and serves as a complement to the propositional *know*. The vocabulary for the derivation is nearly the same, except that we require the explicit addition of $know_{prop}$. In addition, the derivation of the relative clause follows the pattern shown in (130), which is repeated for convenience below:

(133) Vocabulary and translations:

i. $[[know_{prop}]] = \lambda p_{\langle s,t\rangle}.\lambda x.\forall w \in Dox_x(w_0).p(w) = 1$

ii. $[[\text{the price that Fred knows}]] = \forall w'' \in Dox_F(w_o).\mathscr{E}c : c \in \wp[c(w'')(w_o) \doteqdot \iota x.\mathscr{E}d \uparrow price(\uparrow d)(x(w'')(w_o))]$

(134) From internally-headed relative clause to proposition:

a. $\downarrow(\ [[\text{the price that Fred knows}]]) =$
$\forall w'' \in Dox_F(w_o).\exists c \in \wp[c(w'')(w_o) \doteqdot \iota x.\exists y price(y)(x(w'')(w_o))]$

b. The lift ^ of (134) yields:
$\lambda w.[\forall w'' \in Dox_F(w_o).\exists c \in \wp[c(w'')(w_o) \doteqdot \iota x.\exists y price(y)(x(w'')(w_o))]](w)$

(135) The proposition above composes with $know_{prop}$:

a. $[[know_{prop}]](\ ^\wedge\downarrow\ [[\text{the price that Fred knows}]]) =$
$[\lambda p_{\langle s,t\rangle}.\lambda x.\forall w \in Dox_x(w_0).p(w) = 1]$
$(\lambda w.[\forall w'' \in Dox_F(w_o).\exists c \in \wp[c(w'')(w_o) \doteqdot \iota x.\exists y.price(y)(x(w'')(w_o))]](w)) =$
$\lambda x.\forall w \in Dox_x(w_0).[\forall w'' \in Dox_F(w_o).$
$\exists c \in \wp[c(w'')(w_o) \doteqdot \iota x.\exists y.price(y)(x(w'')(w_o))](w)] = 1$

b. $(\ [[know_{prop}]](\ ^\wedge\downarrow\ [[\text{the price that Fred knows}]]))(\ [[John]]) =$
$\forall w \in Dox_J(w_0).[\forall w'' \in Dox_F(w_o).\exists c \in \wp[c(w'')(w_o) \doteqdot \iota x.\exists y.price(y)(x(w'')(w_o))](w)] = 1$

c. The interpretation of the above is just that for all of John's doxastic alternatives w accessible from w_0, the proposition that *in all of Fred's doxastic alternatives* w''*, Fred knows the price of something* holds in w.

With the derivation of Reading B, we thus conclude our exposition of the present account of concealed questions, and turn our attention to the problems and prospects associated with the account.

3.6 Problems and Prospects

Naturally, several loose ends remain to be discussed. As we cannot address all of them here, we instead focus on just a few pertinent ones. We re-examine how CQs relate to propositions and free relatives, offer a general pragmatic solution for why CQs (a) must be interpreted as sufficiently specific, and (b) do not licenses names such as Rome as the CQ NP. We conclude this chapter with a puzzle that no present approach, including the one endorsed above, can transparently account for. Two types of evidence – grammatical and semantic – suggest that the interpretation of abstract and concrete relational nouns differ significantly in certain contexts.

3.6.1 CQs and other Grammatical Objects

3.6.1.1 CQs and Propositions

Recall Nathan's (2006) observation that verbs that select propositions in English typically select for CQs as well. The correlation (PCQC) is indeed suggestive and provides a powerful argument for a relation between CQs and propositions. Given the semantic approach introduced in this chapter, it might be argued that the existential binding of the concept *c* could just as well occur as part of the CQ meaning itself, rather than occurring as part of the verb denotation. Evidence against this view comes from unifying the paraphrase with free relatives, as in the next section.

Nathan (2006) acknowledges that the correlation makes extensive predictions that do not appear to fit the predictions of his account, and spends significant time addressing how such predictions might be explained. One particularly problematic verb is *depends* which selects for a wide variety of arguments, but not propositions. The examples below indicate that *depends* takes CQs and free relatives, but not propositions (examples (136a–b) are from Heim (1979)):

(136) a. The temperature of the lake depends on the season
b. What the temperature of the lake is depends on what season it is
c. * That temperature of the lake is 58°F depends on the season

The correlation is also problematic for cross-linguistic reasons. It fails to explain why languages with propositional complements could fail to license CQs. Numerous languages do *not* admit concealed question intepretations (Caponigro and Heller, 2007), or severly limit the appropriate contexts (*e.g.,* Hungarian; A. Szabolsci pers. comm.).

Yet, it is clear that the correlation that Nathan proposes offers a genuine insight into a close relationship between CQs and propositions. Perhaps the relationship is similar to Rescher's distinction between two types of knowledge: *performatory know-how* and *procedural know-how*. In our case, performatory know-how might describe the cases in which we have the ability to properly individuate an individual when answering the question of who or what that individual is. Procedural know-how, in Rescher's terms, might apply to the case of knowing the fact that the answer to the question is properly individuated by the answer. On the relation between the performatory and procedural know-how, Rescher (2005, 7) writes "... the different modes of knowledge are inextricably interconnected. To know (propositionally) *that* a cat is on the mat one must know (adverbially) *what* a cat is. And this knowledge rests on knowing how to tell cats from kangaroos."

So why might many of the same verbs select for both propositions and CQs, if the latter do not reduce to the former? I speculate that the answer to this question resides in how individual languages organize their lexicons, so to speak. If a language is willing to allow homophonic variants that select for propositions, then they might be more likely to likewise allow additional variants that select for CQs. So I suggest that propositions and CQs are sometimes paired together because once a language allows vagueness into the lexicon by permitting homophonic variance, the cost of allowing additional homophones, such as those selecting for CQs, is reduced. The fact that CQs are not universal supports the idea that the operations $\mathcal{Q}$ forming CQs is perhaps lexical and requires that special homophonic variants to successfully combine. Naturally, the plausibility of this sort of speculation is subject to how widely it is supported by empirical fact.

3.6.1.2 CQs and Pseudo-clefts

In addition to the relationship between CQs and propositions, we ought to re-consider how CQs and their indirect question paraphrases correspond. Recall that the indirect paraphrase is syntactically a headless relative, such as those observed in pseudoclefts (137).[13]

(137) a. The temperature of the lake is 58 degrees F
b. What the temperature of the lake is is 58 degrees F

The relationship between CQs and specificational sentences has been explored in great detail by Romero (2004, 2005). Similarly to Romero, we can unify CQs and pseudoclefts in the present account. Without presenting a full formalization, suppose that relative phrases of the kind above are of the same type of CQs, that is, of type $\langle\langle s,e\rangle,\langle s,\langle s,t\rangle\rangle\rangle$. Then, the relative phrase *what the temperature of the lake is* combines directly with the post-copular noun phrase *58 degrees F*, to yield a propositional concept: $\lambda w.\lambda w' : 58(w)(w') \doteqdot \iota x.\text{temp}(\text{lake})(x(w)(w'))$, to be evaluated at the relevant worlds, depending on context.

Note that we cannot claim that *all* pseudoclefts are of the above form. In particular, only specificational pseudoclefts, those that exhaustively specify the argument of the relative phrase, involve the CQ semantics. Other types, such as predicational pseudoclefts (138), must be analyzed in other ways.

(138) a. What John wants to eat for dinner is lasagna
b. What Mary is is very intelligent

3.6.1.3 CQs and Names

Recall our discussion of Romero's criticism of Heim's context-dependence approach in §2.2.3. Romero observed that proper names, such as *Rome*, do not easily admit a CQ interpretation:

[13] See den Dikken (2001) and Higgins (1973) for a review of pseudocleft types.

(139) Names do not favor CQ paraphrases

a. John knows Rome

b. ≠ John knows what Rome is

There are at least two ways of viewing this contrast. First, it might be that names are independently incompatible in free relative phrases, at least when used in specificational contexts. This generalization appears to be in part correct, as the free relative *what Istanbul was* is licensed in predicational contexts (140a), but not specificational contexts (140b), where identity is at issue:

(140) a. What Istanbul was will never be achieved again

b. ?? What Istanbul is is Constantinople

A second way to interpret Romero's observation offers a more pragmatic explanation. Assuming that a CQ interpretation is created by a type-shifting operation $\mathscr{Q}$, it is possible that the concept *c* which answers the indirect identity question associated with *what Istanbul is* must be more specific than the CQ NP itself. Intuitively, this requirements makes good sense, as we typically require legitimate answers to be informative in the relevant contexts. Thus, names such as *Rome*, *Istanbul*, and *John*, do not license CQ interpretations precisely because these names typically provide the most direct method of identification.

3.6.1.4 CQs and Determiners

For reasons of space, we have only considered CQs with definite descriptions. Although Heim (1979) and others have argued that other determiners/ quantifiers make for acceptable CQ paraphrases, I think that there are severe limitations on which quantifers in fact permit a CQ interpretation. For instance, putting the indefinite *a mayor of Amherst* into a CQ paraphrase (*who/what a mayor of Amherst is*) cannot be interepreted as a quesiton about the *identity* of the individual denoted by the DP. Instead, the paraphrase seems to involve attibution of a property to a kind *mayors of Amherst*.

(141) a. John knows a mayor of Amherst

b. John knows who/what a mayor of Amherst is

Informally, the determiners that typically allow a 'genuine' CQ interpretation are just those that support a distributive reading of the identity conditions picking out the individuals in question. For instance, the DP *most senators in the 110th Congress* can surely be paraphrased with a question reading. The most immediate interpretation is the one in which John can identify *each* member of Congress in the group (142b). We do not get the reading where John knows *the property* that picks out most members of Congress (142c).

(142) a. John knows most senators in the 110th Congress

b. John knows who most senators in the 110th Congress are
e.g., John knows that Maria Cantwell is a senator for Washington State, that John Kerry is a senator for Massachusetts, etc.

c. ≠ John knows what most senators in the 110th Congress are
e.g., John knows that most senators in Congress are self-proclaimed fiscal conservatives

Although the predictions that the present theory of CQs makes for how to interpret a wider range of quantifiers in CQ contexts cannot be elaborated here, it is expected that it will be able to capture the fact that such quantifiers are interpreted as distributing the identity conditions across members in the set.

3.6.2 Abstract *versus* Concrete Relational Nouns

In the next two sections, I provide evidence for differentiating abstract from concrete relational nouns. First, relational nouns in contexts of 'property specification' are shown to allow abstract but not concrete relational nouns. Second, the two types of relational noun are shown to differ with respect to how they are modified by *actual* and *actually*, making for a difference in potential CQ paraphrases. These issues are discussed in turn, but no attempt to incorporate these facts into the present approach is made.

3.6.2.1 CQs in Contexts of Specification

The construction in (143) contains the postverbal light quantifier *something*, which functions as a placeholder for a property or proposition (see Moltmann (2003) for discussion of *something* in complement position). For lack of a better term, I call constructions which specify the content of the property or proposition immediately following the placeholder a *property specification* environment. Curiously, a CQ from a Concrete Relational Noun is not admissible in such environments (143a), even though an overt identity question (143b) is allowed. This pattern contrasts with CQs from Abstract Relational Nouns; both the CQ and the overt question are acceptable (143c–d).

(143) John knows something

a. * the mayor of the town
b. who the major of the town is
c. the price of milk
d. what the price of milk is

As a whole, the pattern observed in data in (144) contrasts with the pattern observed in (143). Here, only the individual interpretation of the DP is supported, and as such gender agreement prohibits the neuter noun phrase in (144c).[14]

(144) John knows someone:

a. the mayor of the town
b. * who the major of the town is
c. * the price of milk
d. * what the price of milk is

At first sight, the opposing patterns might be thought to be easily explained away by appeal to an ambiguous reading of *something*. It could be said that one reading supports the individual interpretation marked for neuter agreement, thereby explaining the contrast between (143a) and (143c). The other reading of *something* stands for proposition-like entities, and thus takes the full question of either Abstract or Concrete RNs. The account might continue that it is a curious fact of English that the animate *someone* is not likewise ambiguous, lacking as it does the second, proposition-like interpretation.

Yet, this fact is still truly confounding; for we expect that DPs which function as CQs in other contexts should be interpreted in contexts that select for properties or propositions. Clearly, this is not possible with DPs denoting persons, or more accurately, with DPs denoting positions which select persons. Thus, any alleged ambiguity of *something* would not be sufficient to explain the contrast (143 – 144).

Similarly, the specificational pseudocleft exhibits a pattern identical to (143)

(145) What John knows is

a. * the mayor of the town
b. who the mayor of the town is
c. the price of milk
d. what the price of milk is

If an ambiguity for *something* is posited to explain the contrast in (143), the very same ambiguity must be operative in pseudoclefts (145). In particular, such a view would stipulate that the entire relative clause shares an ambiguity which was thought to be lexical in nature. Yet, these relative clauses are not lexical items, nor lexically fixed; any ambiguity proposed for pseudoclefts should be cashed out systematically, not on the basis of an exception.

[14]This contrast developed out of a helpful conversation with Angelika Kratzer.

3.6.2.2 The Interpretation of *actual*

Although the paraphrase conditions of CQs are often treated as uniform across various types of noun phrases (cf. Jackendoff (1997)), certain abstract NPs diverge from concrete relational NPs with respect to the interpretation of the intensional adjective *actual*.[15] In brief, the adnominal modifier *actual* enforces different truth conditions on its complement than its adverbial counterpart *actually*, when they modify relational nouns. Furthermore, the paraphrasability conditions between *abstract* and *concrete* are shown to differ. The interpretation of the question paraphrase with *actual* and *actually* reduce to the same interpretation if the noun is an abstract relational noun, but differ significantly if the noun is concrete. Thus, there is a choice of paraphrase between the adjective and the adverb in one type of relational noun, but not the other.

For illustration, suppose that John is an economist on the House Budget Committee and that the House Republicans have pushed a new bill that calls for a reduction of government spending for retirement. Being an expert on taxation and retirement, John recalculates the proposal under a more realistic economic model and finds that, lo and behold, his calculations diverge from those quoted by the drafters of bill. Example (146) can be paraphrased *two ways*, either as (146a) in which the adjective *actual* modifies *cost*, or as (146b) in which case the adjective *actual* reinterpreted as an adverb, modifying the *way in which the cost instantiates*.

(146) John predicted the actual cost of the reform

a. John predicted what the *actual* cost of the reform was
b. John predicted what the cost of the reform *actually* was

Do the truth conditions of (146a–b) differ? The difference, should it truly exist, is not entirely intuitive. Let S_1 be a situation in which w' is thought to be the actual world by the addressee(s) of (146). In w' the cost of the reform equals $26 million dollars. John's recalculations show that the cost of the reform in w'' is instead $28 million dollars.

In this situation, example (146) has the force of asserting that w'', and not w', is the actual world with respect to the cost of the reform. It seems that either (146a–b) captures this reading.

Let S_2 be the situation in which John has uncovered a case of attempted graft in which the monies from the proposed reform was, in reality, earmarked for kickbacks to government

[15] *Actual* has many uses, including one for signalling negation or revision in discourse. It is important not to interpret the examples with such a use. Other intensional-like adjectives such as *real*, suffer from the same difficulties in interpretation. Another candidate is the non-intersective *alleged*. However, it might invoke a hidden speech act, say from an implicit accuser, which may degrade the intended adverbial reading. At any rate, the productivity of such examples is not at all certain. See also Zimmerman (2000) for an account of *occasional* that displays a similar type of ambiguity, although on its adverbial reading *occasionally* no longer modifies the DP, but instead the entire sentence. Further investigation is required to determine the relationship between these cases.

World		Reform cost
w'	$\mapsto$	\$26m
w''	$\mapsto$	\$28m

Table 3.5: Cost Concept in situation S_1

contractors. John may not even be able to calculate the exact number, and yet might *still* predict that the cost of the reform is always equal to the funds diverted towards another fund (or even a part thereof). John need not know or predict any particular number, only a relation that sufficiently establishes (a relative) identity. In this situation, the cost in dollars of the reform might not differ between possible worlds, say w' and w'', but the way in which John identifies it.

Worlds		Reform cost	Worlds		Kickback amount
w'	$\mapsto$	\$26m	w'	$\mapsto$	\$10m
w''	$\mapsto$	\$26m	w''	$\mapsto$	\$26m

Table 3.6: Cost concept in situation S_2

The point is that in S_2 example (146) seems to require not that the costs in w' and w'' differ and that w'' is the actual world, but rather that the cost of the reform shares a property or is identical with another concept in the actual world. In this situation, w' would be ruled out on these grounds.

It is important to determine whether both paraphrases (146a–b) are appropriate in situation S_2. While it is not completely clear to me, I am strongly inclined to suggest that they are. That is, I maintain that there is no crucial difference between the question paraphrases in the CQ context.[16] So far, this conclusion is unsurprising if we accept the common view that *actual* and *actually* are mere syntactic variants of one another. The contrast that appears with concrete relational nouns clearly makes this view untenable.

Suppose that John has been investigating a vast conspiracy instituting a shadow government in America. He uncovers that the presidency is, in reality, occupied by another politician. In this scenario, (147) can only be paraphrased as (147a).

(147) John discovered the actual president

a. John discovered who/what the actual president was

b. # John discovered who/what the president actually was

[16]The slight difference could be cashed out as follows: it seems that (146a) might be appropriately used when there was a slight miscalculation and the *actual cost* is another value. In contrast, (146b) has a stronger reading, namely that the quality constituting the cost is different in the actual world then expected. In such an instance, the numerical value could be the same or not, but the character of the way in which it constitutes a cost must differ.

To see the difference, suppose that S_3 is the situation in which the addressee(s) of (147) are thought to believe that they are in w' and the president (of the United Sates) in the actual world denotes some individual: George W. Bush. However, John discovers the shadow government's plot, in which the president in w' is not president at all, but rather a political puppet. Some other individual – say Dick Cheney – is in fact president, as John discovers.

World		President
w'	$\mapsto$	GWB
w''	$\mapsto$	DC

Table 3.7: President concept in situation S_3

It is a striking contrast to (146) that the paraphrase in (147b), containing the adverb *actually*, is not an appropriate paraphrase. By itself, (147b) means something quite different than either (147) or (147a). Let S_4 be a situation in which John discovers that in reality the president occupies another position, such as a renegade semanticist, say Aaron Aaronson, who writes mysterious yet brilliant papers anonymously. In this very counterfactual world, it is not the referent of *the president* that changes. Instead, the extension of the president concept is equivalent to the extension of the renegade semanticist concept in whatever world is actual.

Worlds		President	Worlds		Renegade Semanticist
w'	$\mapsto$	GWB	w'	$\mapsto$	AA
w''	$\mapsto$	GWB	w''	$\mapsto$	GWB

Table 3.8: Cost concept in situation S_4

In this case, we could utter (147b), but not (147) or (147a). The point is that while (147b) is felicitous on other occasions of use, it is not a readily available paraphrase of (147). The difference in interpretation here is surprisingly robust, especially in contrast to the pattern illustrated by (146).

The contrast between (146) and (147) suggests that CQs containing *actual* and an abstract relational noun may be paraphrased in either (a) the adverbial reading or (b) the adjectival reading of *actual*, where as those CQs containing concrete relational nouns do not license the adverbial paraphrase under similar circumstances.

Note that this contrast is not limited to CQ environments. The basic pattern is replicated in simple matrix clauses, such as those below.

(148) Abstract Relational Noun

a. The actual cost of the reform is \$26m

b. The cost of the reform is actually \$26m

(149) Concrete Relational Noun
  a. The actual president is Dick Cheney
  b. The president is actually Dick Cheney

Although the different senses of (148a–b) is subtle, the distinct interpretations of (149a–b) is non-trivial. While (149a) claims that the person who truly holds the presidency is not Bush, but Cheney, (149b) claims something much stronger. It means that the referent of the term *the president*, George Bush, is in fact Dick Cheney. So, it appears that the relation of identity which is asserted by (149a) is about who holds office in the actual world – a world implicitly other than how we had imagined. Example (149b), on the other hand, is about the identity of the person who holds office in the worlds we believe to be actual. This is perhaps why (149b) sounds like a strange assertion; it equates two objects we believe to be distinct.

The pattern is again replicated with descriptions in the post-copular position:

(150) Abstract Relational Noun
  a. The actual cost of the reform is the amount of the kickback
  b. The cost of the reform is actually the amount of the kickback

(151) Concrete Relational Noun
  a. The actual president is the man in the brown hat
  b. The president is actually the man in the brown hat

The full ramifications of this contrast are far from clear. It is possible that the difference between types of relational noun could be fomulated within a two-dimensional semantic theory (see Chalmers (2007) for an excellent overview article, as well as Soames (2005)), in which the adverbial *actually* modifies the 1-intension, or the standard Carnap/Fregean sense of an expression, whereas the adjective *actual* modifies the 2-intension, which determines how the reference of the term is fixed. All the same, the exact solution of this puzzle is left as a challenge for future, and more comprehensive, theories of concealed question. Before ending this chapter, we briefly review the material discussed above.

3.7 Summary

In this chapter, a partially pragmatic approach to concealed questions was introduced and defended. In particular, CQs were argued to be formed via a type shifting operation $\mathcal{Q}$ lifting nouns of type $\langle s, e\rangle$ to type $\langle\langle s, e\rangle, \langle s, \langle s, t\rangle\rangle\rangle$. The result was interpreted as a relation between a concept from a conceptual cover (Aloni, 2001) and an identity question. Two inter-definable conceptions of the proposed denotation were explored and discussed in detail. The formal semantics of the approach was developed within the theory of Dynamic Montague Grammar Groenendijk and Stokhof (1990a) with Existential Disclosure (Dekker, 1993a). Complete derivations of basic CQ NPs provided the basis for an account of Heim's ambiguity that exploited the syntactic evidence for analyzing the ambiguous readings in terms of distinct sorts of relative clauses.

As the account sketched above depends on the re-interpretation of a nominal as a higher-order relation, the next chapter presents extensive and novel psycholinguistic evidence which is consistent with such an analysis.

4

The Cost of Concealed Questions

In chapter 2, a review of the current literature on the semantics of concealed questions was presented, along with a few novel facts, both interpretive and syntactic, that challenged the empirical coverage of several competing accounts. In chapter 3, a proposal specifying how CQs are interpreted was offered. Special care was made to incorporate these novel facts into this analysis. The proposal used the theory of conceptual covers (Aloni, 2001) to account for the counter-examples to Greenberg's generalization (Greenberg, 1977), and a dynamic extension of Montague Grammar (Groenendijk and Stokhof, 1990a) to provide an account for Heim's ambiguity (Heim, 1979) that observed the syntactic evidence gathered by examining CQs with relative clauses in ACD and sluicing contexts.

The semantic core of this account centered around positing a type-shifting operation, $\mathcal{Q}$, which maps an individual concept to a relation between identifying concepts from a conceptual cover to an identity question, thus indirectly accounting for the covert question paraphrase, characteristic of concealed questions.

This chapter presents psycholinguistic evidence that is broadly compatible with a view of concealed questions that involves a shift of interpretation.[1] The evidence here does not claim to support any particular formal theory of interpretation associated with concealed questions, and as such is consistent with many, if not most, current accounts of CQs, in particular Nathan (2006) and Romero (2004, 2005, 2007a, 2007b). Further experimentation and inquiry would have to be undertaken to determine whether these theories make different

[1]This chapter is based in large part on Harris et al. (2007), to which the methods and result sections are nearly identical. Given the different audience, introductions to the experiment and the methodologies have been altered. The conclusions have also been revised and somewhat simplified. I thank my co-authors Liina Pylkkänen, Brian McElree, and Steven Frisson for their work on the project.

predictions for the measurements employed here.

In brief, this chapter presents the findings from an eye-movements and magnetoencephalography (MEG) study, which found, on both measures, that concealed questions were more costly to interpret during online sentence comprehension than were matched controls. Numerous controls of the materials were performed to ensure that CQs were not independently less plausible or less preferred constructions. Further, a separate lexical decision task in which question-selecting and control verbs were presented in MEG provides additional evidence that the critical verbs were not independently more difficult than controls, and that the effect observed when processing sentences with concealed questions was indeed a reflection of sentence, and not lexical, level processing. Before further discussing the results, however, we motivate the logic behind the design of the experiment.

4.1 The Cost of Non-Compositionality

Many linguists and philosophers have hypothesized natural language to be, by and large, compositional in nature (see discussion in Dowty (2007), Fodor and Lepore (2002), Montague (1970)). To say that a language is compositional is just to say that the meaning of a complex term results from a function combining the meanings of its constituent parts. The origin of compositionality as a theoretical principle has sometimes been attributed to Frege, although Janssen (1997) argues that there are significant differences between Frege's writings and the current thesis, or rather classes of theses, regarding the compositionality of language.

Compositionality itself is not a monolithic concept, as there are numerous ways in which a language could obey the principle, general as it is. In particular, a major dividing line has been drawn between *strong compositionality*, in which the meaning of a complex phrase is *exclusively* a function of its overt syntactic parts, and *weak compositionality*, in which overt constituents contribute *inclusively* to the meanings of the phrases they comprise (see Groenendijk and Stokhof (1989), Janssen (1997) and Partee (1995) in particular). The differences between the two visions of compositionality are far from trivial. Most current semantic theories include extra-syntactic, *i.e.,* syntactically covert, operations that contribute to meaning composition. Under this view, language is *weakly compositional*; the fundamental principles of compositionality are maintained even though the meaning of constitutive expressions cannot be transparently recovered from their overt subexpressions. Thus, we speak of these covert operations as *enriching* the semantic content of an expression. Although there is much more to be said about compositionality in general, I will instead refer the reader to the sources cited above, as a complete discussion would prevent us from reaching the experimental results to be presented in the chapter.

How does the way in which natural language is compositional relate to psycholinguistic research? The hypothesis that constructions requiring semantic enrichment elicit observable processing costs has been the subject of many recent studies in the psycholinguistic literature (see Pylkkänen and McElree (2006) for review). For clarity, the broadest working assumption is condensed into the following hypothesis:

Hypothesis 4.1.1. *[Enriched meanings tax the language processor] Constructions which require additional covert operations in the semantics are taxing on the comprehender.*

Of course, such a hypothesis does not specify *why* enriched composition might tax the processor. Therefore, another hypothesis regarding the relationship between natural language structure and comprehension is posited:

Hypothesis 4.1.2. *[The linking hypothesis] Enriched meanings require additional cognitive resources for interpretation than are signaled in the syntax.*

It is clear that the second of these hypotheses is far stronger than the first and that the correctness of the former in no way determines or hinges upon the latter. Nonetheless, it is a useful hypothesis as it explicitly addresses *why* enrichment is taxing and makes a serious, and falsifiable, claim about *what* psycholinguistic studies intend to measure. This is not to commit to the view that *only*, or even that *all*, enriched meanings are taxing, for there is ample evidence that lexical level factors, such as frequency, length, probability, plausibility and the like all contribute to natural language processing. Additionally, far too little is known about the brain or the measurements used to explore it to make sweeping assertions about what a particular study or paradigm tells us about language processing, at least if we are to remain conservative and modest in our generalizations.

Even within the field, many researchers have expressed concerns over the connection between the measurements of an experiment and the conclusion about cognitive processes that the results are thought to support. Such doubts have been articulated with regard to comparing results across methodologies, *e.g.,* hemodynamic versus electromagnetic measures (see Horwitz and Poeppel (2002)), and even to the entire localization enterprize in general (see Uttal (2001) for extensive discussion). Heeding these concerns as much as it is fruitful in the present context, we now discuss a construction which bears some superficial relation to concealed questions in order to establish a basis of comparison for similar experiments having precedent in the literature.

4.1.1 Complement Coercion

In previous psycholinguistic research, the process of semantic enrichment known as *coercion* has gained much attention. Coercion is a general term within linguistics that describes the situation in which linguistic elements, thought to be incompatible on their default meanings, nonetheless combine successfully. Two specific forms of coercion have attracted the attention of psycholinguistic researchers thus far. The first of these is *aspectual coercion* (Pustejovsky and Bouillon (1995), Jackendoff (1997)), whereby a punctual action (e.g., *hop*) takes on an iterative meaning in durative contexts (*The robin hopped all afternoon*). Although aspectual coercion has been found to be difficult for aphasic subjects to interpret (Piñango et al.

(1999) and Piñango and Zurif (2001)), corresponding results have not been systematically replicated with normal subjects in online sentence comprehension tasks (see Todorova et al. (2000), and Pickering et al. (2006), as well as Pylkkänen and McElree (2006) for discussion). The second type of coercion studied by psycholinguists is *complement coercion* (Jackendoff (1997), Pustejovsky (1995)) in which an object-denoting noun (e.g., *book*) is interpreted as an event when combining with an event-selecting verb (e.g. *The author began the book*). For instance, example (152a) pairs an object-selecting verb (*read*) with a typical object-denoting complement (*book*). When the same complement is interpreted with an event selecting verb (*began*), as in (152b), the interpretation involves additional composition, often formalized as shifting the meaning of the object-denoting noun into an event (Jackendoff, 1997). Further, the type of event is determined in part by context, as shown in (152c):

(152) a. The author read the book
b. The author began the book
c. The author writing/reading/editing the book

4.1.1.1 The cost of complement coercion

Complement coerion was found to be costly across several different types of method, including self-paced reading (McElree et al. (2001), Traxler and McElree (2002)), speed-accuracy tradeoff studies (McElree et al. (2006)), eye-tracking paradigms (Traxler and McElree (2002)) and MEG (Pylkkänen and McElree (2007)). Focusing on the second two for comparison to our own study, eye-tracking studies found that phrases with coerced complements (152b) engendered longer eye gaze duration fixations on the noun (McElree et al., 2001) or spill over regions directly following the noun Traxler and McElree (2002), than did controls (152a), despite being rated as equally plausible.

Further studies suggest that the additional time spent on coercion constructions reflects a genuine interpretive cost, rather than underspecification or plausibility differences. Pickering et al. (2005) compared coercion constructions, as in (152), and counterparts that included the gerund expressing the eventive interpretation (152c). Interestingly, the results showed longer reading times on the complement noun on both of the coerced constructions (152b–c), than on controls. Importantly, event-denoting complements when composing with event-selecting verbs do not show any processing delays from object-denoting complements combining with typical object-denoting complements. This pattern suggests that the interpretative cost recorded from complement coercion is related to accommodating the mismatched selection, presumably by an interpretative shift applied to the noun (see Pylkkänen and McElree (2006) for overview of existing analyses).

In MEG, Pylkkänen and McElree (2007) tested similar materials and found that coerced noun complements elicited larger amplitudes in the Anterior Midline Field (AMF), a magnetic field pattern generated in the ventromedial prefrontal cortex at 350–400 ms. This area has sometimes been implicated in EEG studies manipulating sentence level processing, although the general role of this area in language processing is largely unexplored (for a review of recent EEG and ERP work, see Kutas and Schmitt (2003)).

Given the experimental precedent set by studies on complement coercion, we chose to measure the expected costs associated with processing concealed questions to those associated with coercion. There are several reasons to expect that the two constructions might elicit similar effects. The constructions are both hypothesized to involve shifts of interpretation on nouns in contexts where the verb does not select for an individual. Thus, it is possible that the parser possesses just one general device for repairing selection mismatch.

Hypothesis 4.1.3. *[General nominal shifting device] The parser employs a general semantic repair mechanism that applies to both complement coercion and concealed questions.*

Under such a hypothesis, we expect that complement coercion and concealed questions elicit qualitatively similar responses in online comprehension. Given the results discussed above, we would then expect that concealed questions provoke longer gaze durations in an eye-movement study and elicit greater amplitudes in the AMF in later time windows of MEG analysis.

In contrast, there are reasons to expect that the language processing system is endowed with multiple repair mechanisms that elicit qualitatively distinct costs. If complement coercion and concealed questions do not share common mechanisms, then it is plausible that the different constructions will prompt qualitatively different processing costs.

Hypothesis 4.1.4. *[Distinct semantic repair mechanisms] The parser employs a different semantic repair mechanism when interpreting complement and concealed questions.*

Provided that the interpretation of concealed questions is indeed shown to tax the parser, it is entirely appropriate to ask how the specific cost recorded compares to other, similar sorts of constructions. The three experiments discussed in depth below were designed to both determine whether concealed questions are reliably more difficult to processes than comparable controls and to what extent the sort of cost found supports one or other of the above hypotheses. Two major experimental methodologies were employed for this purpose: eye-tracking and MEG. We present the findings of the experiment in that order.

4.2 Eye-tracking

The eye-tracking paradigm is an established way of monitoring eye-movements of a subject as he or she attends to a stimulus. Several types of eye-movement have been identified, some of which are thought to involve independent oculomotor control, and others cognitive processes related to incorporating the information relayed by the stimulus. The coarsest categorization is between *fixations* and *saccades*. Fixations are periods when the eye attends to a stimulus without making any major movements and tend to last 200 – 300 ms, depending on the kind of stimulus presented. Saccades are quick ballistic movements, with velocities of up to 500° per second, made approximately 3–4 times in a single second (Matin (1974), Richardson and Spivey (2004a)).

However, to say that there are genuine fixations is somewhat misleading, as the eye is continuously in motion, whether by the constant tremor of the eye, known as *nystagmus*, or larger movements, such *drifts* and *microsaccades*, which are thought to reflect the imperfections of the human oculomotor control system. Nonetheless, we use the term fixation to refer to periods of time when the eye attenuates to a point of regard for a fixed period of time. Fixations are generally used to measure attention which in turn are thought to sometimes reflect cognitive processing.

A relatively small portion of the eye is responsible for the majority of information uptake within the overall visual field. The most detailed information comes from the *foveal region*, which constitutes only about 2° of the 200° that the eye monitors (see Richardson and Spivey (2004a) for review). The area extending 5° in either direction of the fovea is known as the *parafoveal region* in which visual acuity drops significantly. Information in the parafoveal area is sometimes active for processing, and is thought to both help determine patterns of saccade in the region and aid lexical access. The rest of the visual field extending outside the parafoveal region is called *peripheral vision*, and the information received by the visual system in this area is dramatically lower than the two innermost regions.

Saccadic movements are also classified into several types, some of which are automatic, thought to be manipulated largely by the oculomotor system, either to correct for the movement of the head and body (*vestibular*) or to focus the eyes by moving inward (*vergence*). It is also thought that little to no information is perceived during an eye saccade (Matin, 1974), a suppression that appears to be in part due to a central inhibition system active during visual processing (Riggs et al. (1974); Rayner (1998)). Note that processes need not have been terminated before a saccade is initiated. Given that some degree of parafoveal preview is possible, information about current or future fixations may be continued during eye movements. Thus, sometimes delays in processing are measured on the regions following the critical regions. Such areas are known as *spill-over regions*, and the effects measured therein as *spill-over effects* (Rayner and Duffy, 1986).

The direction of eye movements may be successive or regressive. Regressive eye-movements account for about 10 – 15% of total saccadic movement. At present, not much is known about the factors determining regressions. They are thought to occur as corrections to oculomotor errors, when a saccade overshoots its target, as well as attempts to integrate difficult texts in reading.

The multiple ways in which eye fixations and movements can be categorized naturally leads to an abundance of measures, sometimes divided into the relative time window they appear in perceptual processing. Early measurements of movements that reflect processing time include *first fixations* (the time spent on the very first fixation of a region, regardless of rereading) and *first-pass durations* (the time of initial reading consisting of all forward fixations in a region). Later measures include rereading times (second-pass measures) and saccades (percentage in or out of a region following initial fixations). Of measures recording eye saccades, *first-pass regressions* report the percentage of regressions made out a region after a first-pass fixation. One of the latest measures that is standardly reported is *total reading time*, the summation of all fixations in a region. Rayner (1998) reviews the last 20 years of eye-movement research studying reading, replete with history and more extensive terminology (see also Ebenholtz (2001) for an extensive overview of the physiological aspects of oculomotor systems, and Richardson and Spivey (Richardson and Spivey (2004a), Richardson and Spivey (2004b)) for a review of eyetracking applications).

4.2.0.2 Participants

40 native English speakers from New York University participated in our eye-tracking study. They were paid $10 for one 45-minute session. All participants had normal or corrected-to-normal vision.[2]

4.2.0.3 Materials

As described above, the critical contrast consisted of 28 pairs of passive sentences such as those listed in (153 – 154). The full set of sentences can be found in Appendix A.

(153) The proof of the theorem was guessed by the mathematician. (CQ)

(154) The proof of the theorem was erased by the mathematician. (Control)

We attempted to match the average length and frequency of the verbs in the critical comparison, but this was difficult given the limited number of appropriate CQ verbs. The CQ verbs used in Experiment 1 were slightly longer than the control verbs (8.1 vs. 7.3 characters respectively, $t(27) = 2.87$, $p < .05$), but also slightly more frequent (74.8 vs. 53.5 occurrences per million, based on Francis and Kučera (1982), ($t(27) = 1.92$, $p = .07$)).

In addition, due to the scarcity of CQ verbs, they were repeated 1.3 times on average (separated by 49 filler sentences on average), which might have speeded up their reading times slightly. Hence, the putative disadvantage CQ verbs have with respect to length is likely being offset by an advantage in frequency and a possible repetition advantage. In any case, we will address this putative length and frequency confound in Experiment 3, where the critical verbs are presented in isolation.

The experimental stimuli were mixed with 216 filler sentences of varying syntactic constructions and length, testing other hypotheses which will not be discussed here. The sentences were presented in a fixed random order. For counterbalancing purposes, it was necessary that each experiment within the entire set of materials contain four conditions. To achieve this, a condition that manipulated the meaning of the subject was added as a pilot (e.g., *the success of the theorem was guessed*) as well as a condition involving an overt question

[2]Special acknowledgment is owed to Steven Frisson, who ran the eyetracking experiment, analyzed the data, and wrote up the results in Harris et al. (2007).

complement although not syntactically or semantically an identity question (*how to solve the theorem was guessed*). The latter could have potentially provided a better control condition than the entity-denoting predicate condition, but unfortunately these overt question stimuli proved excessively difficult to interpret. Reading times were extremely slow and a number of participants found these sentences grammatically incorrect and/or non-sensical. The experimental stimuli were divided into 8 lists (4 lists were the reverse of the other lists), with 5 participants randomly assigned to each list. This way, each participant saw 7 CQ sentences of the form in and 7 control items of the form in (153 – 154).

4.2.0.4 Procedure

Participants were run individually on a SensoriMotor Instruments EyeLink I head-mounted eye-tracker apparatus and presentation software. The eye cameras recorded eye movements and fixations every 4 ms. Viewing was binocular, but only data from the eye that was calibrated best was used in the analyses. Screen resolution was 1600 x 1200 pixels and sentences were presented in fixed font, each letter being 18 pixels wide and 33 pixels high, with a maximum of 80 characters per line. Participants were seated 71 cm. from the display monitor. With this setup, 1° of visual angle corresponded to 2.7 characters. A chin rest was used to reduce head movements.

Participants were encouraged to read for understanding and to read at a normal pace. Once they finished reading a trial, they pressed a button on a button box to make the sentence disappear from the screen. Comprehension questions were asked after 50% of the trials (counterbalanced across conditions). Half of the questions required a yes response, half a no response. Participants answered the questions by means of two buttons on the game console. Accuracy was high at 90%.

A calibration procedure was performed at the beginning of the experiment and was repeated whenever the experimenter felt necessary. Before each trial, a fixation box coinciding with the position of the first letter of the upcoming sentence was presented on the screen. This box served as a trigger, with the sentence only being displayed if the fixation was close enough to the middle of the fixation box. After each trial, a drift correction was performed for a fixation point in the middle of the screen. The entire experiment lasted about 45 minutes.

4.2.0.5 Analyses

We report analyses on 2 regions: the past participle of the verb (e.g., guessed) and a spill-over region consisting of the following two words (by the). The following standard measures are discussed: first-pass duration (the summed fixation times on a region before leaving that region to the left or the right), first-pass regressions (percentage of regressions out of a region following a first-pass fixation), and total reading time (the sum of all fixations in a region). Fixations less than 100 ms and over 1200 ms were excluded from the analyses. We set the maximum cut-off at 800 ms for first-pass duration and at 1200 ms for total time duration. Analyses with higher cut-offs did not change the pattern of results.

4.2.0.6 Results and Discussion

Prior to all analyses, we eliminated sentences with major track losses, due to head movements or blinks, and sentences for which the first part (i.e., the subject plus the auxiliary verb) was skipped. This resulted in the elimination of less than 2% of the data. For each measure and each region, we subjected the data to separate one-way Analyses of Variance (ANOVAs), treating participants ($F1$) and items ($F2$) as random effects. All analyses are within-participants and items. Table 4.1 presents the averages, using participants' means.

Measure	Verb Region	Spill-over
First-pass reading time		
Control	278 (9.9)	283 (12.8)
CQ	304 (11.7)	287 (15.3)
First-pass regression		
Control	7.9 (1.9)	11.0 (2.4)
CQ	9.8 (2.7)	15.7 (2.5)
Total reading time		
Control	371 (13.5)	371(20.1)
CQ	420 (17.1)	396 (19.6)

Table 4.1: Experiment 1. Mean reading time durations and percentage of first-pass regressions. Note: CQ = concealed question condition. Reading times are in milliseconds, first-pass regressions in percentages. Standard errors are in parentheses.

First-pass duration analyses revealed a significant difference between the CQ and the control condition on the verb region, with the CQ verb taking on average 26 ms longer to process than the control verb $[F1(1,39) = 9.91, p < .01, MSe = 1383; F2(1,27) = 5.29, p < .05, MSe = 1498]$. The 4ms difference in the spill-over region was not significant. First-pass regression analyses did not show a significant difference in the percentage of first-pass regressions out of the verb region. For the spill-over region, a higher percentage of first-pass regressions was observed for the CQ sentences, but this difference was only marginally significant in the participants' analysis $[F1(1,39) = 3.65, p < .07, MSe = 124; F2(1,27) = 1.44, p > .24, MSe = 199]$. Finally, the total time analyses showed a significant effect on the verb region, with reading times for the CQ verbs 49 ms longer than the control verbs $[F1(1,39) = 10.41, p < .01, MSe = 4583; F2(1,27) = 4.77, p < .05, MSe = 6562]$. The 25 ms difference on the spill-over region was not significant.

The results clearly indicate that, similarly to complement coercion, concealed questions require additional processing time compared to control sentences. However, the time-course of the effect was earlier than what has been reported for complement coercion, which has generally been found to affect second pass reading time. The concealed question effect, however, emerged immediately during first-pass processing on the question-selecting verb itself. This provides initial evidence for distinct mechanisms for the two different types of type-mismatch. In Experiment 2, we used MEG to determine to what extent the concealed question processing cost can be localized in specific areas of the brain and to what extent these areas overlap with the ventromedial prefrontal areas implicated for complement coercion.

4.3 Magnetoencephalography (MEG)

Magnetoencephalography (MEG) is a brain-imaging technique that measures the magnetic flux pattern generated by the flow of electrically charged ions which occurs as a result of the signalling patterns among neuron groups. The loci of the electromagnetic current are referred to as *current sources*, which are represented in source models as *dipoles*, localized and directed currents. The current sources give rise to two types of electromagnetic signal that are used to estimate the direction of the electrical current associated with the sources. The first type of current, *volume currents* or *electrical potentials*, is propogated from the source throughout the volume of the brain and is measured by the closely related electroencephalography (EEG) instrument. These current patterns are electrical and follow the lines of least resistance in the brain and out of the brain casing, ultimately giving rise to irregular and distorted patterns outside the skull, where the EEG can measure them. The second type of current pattern, *magnetic flux*, propagates along a far more regular path as the resistance of the magnetic field to the bone and tissue in the head is practically non-existent. In MEG, magnetic flux (field) patterns are then recorded on the surface of the skull by magnetometers, superconductive sensors placed parallel to various positions around the head. Magnetometers measure the various minute fluctuations in the field pattern, giving rise to generalization of the magnetic current sources generated by the sources of primary current (Hämäläinen et al., 1993).

Once the field pattern is recorded, the location of cerebral current sources that underlie the magnetic field distribution are estimated from several potential hypothetical sources consistent with the distribution. Known as the *inverse problem*, there is no unique solution for deriving the actual source from the flux recordings. Several parameters defining the head shape contribute to the goodness of a particular solution to the inverse problem, including estimates of the strength of the source, its orientation, and an approximation within Cartesian space by three fiducial points, functional anatomical landmarks collected prior to recording (see Papanicolaou (1998)). Individual dipole sources were fit on a grandaveraged condition (the average of all averaged conditions) by hand using solutions provided by the

BESA brain-imaging software. The locations of the dipoles were checked against an estimate of the current density (the minimum norm), and the directionality was confirmed with a visual representation of ingoing and outgoing field patterns. Individual decisions for different events in the time course of activity examined are discussed in more depth below.

Having briefly introduced the MEG method, we now discuss the two experiments performed on concealed question materials using the MEG instrument. In the first experiment, subjects performed a sentence level sensicality judgement task, in which they judged whether sentences were grammatical or not immediately following visual presentation. In the second experiment, they performed a lexical decision task, in which they judged whether a string of letters was a real word immediately upon presentation. The stimuli in the second experiment included the CQ and control verbs from the first, along with additional non-word fillers, to determine whether any effect observed during online processing could be attributed to lexical level properties of the verb. The experiments and results are now presented in turn.

4.3.1 Off-line sensicality task

4.3.1.1 Procedure

Nineteen monolingual native English speakers with normal or corrected-to-normal eyesight participated. They were paid $40 each for one 2-hour session.

4.3.1.2 Materials

The materials of Experiment 2 were identical to those of Experiment 1, except for two minor changes. First, in order to achieve a consistent placement of triggers across our critical materials and the filler items (testing other hypotheses), the Control and CQ stimuli from the eye-tracking experiment were modified to include additional auxiliary verbs (such as must be). Second, one item from Experiment 1 was excluded from Experiments 2 and 3 because it contained the verb asked, which has recently been argued not to take a CQ complement (Nathan, 2006).[3] In addition, its exclusion allowed the exact balance of repetition of verb presentation; there were 9 verbs in the CQ condition and each verb was presented 3 times. An anomalous control condition was added to the design. The experimental task was an off-line sensicality judgment. The materials of Experiment 2 are listed in full in Appendix B.

(155) The proof of the classic theorem must be guessed by the brilliant mathematician (CQ)

(156) The proof of the classic theorem must be erased by the brilliant mathematician (Control)

[3]Taking this verb out from the eye-tracking data resulted in slightly higher levels of significance as the reading time patterns for this verb were the reverse of the standard pattern found.

(157) The proof of the classic theorem must be docked by the brilliant mathematician (Anomalous)

As with the previous materials, CQ verbs were longer (8.2 characters) than either Control (7.4 characters) or Anomalous verbs (7.5 characters) $[F(2,26) = 6.15; p < .01]$. Frequency differed between conditions overall $[F(2,26) = 3.26; p = .05]$, CQ verbs being more frequent in pair-wise comparisons than Anomalous (Scheffe test, $p < .05$), but not Control (Scheffe test, $p = .22$) verbs. Again, we expected that the lower frequency (Francis and Kučera (1982)) of Control (45.52 occurrences per million) and Anomalous (38.1 occurrences per million) verbs over CQ verbs (63.8 occurrences per million) would compensate for the shorter letter string length. Further, there was no difference in cumulative root frequency between any of the conditions $[F(2,26) = .02; p = .98]$.

Additionally, we examined whether there were differences in plausibility between the CQ and Control sentences. Thirty monolingual native speakers of English participated in a separate study that rated the plausibility of the materials on a scale of 1 (implausible) – 7 (very plausible). Materials were counterbalanced into 3 randomized lists containing equal numbers of items per condition. Analysis revealed a reliable effect of condition $[F(2,28) = 56.38; p < .0001]$. Crucially, pair-wise comparisons confirmed that Control (M = 5.3) and CQ (M = 5.2) sentences were rated as equally plausible as the CQ sentences (Scheffe, $p = .99$). Not surprisingly, anomalous (M = 2.2) sentences were rated less plausible than Control (Scheffe, $p < .0001$) and CQ (Scheffe, $p < .0001$) sentences. Thus, a difference in plausibility could not explain an effect between CQ and Control stimuli.

Finally, we also assessed the cloze probabilities of the critical verbs in their sentential context. Cloze probability is known to be one the main factors explaining variance in the amplitude of the N400 ERP (Kutas and Hillyard, 1984) . Twenty participants wrote sentence completions for each sentence type used in the experiment. The participants were presented with the experimental sentences up to the point of the critical verb (e.g., *the name of the furry animal had been*). Participants' completions were compared with the actual experimental sentences to assess how predictable the critical verbs were. Unsurprisingly, the mean cloze probability of our anomalous items was 0, but the cloze probabilities of the CQ and Control verbs were also extremely low, 0.004 and 0.002 respectively. Consequently, there was no overall effect of condition on cloze probability ($F < 1$). Thus our critical verbs were in general unpredictable, ruling out cloze probability based explanations of any measured MEG effects.

4.3.1.3 Procedure

Subjects lay in supine position in a dimly lit magnetically shielded room and viewed the stimuli through fiberoptic goggles (Avotec, FL), which were adjusted to their vision. The sentences were preceded by a fixation cross in the middle of the screen which was presented until the participant initiated the trial by pressing a button. Sentence stimuli were presented word by word (300 ms on, 300 ms off) in non-proportional Courier font (size = 70). A question mark appeared in the center of the screen after each sentence, at which point the participant had four seconds to indicate whether the sentence made sense or not. Subjects were instructed not to move their head or blink while reading and were encouraged to rest as often as necessary.

Neuromagnetic activity was recorded with a whole-head, 148-channel neuromagnetometer array (4-D Neuroimaging, Magnes WH 2500). The sampling rate was set to 679 Hz in a band between 0.1 and 200 Hz. The recording lasted approximately 45 minutes per participant. Afterwards, a short auditory baseline test was conducted, in which one hundred 1kHz tones were presented through earpieces. Collecting these data permitted us to establish the source location of the auditory M100 response, which was used as a functional landmark in further data analysis.

4.3.1.4 Analyses

Prior to source modeling, MEG data were cleaned of artifacts and trials with incorrect responses. The data were averaged by condition using an epoch length of 900 ms, with 100 ms of prestimulus interval. The data were high-pass filtered at 1 Hz and low-pass filtered at 40 Hz prior to source analysis.

The generators of the magnetic fields were modeled as equivalent current dipoles. A multiple source model (BESA) was created for each participant on the basis of that individual's grandaveraged data across conditions. Sources were first fit at the peaks of all prominent response components between 0 and 500 ms after the onset of the critical verb and then combined together into a multidipole model. This multidipole model was then kept constant across conditions. All sensors were used in localization.

Source localizations and orientations were evaluated on the basis of two visual representations of magnetic activity: the magnetic flux pattern, representing the surface distribution of electromagnetic energy around the head, and the minimum norm, estimating current density. Only those sources consistent with the location and orientation of the magnetic field pattern and the minimum norm visualizations were accepted for analysis. Goodness of fit (GOF) did not vary significantly between conditions $[F(2, 16) = 1.3; p = .29]$; on average, the multi-dipole solutions explained approximately the same percentage of activity for the interval of 0 – 500 ms for the Control (84%), CQ (83%), and Anomalous (83%) conditions. Multi-dipole solutions explained approximately 87% of the activity from 0 to 500 ms after stimulus-onset in the grandaverage. On average, stimulus categories contained the same percent of trials (78%) per participant for all conditions.

4.3.1.5 Results and Discussion

Sensicality judgment data. For inclusion in the statistical analysis, we required that participants perform with a minimum of 75% sensicality accuracy on the main two experimental conditions (CQ and Control). Two participants were excluded on these grounds from Experiment 2. Due to a button-box failure, response times from one participant were not recorded.

Sensicality judgments were recorded at the end of each trial at the question mark and are summarized in Table 2. Analysis of sensicality judgment response time revealed a main effect of condition $[F(2,15) = 4.06; p < .05]$. The speed of sensicality judgments did not differ significantly between Control and CQ conditions (Scheffe test, $p = .97$), but was faster in the Anomalous condition over controls (Scheffe test, $p = .05$). Crucially, Control and CQ conditions did not differ in sensicality judgment accuracy (∼86%). This suggests that participants did not find the CQ condition less well-formed than the Control condition, confirming the previous plausibility ratings. As expected, anomalous conditions were judged as sensical in only a fraction of the trials (7%).

Condition	Sensicality judgement time (ms)	% Sensical
Control	747	86.5%
CQ	725	85.6%
Anomalous	656	7.6%

Table 4.2: Experiment 2. Mean sensicality judgment times and percent materials judged sensical.

Multiple source models. As described above, a multiple source model was created for each participant's data based on the grandaverage of that individual's responses to all critical verbs. We first describe the characteristics of these multiple source models, before reporting on the effects of condition on source strengths and latencies.

Figure 4.1 displays all individual source localizations (grey dipoles) and average localizations (black dipole). Auditory M50/M100 dipoles are included for reference. Figure 4.2 illustrates for a single individual the typical field patterns associated with all the major response components as well the distributed source solutions of each component, as obtained by minimum norm estimates, which were used to guide the number of dipoles that were entered into the source modeling.

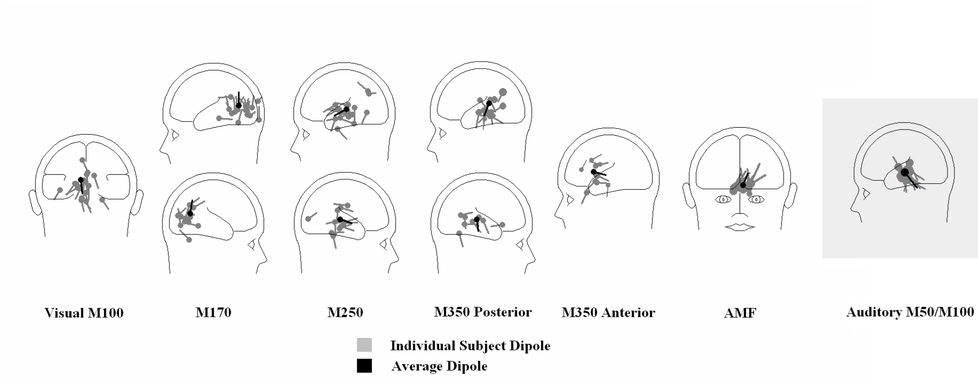

Figure 4.1: Experiment 2. Individual dipole (grey) and averaged dipole (black) locations of major MEG response components. Auditory M50/M100 dipoles were localized from a separate auditory baseline test.

The MEG signals showed a pattern of response components largely familiar from previous MEG studies of language comprehension in the visual modality (Embick et al. (2001); Fiorentino and Poeppel (2004); Helenius et al. (1998); Pylkkänen et al. (2001); Pylkkänen et al. (2002); Pylkkänen et al. (2004) ; Pylkkänen and McElree (2007); Stockall et al. (2004); Tarkiainen et al. (1999)). The only difference to previously reported activation patterns was the spatial distribution of left hemisphere activity at 300-400 ms. In this time window, previous MEG studies have consistently found a source cluster in the posterior temporal lobe, the so-called M350 (or N400m) (Embick et al. (2001); Helenius et al. (1998); Pylkkänen and Marantz (2003); Pylkkänen et al. (2004)), which has been hypothesized to index aspects of lexical access (Pylkkänen and Marantz (2003)). In addition to the posterior M350, several of our participants exhibited more anterior left lateral activity at 300-400 ms. This finding is consistent with a large body of neuroimaging studies on language comprehension, including syntactic (Kuperberg et al. (2003); Just et al. (1996); Just et al. (1996); Stowe et al. (2005)), sentence-level semantic (Baumgärtner et al. (2002); Friederici et al. (2003); Kuperberg et al. (2003)), and lexical-level manipulations (Kotz et al. (2002); Noppeney and Price (2002)), which have also reported activity in left frontotemporal regions.

Activity at 0– 300 ms. MEG data from all participants (N = 17) exhibited a magnetic flux distribution typical to the visual M100 field pattern, characterized by an out-going magnetic field over the right-occipital sensors and an in-going field over the left-occipital sensors. This activity was best explained by a single midline dipole over occipital sensors in all participants. After the visual M100 peak, activation spread to occipito-temporal areas bilaterally, peaking at $\sim$150 – 200 ms (M170). In general, the M170 peak was associated with an out-going field over the left-occipital sensors, re-entering over the right-occipital sensors. Bilateral sources in the left and right occipital areas best explained the M170 activity for most participants (N

= 9), and in only left but not right occipital areas for others (N = 5).

Activity then spread to posterior temporal areas, peaking between ~200 – 300 ms (M250). Bilateral two-dipole solutions best accounted for activity within the M250 time window for most participants (N = 9). Some participants displayed activity only in the left (N = 2) or right hemisphere (N = 2). Additional posterior sources were fit in the left (N = 5) and/or right (N = 2) hemisphere to create physiologically accurate multi-dipole solutions.

Midlatency Left Hemisphere Activity. Previous research indicates that areas clustered around the left hemisphere auditory cortex are sensitive to semantic congruence (Helenius et al. (1998); Helenius et al. (2002)). We used the localization of the auditory cortex as a functional landmark to classify midlatency left hemisphere activity as either 'posterior' or 'anterior'. Left hemisphere sources at 300-400ms was labeled as 'anterior M350s' (M350-A), if they localized anterior to the auditory cortex and as 'posterior M350s', if they localized posterior to the auditory cortex. The auditory cortex was located on the basis of the neural response to 100 1kHz tones in an auditory baseline test. Localization was performed either at the auditory M100 component (N = 11) or at the M50 component (N = 1), whichever field pattern was the clearest (for evidence that the two components are generated by similar cortical regions, see Mäkelä et al. (1994)). Auditory baseline data were not available for 5 participants; for those participants we used the mean auditory cortex localization obtained from the other 12 participants.

In total, 14 participants were modeled with left hemisphere sources in the M350 time window. The data from 6 participants were most accurately explained by both an M350-P and M350-A source in the left hemisphere, whereas 5 participants showed exclusively posterior and 3 participants exclusively anterior activity. As in previous studies, the M350-P was associated with an out-going field over left posterior areas, re-entering over sensors covering or anterior to the auditory cortex. For seven participants, accurate localization of the M350-P field pattern required an additional right hemisphere dipole (cf., Pylkkänen and McElree (2007)). In some participants, it was necessary to model other sources as co-active with the M350-P source: a parietal midline (N = 2), occipital (N = 4), or a frontal right lateral (N = 2) source. The M350-A, which closely followed the M350-P, was typically associated with an in-going magnetic field in areas superior to the left anterior temporal lobe and an out-going field over inferior left-hemisphere sensors.

Midlatency Right Hemisphere Activity. Eight participants showed right hemisphere activity at 300 – 400ms. This activity was most often coactive with the M350-P left temporal source, consistent with the findings of Pylkkänen and McElree (2007). Right hemisphere activity ranged over the entire right temporal and frontal lobes and was generally characterized by an outgoing field over sensors covering the lateral sulcus and an ingoing field in areas superior to this, although the precise orientation varied between participants.

Anterior Midline Field (AMF). Fourteen participants showed an Anterior Midline Field (AMF) at ~350 – 500 ms, which was characterized by an out-going field over frontal right hemisphere sensors, and a re-entering field over frontal left hemisphere sensors. As discussed above, this component has been implicated in the cost of complement coercion (Pylkkänen and McElree (2007)).

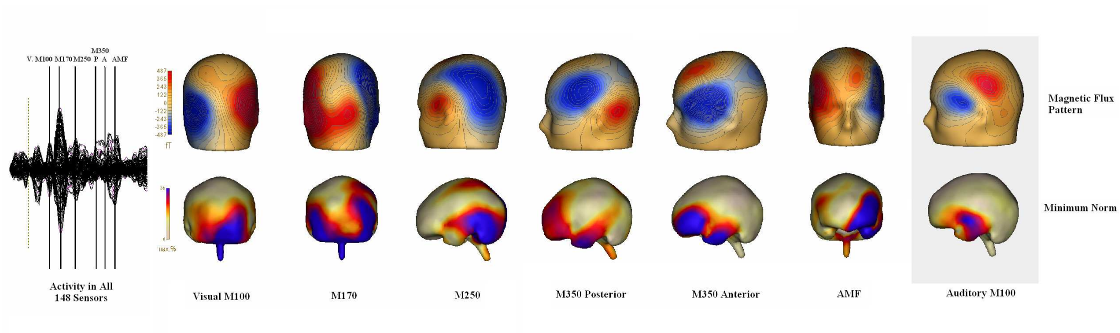

Figure 4.2: Experiment 2. Major component responses to visually presented words in MEG from a single subject. Estimates of magnetic flux distributions are shown on top and minimum norm activations on the bottom. A typical auditory M100 response is shown separately.

No effect of condition for sources from 0 – 300 ms. No reliable effect of question concealment or semantic anomaly was obtained either in peak latency or amplitude for sources at 0 – 300 ms. This is consistent with previous research where effects of semantic manipulations have only been found within later time windows (Helenius et al. (1998); Pylkkänen and McElree (2007)). The only early source that showed any effect of the stimulus manipulation was the left hemisphere M250 source, where pair-wise analyses revealed a marginal delay for concealed questions as compared to controls (Scheffe test, p = .11), although the main effect of condition was not significant $[F(2,12) = .82; p = .45]$.

Instead of the AMF, left lateral effect of question concealment in the extended M350 time window. Contrary to the hypothesis that complement coercion and question concealment involve the same semantic repair mechanisms, the AMF showed no effect of condition in either latency or amplitude (both $Fs < 1$). Instead, a main effect of condition was found in the peak latencies of both the M350-P $[F(2,10) = 5.99; p < .01]$ and the M350-A $[F(2,8) = 4.67; p < .05]$. Pair-wise analyses revealed that the M350-P component peaked later for both CQ (Scheffe test, $p < .01$) and Anomalous (Scheffe test, $p < .05$) conditions than controls. The M350-A component, on the other hand, peaked reliably earlier for the CQ than for the Control condition (Scheffe test, $p < .05$). The Anomalous condition did not differ reliably from controls (Scheffe test, $p = .16$). Figure 4.3 summarizes the M350 left hemisphere latency results. In addition, there was a trend towards diminished M350-A amplitudes $[F(2,8) = 3.23; p = .07]$ in the Anomalous condition as compared to sensical conditions, which showed no difference in a pair-wise comparison (Scheffe, $p = .99$). The stimulus manipulation had no reliable effect on the amplitudes of the M350-P source, and no effect on either peak latencies or amplitude of right hemisphere sources. Table 4.3 shows the means for peak latency and amplitude for all sources in all conditions.

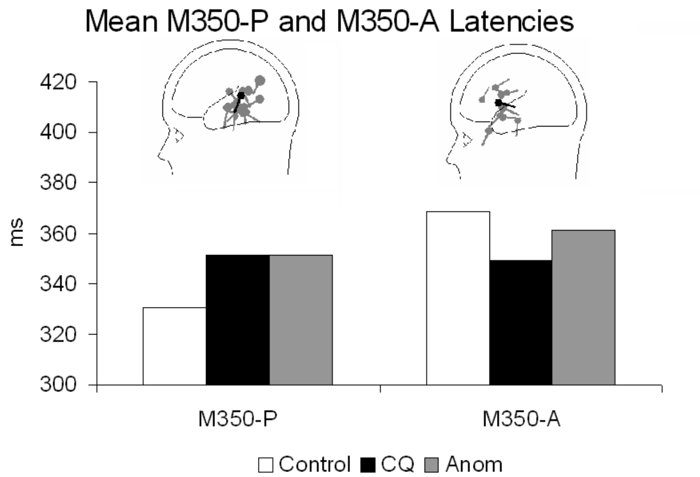

Figure 4.3: Experiment 2. Mean peak latencies of M350-P and M350-A components. Represented inside a model of a head, individual (grey) and average (black) dipole locations are shown. We observed a reliable effect of CQ peak latency for both dipole locations.

In sum, Experiment 2 aimed to identify the specific neural components underlying the processing cost elicited in Experiment 1. We found a reliable delay at ~350 ms in the left posterior temporal lobe. This effect did not extend to anterior temporal regions, which in fact showed the opposite effect for CQs. As shown in Table 3, in the control and anomalous conditions, the left anterior temporal sources peaked somewhat later than the left posterior temporal sources. For CQs, however, the anterior and posterior sources peaked approximately at the same time. It is possible that the simultaneity of the anterior and posterior activity for the CQs reflects the involvement of both regions in CQ-type-shifting. However, at this point this hypothesis remains a speculation only, in particular since only 6 of our subjects exhibited both an anterior as well as a posterior left temporal source. Importantly, the AMF showed no sensitivity to question concealment. Thus, the overall pattern of our results suggests distinct processing mechanisms for concealed questions compared to complement coercion.

Given that our manipulation varied the target verb, Experiment 2 cannot, however, conclusively determine whether the latency effect in the posterior M350 reflects an increased processing cost associated with type-mismatch or difficulty accessing relevant lexical properties of the question-selecting verbs. We addressed this question in Experiment 3.

Source Dipole	Latency (ms)			Amplitude (nAm)		
	Control	CQ	Anomalous	Control	CQ	Anomalous
Visual M100	117	118	119	38.11	35.21	32.94
M170 LH	176	175	176	16.71	20.77	18.32
M170 RH	173	175	184	19.88	15.32	19.52
M250 LH	250	157	255	21.70	22.54	21.40
M250 RH	265	267	270	21.05	24.67	20.83
M350-P LH	331	351	351	10.49	12.45	13.05
M350-P RH	355	357	354	20.44	19.14	11.34
M350-A LH	369	349	361	12.57	12.85	9.00
AMF	405	404	398	27.97	24.56	26.74

Table 4.3: Mean source latencies and amplitudes in all conditions for all components.

4.3.2 Lexical Decision Task

4.3.2.1 Participants

The same 19 participants who participated in Experiment 2 also participated in Experiment 3.

4.3.2.2 Materials and Procedure

The lexical decision study was initiated approximately 30 minutes after the sentence-level study had been completed. The two experiments were separated by the auditory baseline test as well as a significant rest. Consequently, repetition priming from Exp 2 to Exp 3 was extremely unlikely.

The verbs of Experiments 1 and 2 served as the word stimuli and the nonwords were generated using the English Lexicon Project (Balota et al. (2002)). Nonwords were pairwise matched to the lexical level factors of each verb. The stimuli were presented after a centered fixation cross was displayed for 500 ms. The items were visible until the button press response. Subjects were given four rest breaks. The recording lasted approximately 20 minutes per participant.

4.3.2.3 Analyses

The same procedure for filtering, cleaning, and averaging the MEG data was followed as in Experiment 2, except that data were averaged on a smaller epoch interval, ranging from -100 to 500 ms. As before, only trials with correct responses were included in analyses.

We focused our analyses on the source that showed a processing delay in the previous experiment, i.e., the M350-P. Since there is no standard method for comparing activity between a sentential context and isolation, we analyzed the data in the following two ways.

First, we imported the multi-dipole models from the M350 time window in Experiment 2 into each averaged condition of Experiment 3. By importing the sources from Experiment 2, we kept source location and orientation identical between the experiments, allowing us to assess to what extent the M350-P, as localized in a sentential context, showed the same effect in isolation. However, while importing sources permitted direct comparison of how context affected the M350-P, the imported solutions were modeled from different data and thus were not necessarily accurate models of activity measured in Experiment 3. Therefore, we also created multi-dipole solutions for the M350 time-window by using the data measured in the lexical decision task. As is shown in Figure 4.4, posterior M350 sources localized very similarly whether the lexical decision or the sentential data were used, suggesting that similar activity was indeed elicited in the two different contexts. Having established that the presence of the M350-P was not context dependent, we proceeded to examine whether the effect it showed in Experiment 2 was.

4.3.2.4 Results and Discussion

Lexical decision data. The behavioral lexical decision data are summarized in Table 4.4. There were no differences in lexical decision time between conditions overall [$F(2,17) = 1.53$; $p = .23$], and the accuracy difference were less than 2% between the Control conditions and the others (97% accuracy on both CQ and Anomalous conditions versus 99% accuracy on Control conditions). Both measures suggest that participants found the two verb classes equally easy to process. As expected, nonwords were significantly slower than all real word conditions (M = 756 ms) and were judged less accurately on average (86% correct).

Condition	Decision Response Time	% Correct
Control	629	99.0%
CQ	653	97.4%
Anomalous	646	97.1%

Table 4.4: Experiment 3. Mean lexical decision times and percent correct.

MEG data. As already noted above, modeling M350-P sources directly from the lexical decision data revealed a close correspondence between patterns of activation and source locations in the two experiments. All but two M350-P sources fit in the sensicality experiment had corresponding activity in the lexical decision task. M350-P dipoles did not differ in location or orientation between types of models. There was no difference in goodness of fit recorded at the M350-P peak [$F(1,9) = .56; p = .58$] or between the M350-P time window (300 – 400 ms), [$F(1,9) = 1.86; p = .20$] between imported dipoles and dipoles modeled directly from the lexical decision data. Figure 4.4 shows the mean latencies and individual locations of both the imported dipoles and the modeled dipoles. Our results clearly show that the latency effect obtained in Experiment 2 was absent when the sentential context was removed: There was no significant effect of condition on either M350 peak latency or amplitude in either analysis.

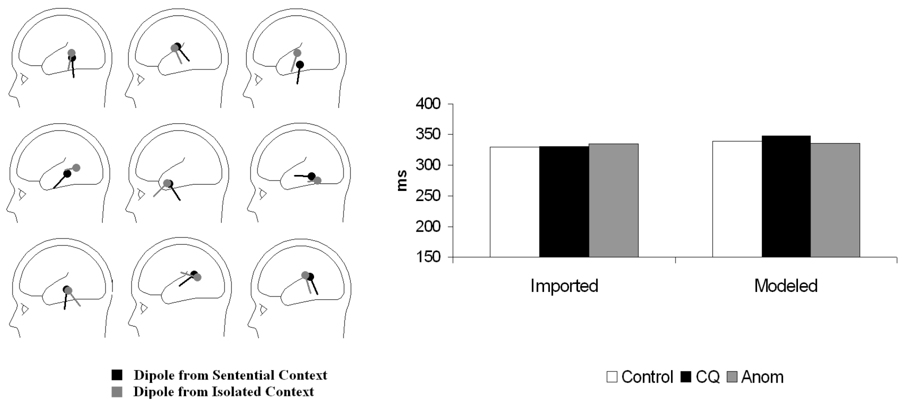

Figure 4.4: Dipoles obtained from Experiment 2 (grey) versus Experiment 3 (black) are shown for each individual subject with M350-P activity in both experiments. Mean peak latencies for dipoles modeled from Experiment 2 (imported) and Experiment 3 (modeled) are shown on right; no effect of verb type on peak latency for any condition within Experiment 3 was obtained.

To further assess the dependence of the latency effect on the sentential context, M350-P latencies from Experiments 2 and 3 were entered into a 2 × 3 ANOVA testing for the interaction between context (sentential vs. isolation) and verb type (CQ, Anomalous, Control) (Fig 4.5). When imported multi-dipole models were used for the M350 values of Experiment 3, the interaction was reliable $[F(2,10) = 4.01; p < .05]$. When the latencies of the M350s modeled on the basis of the lexical decision data were used instead, the interaction did not quite reach significance $[F(2,14) = 2.93; p = .09]$.

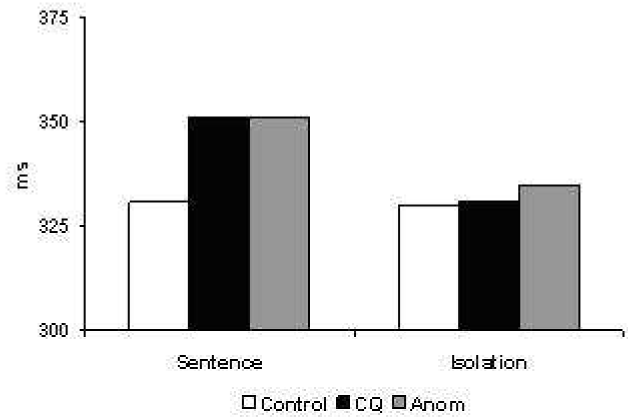

Figure 4.5: Interaction between context and verb type (Experiments 2 and 3

In conclusion, the results of Experiment 3 show that when the critical verbs of Experiments 1 and 2 are presented in isolation, they elicit no effect in the latency of the posterior M350 source. This suggests that the effect obtained in Experiment 2 was due to the type-mismatch between the CQ-verbs and their object, and not to the meaning difference between the CQ-verbs and the control verbs themselves.

4.4 General Discussion

To briefly recapitulate the methods and results of the study, concealed questions were found to elicit a processing cost in both eye-tracking and MEG experiments during online comprehension. In eye-tracking, there were significantly prolonged fixations during first-pass duration measures on question-selecting verbs (*guessed*) following a CQ complement than for comparable controls (*wrote*). In the MEG sensicality task, the same verbs elicited a delay in the left hemisphere M350 component, an area thought to involve semantic integration during that time window. In the MEG lexical decision task, these verbs did not elicit any pattern distinct from control verbs when presented in isolation. The consistency of the findings across experimental paradigms, along with the MEG control, suggests that CQs constructions are costly to interpret. Table 4.5 displays the observed results:

Method	Summary of Effect	Time Course of Effect
Eye-tracking	Longer first-pass fixations on verb	Early, around 300 ms
MEG	Delayed peak of left posterior M350 source when verbs were presented in sentence, none when in isolation	Late, around 350 ms

Table 4.5: Eye-tracking and MEG results compared

As this is the first study on concealed questions, and one of the first on semantic repair using the same materials in eye-tracking and MEG, these results are interesting and informative in their own right. However, several issues regarding how the results are to be interpreted remain at large, including the compatibility of the results obtained from eye-tracking and MEG, the relation of CQs to complement coercion, and the role of the left posterior M350 in language processing. These topics are briefly discussed in turn.

4.4.1 Compatibility between paradigms

In the study above, eye-tracking effects were obtained at approximately 50 ms earlier than those obtained in MEG. Comparing results obtained from eye-tracking and those obtained from MEG does not, to my knowledge, have an established precedent in the psycholinguistic literature. The correlation between different methods is, in general, quite difficult to establish, as each tool used to explore cognition has its own limitations. For instance, Horwitz and Poeppel (2002) warn that it may not be possible to show that hemodynamic imaging methods, such as fMRI and PET, measure events or cognitive processes within the brain that are comparable to the temporal data collected by electromagnetic measures, such as MEG and EEG. The success of multimodal imaging is mitigated by the fact that these methods do not meaningfully record the same neurophysiological phenomena.

We might expect that behavioral methods, such as eye-tracking, are confounded by recording a relatively indirect measure of cognitive processes via a system that is mediated by independent motor control factors. MEG technologies, on the other hand, provide a more 'direct' measurement, by recording the minute fluxations of the magnetic field around the head, which are thought to echo the patterns of primary curent sources. So, it may seem surprising that an 'indirect' measure should precede a more 'direct' one, in that the effect in eye-tracking appeared in an earlier time window than the effect in MEG.

In fact, the latency and relative order of the effects reported in our results are consistent with a general trend observed in recent studies. For instance, lexical level factors, such as semantic priming and frequency, are known to modulate the N400 pattern in ERPs (Kutas and Federmeier, 2000). Similar lexical level effects are found in first fixation times in eye-tracking measures (see Staub and Rayner (2007)). Thus, manipulation of lexical level factors correlates with different effects in different experimental methods. In this case, as with our own, eye-tracking effects appeared earlier than those observed in electromagnetic recordings.

Confirmation from other work aside, it is still unclear what the relationship could entail for models of lexical access and semantic integration in real-time processing studies. Several hypothetical models are plausible here, and we cannot review them all. One aspect of the relationship to consider is the difference in visual presentation of the material in the experiments. In eye-tracking, the entire sentence was available to subjects all at once. The parafoveal preview available on pre-target words might have aided lexical access of the target, thought to be active 70–140 ms after making the first fixation (Inhoff et al., 2005), thereby inducing an early effect. As sentences were presented one word at a time in the MEG experiment, subjects would not have been able to exploit preview in this fashion. Another possibility to consider is that the early eyetracking effects need not correspond to component peaks in the MEG signal, but rather could be related, at least in part, to the slope of summation of these peaks. Whatever the relation between these different sorts of effects in the present study, the relationship requires much further inquiry before we can make any claims about how effects are related across paradigms, and remains an important goal for future study.

4.4.2 Concealed Questions and Complement Coercion

This chapter began by presenting two plausible hypotheses regarding the processing of non-compositional constructions in general. Previous studies found that a superficially similar construction, complement coercion, which has been posited to require semantic repair for interpretation, elicit relatively late effects in prior eye-tracking studies, and high amplitudes in a component localized in the ventromedial prefrontal cortex. Table 4.6 compares the results from the present experiment on concealed questions and those previous experiments on complement coercion.

Concealed questions and complement coercion did not elicit qualitatively similar sorts of effects in either eye-tracking or MEG measures. In particular, the *latency* of the effect for CQs was earlier than for those observed for the processing of complement coercion nouns

	Manipulation	
Measurement	Concealed Questions	Complement Coercion
Eye-tracking	earlier delayed processing	later processing effects
MEG	delayed M350-P peak	larger AMF peak

Table 4.6: Concealed questions and complement coercion compared

in eye movement studies. In the MEG not only was the type of effect different in that CQs elicited a delayed peak in the response signal, while coercion nouns elicited greater amplitudes, but the actual *location* of the effect was distinct: CQs affected the left posterior region of the brain, while coercion affected a later, frontal component. Much remains to be said about how these areas are involved in language processing, although converging evidence suggests that the left posterior area may be involved in lexical access, our findings clearly do not support the view that there is but one semantic repair mechanism responsible for both complement coercion and concealed questions. In terms of the specific, if narrow, hypothesis presented in §4.1.1, the evidence collected here supports Hypothesis 4.1.4 over Hypothesis 4.1.3, although much more evidence is required to give confirm the generality of the claim.

That the evidence patterns against a unified type-shifting mechanism for nominals may not surprise semanticists, as the type resulting from the shift is thought to be quite different. In the case of CQs, a noun is interpreted with a question-like meaning, whereas in complement coercion the object denoted by the noun is reinterpreted as an event.

Viewed differently, however, we might interpret these results as showing that the interepretation of CQ does not involve type-shifting or semantic repair after all. Simply because there was an delayed processing effect does not automatically indicate that the supposed difficulty encountered must be cashed out in terms of semantic repair. Without further evidence to the contrary, it might be thought that the difficulty the parser encounters when interpreting CQs could be due to a delay elicited when the parser selects the appropriate sense of a polysemous word, rather to a delay invoked when computing a lexicalized type-shifting operation.

Two types of experiment speak against this interpretation. First, eye-tracking studies on polysemous words have not shown familiar polysemy to be costly. For instance, Frisson and Pickering (1999, 2007) found that productive metonymic relationships did not tax the parser, unless they exploited unfamiliar terms in the metonym. In particular, they examined the *producer-for-product* metonymy, e.g., the name of an author to represent his work (*Dickens*), and *place-for-event* metonymy, e.g., a place name for a salient event occurring there (*Vietnam*).[4] Here, polysemous senses did not elicit any delay over controls. The lack of effect observed in these studies does not explain the early effect found in the CQ materials. Second, Pylkkänen et al. (2006) recently found that when two senses of polysemous string were in competition, the effect localized in the *right* temporal lobe. Given that the CQ latency effect

[4]See Nunberg (1995) for a brief overview of standard metonyms.

was found only in the left hemsiphere, the MEG evidence does not support the polysemy account either.

All this discussion brings us to our final question to consider: what role does the localized area contribute to language processing in general, and why would it be implicated in the processing of concealed questions?

4.4.3 The role of the left posterior area in semantic processing

In the sensicality experiment reported in §4.3.1, processing concealed question constructions correlated with a delayed M350 component response in the posterior left hemisphere. The role of the left posterior region of the brain has long been implicated in language processing. The relation of the area to semantic processing was made famous in early studies by Wernicke in the late 19th century, in which aphasic patients with damage to the posterior left hemisphere regions displayed increased difficulty understanding and producing well-formed utterances (see Poeppel and Hickok (2004) and Stowe et al. (2005) for history and recent challenges to the traditional view).

More recently, Piñango and Zurif (2001) found that patients with Broca's and Wernicke's aphasia had increased difficulty comprehending semantically complex constructions, in particular aspectual and complement coercion. In subjects without brain-lesions or deficit, ERP studies have found that semantically anomalous materials, including words, pictures, sounds, are accompanied by a sustained negativity at 400 ms (the N400).

In addition to the effect of question concealment, the left posterior M350 was also modulated by semantic anomaly, traditionally associated with the N400 response in ERPs (Kutas and Hillyard, 1980). The primary generator of the N400 has been localized to the posterior parts of the left superior temporal gyrus in previous MEG studies as well (Halgren et al. (2002); Helenius et al. (1998)). This finding is further corroborated by deficit-lesion studies (for review see van Petten and Luka (2006)), which have found evidence of a dependency of the N400 effect on the left hemisphere in "split-brain" epileptics (Kutas et al., 1988). Also, several studies have found that damage to the right hemisphere reduces but does not eliminate the N400 (Hagoort et al. (1996); Kotz and Friederici (2003); Kotz et al. (1999); Swaab et al. (1997)). Finally, Friederici and Kotz (1999) reported that damage to anterior or frontal areas does not eliminate the N400.

Our results are consistent with a large body of evidence implicating the posterior left temporal region of the brain in language processing. More specifically, there is reason to believe that activation in the region during later time windows is involved in lexical access. However, only materials from the sensicality experiment modulated effected the M350-P component, suggesting that the M350 might reflect the application of semantic repair mechanism, in addition to activations of lexical representations. Naturally, further research is required to better determine the role of the M350-P in online language processing.

4.5 Summary of Experimental Findings

This chapter presented results from the first psycholinguistic experiment testing online comprhension in two methodologies: eyetracking and MEG. It was discovered that concealed questions elicited delayed reading times in eyetracking, and were recorded with delayed M350-P components in MEG. The results discussed here are consistent with the view that concealed questions involve a semantic repair operation, which elicits different behavioral and neural responses from those costs found in complement coercion. Further, a separate lexical decision task in MEG confirmed that the cost localized in the posterior left temporal lobe reflected online comprehension, and not independent differences between the critical verbs.

5
Concluding Remarks

In the preceding four chapters, we have covered much ground regarding (i) a fuller description of concealed question contexts, (ii) the comparison of past and present semantic proposals for concealed questions, evaluated in terms of how well they hold up to (ii.a) linguistic and (ii.b) philosophical considerations, and (iii) evidence from a psycholinguistic experiment, which used a variety of methods to measure the cost of interpreting concealed questions online. For all these separate areas of research, we have devoted little space to shoring up the connections between them. As such, there are numerous outstanding issues to address, most of which are far beyond the scope of this chapter.

For instance, a specific semantic interpretation of the syntactic concept of Late Merger, in which a relative clause is adjoined "late" to a nominal head, was provided in chapter 3. It was assumed that the late addition of syntactic material likewise corresponded to a delayed availability of semantic material in semantic composition, and which was then cashed out in terms of dynamic conjunction of sentences. This would perhaps imply that semantic material is made available at different times in the interpretation of a sentence. Such a claim is not uncontroversial, and was not explicitly defended.

Another remaining issue is the relation between semantic compositionality and language processing. Chapter 4 presented psycholinguistic evidence that concealed questions are more difficult to process than comparable controls. It was then argued that the observed processing cost reflected additional representational structure of the CQ. Note, however, that the psycholinguistic effects were never argued to directly reflect the computation of the type-shifting rule $\mathcal{Q}$, although such a correspondence should not be ruled out without further evidence. The fact of the matter is that the MEG results merely indicate that components active during online comprehension of CQs had a different distribution than controls or complement coercion constructions did. Although we speculated about how these results might be incorporated into what is already known from existing research, we cannot yet distinguish more acutely what sorts of processes might have been engaged as subjects interpreted CQs, or why these might have influenced cognitive resources differently. Nonetheless, the results of the experiment were broadly consistent with an account that requires CQs to shift their semantic type.

In general, this thesis has taken the middle ground on the continuum between two extremes. On one end, semantics is autonomous from processing; as a matter of method, it might be argued, experimental evidence does not weigh in heavily on semantic theorizing. On the other end, semantics and psycholinguistics are fully incorporated disciplines. According to this view, our best semantic theories posit operations that are clearly reflected in experimental measurements. The position I take is somewhere between these two extremes, for while I believe that psycholinguistic evidence can be marshalled in semantic argumentation, I also believe that the semantic argumentation can and should exist on its own. Hopefully, careful consideration of *when* and *how* psycholinguistic evidence can best inform our semantic theories will determine the extent to which semantics and psycholingusitic can interface. Demarcating such a program was well beyond the scope of the present task.

Lastly, the approach makes several empirical predictions, and it is unclear whether they obtain. For instance, we expect that if the semantic account of Late Merger is correct, then *no* late merged adjunct or relative clause will be interpreted as a propositional attitude. Assuming Fox's (2002) account of ACD, this means that embedded clauses with ACD structures do not support propositional attitude interpretations. In terms of cross-linguistic typology, we also expect that languages *could* have two separate phonological forms for the $know_{CQ}$ and $know_{prop}$. If such a language had VP ellipsis, it would be expected that $know_{prop}$ would not be licensed in ACD contexts, assuming that the syntactic distinction between internally and externally headed relative clauses is correct and universal across languages. Data for these predictions were not provided, and remain an invitation for future research.

In conclusion, a brief summary of the most important points is compiled below. The core properties of concealed questions were critically examined and, in certain cases, modified. Table 5.1 summarizes the important descriptive properties and their analyses.

	Characteristic	Analysis
1	CQs are interpreted as indirect identity questions	CQs and identity questions partition logical space in similar ways
2	CQs fail entailments in opaque contexts	CQs are subject to pragmatic restrictions similar to concepts in conceptual covers
3	CQs do not support gendered anaphora	CQs do not denote persons
4	CQs do not conjoin with individuals under individual-selecting verbs	CQs are not individuals; cannot conjoin with individuals
5	CQs prefer names as answers to their identity question paraphrases	CQs subject to pragmatic constraints such that names preferred to descriptions as names are more restrictive; also explained exceptions to rule
6	ACD environments do not allow Reading B	Relative clause in Reading B is internally headed; yields a proposition
7	Sluicing environments do not allow Reading A	Relative clause in Reading A is externally headed and Late Merged

Table 5.1: Central Characteristics of CQs and their Analyses

In chapter 3, a semantic and pragmatic account of CQs was presented that crucially used Aloni's (2001) notion of conceptual covers, sets of individual concepts that uniquely and exhaustively specify the domain of individuals. In short, it was argued that CQs are formed via a specific type-shifting operation $\mathscr{Q}$ on the CQ NP α, as shown in (158). The semantics for CQs were expressed according to two different conceptualizations, highlighting distinct aspects of their meaning (159).

(158) CQ Type-shifter $\mathscr{Q}$

$$\langle s,e\rangle \mapsto \langle\langle s,e\rangle,\langle s,\langle s,t\rangle\rangle\rangle$$
$$\alpha_{\langle s,e\rangle} \mapsto \lambda c_{\langle s,e\rangle}.\lambda w.\lambda w'.c \in \wp_i(a)[c(w)=\alpha(w) \wedge c(w')=\alpha(w')]$$

(159) Two Conceptualizations of CQ Interpretations:

I. CQs are relations between individuating concepts c and a question about the identity of the CQ NP, such that c identifies the CQ NP in worlds w and w'. CQs are of type $\langle\langle s,e\rangle,\langle s,\langle s,t\rangle\rangle\rangle$.

II. CQs are sets of individual concepts: $(\!|c|\!)_{\alpha,w,w'}$ the set of *concepts generated by the CQ NP α in w and w', i.e.,* concepts that can be used to specify the individual denoted by the CQ NP, given a pair of worlds w and w'. CQs are of type $\langle s\times s,\langle\langle s,e\rangle,t\rangle\rangle$

We further developed a completely compositional and partially dynamic semantics for concealed questions using Dynamic Montague Grammar (Groenendijk and Stokhof, 1990a) with existential disclosure (Dekker, 1993a). As further evidence for this approach, we found that different anaphoric continuations were possible in cases in which the extrinsic argument of the relational noun was specified overtly and cases in which the argument was not, as predicted by Dekker (1993a).

Lastly, a novel psycholinguistic experiment using both eye-tracking and MEG methodologies was presented. The results suggest that CQs are more difficult to process than controls, and were interpreted as supporting a theory of CQs which involves additional semantic structure, such as the kind generated by application of the type shifting operation $\mathscr{Q}$.

As discussed, numerous outstanding issues and puzzles remain to be addressed. It is my hope that the work here will have provoked further inquiry into a pragmatic account of concealed questions, as well as experimentation on similar non-compositional constructions, even though the present thesis cannot hope to have satisfactorily solved all the issues it has raised.

Bibliography

Aloni, M. (2001). *Quantification under Conceptual Covers*. PhD thesis, Universiteit van Amsterdam, Amsterdam, The Netherlands.

Balota, D., Cortese, M., Hutchison, K., Neely, J., Nelson, D., Simpson, G., and Treiman, R. (2002). The English lexicon project: A web-based repository of descriptive and behavioral measures for 40,481 English words and nonwords. http://elexicon.wustl.edu/. Washington University.

Barker, C. and Dowty, D. (1993). Non-verbal thematic proto-roles. In Schafer, A., editor, *Proceedings of NELS*, volume 23, pages 49–62, Amherst, MA.

Barker, C. and Jacobson, P., editors (2007). *Direct Compositionality*. Oxford University Press, Oxford, UK.

Baumgärtner, A., Weiller, C., and Büchel, C. (2002). Event-related fMRI reveals cortical sites involved in contextual sentence integration. *Neuroimage*, (16):736–745.

Beck, S. and Rullman, H. (1999). A flexible approach to exhaustivity in questions. *Natural Language Semantics*, 7:249–298.

Beck, S. and Sharvit, Y. (2002). Pluralities of questions. *Journal of Semantics*, 19:105 – 15.

Berman, S. and Hestvik, A. (1991). LF: A critical survey. Lecture notes: *Arbeitspapiere des Sonderforschungsbereichs 340* "*Sprachtheoretische Grundlagen für die Computerlinguistik*".

Bhatt, R. (2002). The raising analysis of relative clauses: Evidence from adjectival modification. *Natural Language Semantics*, 10:43–90.

Bhatt, R. and Pancheva, R. (2006). Implicit arguments. In Everaert, M. and van Riemsdijk, H., editors, *The Blackwell Companion to Syntax*, volume 2, pages 554–584. Blackwell.

Bianchi, V. (2002a). Headed relative clauses in generative syntax. Part I. *GLOT International*, 6(7):197 – 204.

Bianchi, V. (2002b). Headed relative clauses in generative syntax. Part II. *GLOT International*, 6(8):235–247.

Bonomi, A. (1995). Transparency and specificity in intensional contexts. In Leonardi, P. and Santamrogio, M., editors, *On Quine, New Essays*, pages 164–185. Cambridge University Press, Cambridge, MA.

Caponigro, I. and Heller, D. (2007). The non-concealed nature of free relatives: Implications for connectivity in specificational sentences. In Barker and Jacobson (2007), pages 237–263.

Carlson, G. (1977). Amount relatives. *Language*, 53:520–542.

Chalmers, D. J. (2007). The foundations of two-dimensional semantics. In García-Carpintero, M. and Macià, J., editors, *Two-Dimensional Semantics*, pages 55–140. Oxford University Press, Oxford, UK.

Chomsky, N. (1995). *The Minimalist Program*. MIT Press, Cambridge, MA.

Crimmins, M. and Perry, J. (1989). The prince and the phone booth: reporting puzzling beliefs. In Ludlow (1997), pages 963–992.

de Bruyn, J. and Scha, R. (1988). The interpretation of relational nouns. In *Proceedings of the 26th annual meeting of the Association for Computational Linguistics*, Buffalo, NY. SUNY.

Dekker, P. (1993a). Existential disclosure. *Language and Philosophy*, 16(6):560–587.

Dekker, P. (1993b). *Transsentential Meditations: Ups and Downs in Dynamic Semantics*. PhD thesis, ILLC, Universiteit van Amsterdam, Amsterdam, The Netherlands.

Dekker, P., Aloni, M., and Butler, A. (2005). The semantics and pragmatics of questions. In Aloni, M., Butler, A., and Dekker, P., editors, *Questions in Dynamic Semantics*, Amsterdam, The Netherlands.

den Dikken, M. (2001). Specificational copular sentences and pseudoclefts: A case study. In Everaert, M. and van Riemsdijk, H., editors, *The Syntax Companion*. Unpublished.

Dowty, D. (2007). Compositionality as an empirical problem. In Barker and Jacobson (2007), pages 23–101.

Dowty, D. R., Wall, R. E., and Peters, S. (1981). *Introduction to Montague Semantics*. Kluwer Academic Publishers, Dordrecht, The Netherlands.

Ebenholtz, S. M. (2001). *Oculomotor Systems and Perception*. Cambridge University Press, Cambridge, UK.

Embick, D., M, H., Schaeffer, J., Kelepir, M., and Marantz, A. (2001). A magnetoencephalographic component whose latency reflects lexical frequency. *Cognitive Brain Research,*, 10(3):345–348.

Fiorentino, R. and Poeppel, D. (2004). Decomposition of compound words: An MEG measure early access to constituents. In Alterman, R. and Kirsch, D., editors, *Proceedings of the 25th Annual Conference of the Cognitive Science Society*, pages 13–42, Mahwah, NJ. Erlbaum.

Fodor, J. A. and Lepore, E. (2002). *The Compositionality Papers*. Clarendon Press, Oxford, UK.

Fox, D. (2002). Antecedent-contained deletion and the copy theory of movement. *Linguistic Inquiry*, 33(1):63–96.

Frana, I. (2007). The *de re* analysis of concealed questions. In *Semantics and Linguistic Theory XVI*.

Francis, W. and Kučera, N. (1982). *Frequency analysis of English usage*. Houghton-Mifflin, Boston.

Friederici, A., Rüschemeyer, S., Hahne, A., and Fiebach, C. (2003). The role of the left inferior frontal and superior temporal cortex in sentence comprehension: localizing syntactic and semantic processes. *Cerebral Cortex*, 13:170 – 177.

Friederici, A. D.and von Cramon, D. Y. and Kotz, S. A. (1999). Language related brain potentials in patients with cortical and subcortical left hemisphere lesions. *Brain*, (122):1033–1047.

Frisson, S. and Pickering, M. J. (1999). The processing of metonymy: evidence from eye movements. *Journal of Experimental Psychology*, 25(6):1366–1383.

Frisson, S. and Pickering, M. J. (2007). The processing of familiar and novel sense of a word: Why reading Dickens is easy but reading Needham can be hard. *Language and Cognitive Processes*. in press.

Gamut, L. (1991). *Logic, Language, and Meaning*, volume 2. University of Chicago Press.

Ginzburg, J. (1996). Interrogatives: Questions, facts, and dialogues. In Lappin, S., editor, *The Handbook of Contemporary Semantic Theory*, chapter 15, pages 361–384. Blackwell Publishers, Oxford, UK.

Greenberg, B. (1977). A semantic account of relative clauses with embedded question interpretations. ms, UCLA.

Grimshaw, J. (1979). Complement selection and the lexicon. *Linguistic Inquiry*, 10:279 – 326.

Groenendijk, J. and Stokhof, M. (1984). *Studies on the Semantics of Questions and the Pragmatics of Answers*. PhD thesis, Universiteit van Amsterdam, The Netherlands.

Groenendijk, J. and Stokhof, M. (1989). Type-shifting rules and the semantics of interrogatives. In Portner and Partee (2002), chapter 17, pages 421 – 456. Originally published in *Properties, Types and Meanings, Vol 2: Semantics Issues*, Chierchia *et al* (eds), pp. 21 – 68, 1989, Kluwer Academic Publishers.

Groenendijk, J. and Stokhof, M. (1990a). Dynamic Montague grammar. In *Proceedings of the Second Symposion on Logic and Language*, pages 3–48, Budapest. Eotvos Lorand University Press.

Groenendijk, J. and Stokhof, M. (1990b). Partitioning logical space. Annotated handout.

Groenendijk, J. and Stokhof, M. (1991). Dynamic predicate logic. *Linguistics and Philosophy*, 14(1):39–100.

Groenendijk, J. and Stokhof, M. (1997). Questions. In van Bentham and ter Meulen (1997), chapter 19, pages 1055 – 1124.

Hagoort, P., Brown, C. M., and Swaab, T. Y. (1996). Lexical-semantic event related potential effects in patients with left hemisphere lesions and aphasia and patients with right hemisphere lesions without aphasia. *Brain*, 119:627–649.

Halgren, E., Dhond, R., Christensen, N., Van Petten, C., Marinkovic, K., Lewine, J. D., and Dale, A. M. (2002). N400-like MEG responses modulated by semantic context, word frequency, and lexical class in sentences. In *Neuroimage*, volume 17, pages 1101–1116.

Hämäläinen, M., Hari, R., Ilmoniemi, R. J., Knuutila, J., and Lounasmaa, O. V. (1993). Magnetoencephalography – theory, instrumentation, and applications to noninvasive studies of the working human brain. *Review of Modern Physics*, 65(2):413 – 497.

Hamblin, C. (1958). Questions. *Australasian Journal of Philosophy*, 36:159–168.

Hamblin, C. (1973). Questions in Montague English. *Foundations of Language*, 10:41 – 53.

Harris, J., Pylkkänen, L., McElree, B., and Frisson, S. (2007). The cost of question concealment: Eye-tracking and MEG evidence. In press. Brain and Language.

Heim, I. (1979). Concealed questions. In Bauerle, R., Egli, U., and von Stechow, A., editors, *Semantics from different points of view*, pages 51 – 60. Springer-Verlag, Berlin.

Heim, I. (1983). File change semantics and the familiarity theory of definiteness. In Portner and Partee (2002), pages 223–248. Originally published in: Bäurle, R., Schwarze, Ch., and von Stechow, A. (eds) *Meaning, use and interpretation of language*, Berlin, De Gruyter.

Heim, I. (1997). Predicates of formulas? Evidence from ellipsis. In *Proceedings of the 7th Conference on Semantics and Linguistic Theory*, number 7, Cornell University: Cornell Working Papers in Linguistics.

Heim, I. and Kratzer, A. (1998). *Semantics in Generative Grammar*. Blackwell Publishers, Oxford, UK.

Helenius, P., Salmelin, R., Service, E., Connolly, J., Leinonen, S., and Lyytinen, H. (2002). Cortical activation during spoken-word segmentation in nonreading-impaired and dyslexic adults. *Journal of Neuroscience*, 22:2936–2944.

Helenius, P., Salmelin, R., Service, E., and Connolly, J. F. (1998). Distinct time courses of word and context comprehension in the left temporal cortex. *Brain*, 121:1133–1142.

Henderson, B. (2007). Matching and raising unified. *Lingua*, pages 202–220.

Higgins, R. (1973). *The pseudo-cleft construction in English*. Garland, New York.

Horwitz, B. and Poeppel, D. (2002). How can EEG/MEG and fMRI/PET data be combined? *Human brain mapping*, 17:1 – 3.

Huang, J. (2001). *Anaphora*. Cambridge University Press, Cambridge, UK.

Hulsey, S. and Sauerland, U. (2006). Sorting out relative clauses. *Natural Language Semantics*, (14):111–137.

Inhoff, A., Eiter, B. M., and Radach, R. (2005). Time course of linguistic information extraction from consecutive words during eye fixations in reading. *Journal of Experimental Psychology: Human Perception and Performance*, 81:979–995.

Jackendoff, R. (1997). *The Architecture of the Language Faculty*. Cambridge, MA. MIT Press.

Janssen, T. M. V. (1984). Individual concepts are useful. In Landman, F. and Veltman, F., editors, *Varieties of Formal Semantics*, Proceedings of hte Fourth Amsterdam Colloquium, pages 171–192. Foris Publications.

Janssen, T. M. V. (1997). Compositionality. In van Bentham and ter Meulen (1997), chapter 7, pages 417–474.

Just, M., Carpenter, P., Keller, T., Eddy, W., and Thulborn, K. (1996). Brain activation modulated by sentence comprehension. *Science*, 274:114–116.

Kamp, H. (1981). A theory of truth and semantic representation. In Portner and Partee (2002), pages 189–222. Originally published in *Formal methods in the study of langauge*, Groenendijk, J., Janssen, T. and Stokhof, M. (eds.), Amsterdam: Mathematical Centre, pp. 277–322.

Kaplan, D. (1969). Quantifying in. In Davidson, D. and Hintikka, J., editors, *Words and Objections: Essays on the Work of W.V. Quine*, pages 221–243. Reidel, Dordrecht.

Karttunen, L. (1977). Syntax and semantics of questions. In Portner and Partee (2002), chapter 16, pages 382 – 421. Originally published in *Linguistics and Philosophy*, **1**, 1977, pp. 45 – 77, Kluwer Academic Publishers.

Kotz, S., Cappa, S., von Cramon, D., and Friederici, A. (2002). Modulation of the lexical-semantic network by auditory semantic priming: an event-related functional MRI study. In *Neuroimage*, volume 17, pages 1761–1771.

Kotz, S. A. and Friederici, A. D. (2003). Electrophysiology of normal and pathological language processing. *Journal of Neurolinguistics*, 16:43–58.

Kotz, S. A., Friederici, A. D., and von Cramon, D. Y. (1999). Auditory word list priming in left and right temporal lobe lesion patients. *Brain and Language*, 69:294–296.

Kripke, S. (1972). *Naming and Necessity*. Harvard University Press, Cambridge, MA.

Kripke, S. (1979). A puzzle about belief. In Ludlow (1997), pages 383–414.

Kuperberg, G., Holcomb, P., Sitnikova, T., Greve, D., Dale, A., and Caplan, D. (2003). Distinct patterns of neural modulation during the processing of conceptual and syntactic anomalies. *Journal of Cognitive Neuroscience*, 15:272–293.

Kutas, M. and Federmeier, K. D. (2000). Electrophysiology reveals semantic memory use in language comprehension. *Trends in Cognitive Science*, 4:463–469.

Kutas, M. and Hillyard, S. (1980). Reading senseless sentences: brain potentials reflect semantic incongruity. *Science*, 207:203–205.

Kutas, M. and Hillyard, S. A. (1984). Brain potentials during reading reflect word expectancy and semantic association. *Nature*, 307:161–163.

Kutas, M., Hillyard, S. A., and Gazzaniga, M. (1988). Processing of semantic anomaly by right and left hemispheres of commissurotomy patients. *Brain*, (111):553–576.

Kutas, M. and Schmitt, B. M. (2003). Language in microvolts. In Banich, M. and Mack, M., editors, *Mind, Brain, and Language: Multidisciplinary Perspectives*, chapter 8, pages 171–210. Lawrence Erlbaum Associates, Mahwah, NJ.

Larson, R. and May, R. (1990). Antecedent contained deletion or vacuous movement: A reply to Baltin. *Linguistic Inquiry*, 21:103–122.

Lebeaux, D. (1988). *Language Acquisition and the Form of Grammar*. PhD thesis, University of Massachusetts, Amherst, Amherst, MA.

Lebeaux, D. (1990). Relative clauses, licensing, and the nature of the derivation. In Carter et al., editor, *North Eastern Linguistics Society*, volume 20, pages 318–332, Amherst. University of Massachusetts, Amherst GLSA.

Ludlow, P., editor (1997). *Readings in the Philosophy of Language*. MIT Press, Cambridge, MA.

Mäkelä, J., Hämäläinen, M., Hari, R., and McEvoy, L. (1994). Whole-head mapping of middle-latency auditory evoked magnetic fields. *Electroencephalography and Clinical Neurophysiology*, 94:414–421.

Matin, E. (1974). Saccadic surpression: A review and an analysis. *Psyhological Bulletin*, 81(12):899–917.

McCarthy, J. (2002). *A Thematic Guide to Optimality Theory*. Cambridge University Press.

McElree, B., Frisson, S., and Pickering, M. J. (2006). Deferred interpretations: Why starting Dickens is taxing but reading Dickens isn't. *Cognitive Science*, 30:113–124.

McElree, B., Traxler, M. J., Pickering, M. J., Seely, R. E., and Jackendoff, R. (2001). Reading time evidence for enriched semantic composition. *Cognition*, (78):B15 – B25.

Merchant, J. (2000). Economy, the copy theory, and antecedent-contained deletion. *Linguistic Inquiry*, 31(3):566–575.

Merchant, J. (2001). *The Syntax of Silence*. Oxford University Press, Oxford, UK.

Moltmann, F. (2003). Propositional attitudes without propositions. *Synthese*, 135:70–118.

Montague, R. (1970). Universal grammar. *Theoria*, 36:373–398.

Muskens, R. (1995). *Meaning and Partiality*. CSLI/FoLLI, Stanford, CA.

Nathan, L. (2006). *The Interpretation of Concealed Questions*. PhD thesis, MIT, Cambridge, MA.

Noppeney, U. and Price, C. (2002). A PET study of stimulus and task induced semantic processing. *Neuroimage*, 15:927–935.

Nunberg, G. (1995). Transfers of meaning. *Journal of Semantics*, 12:109–132.

Papanicolaou, A. C. (1998). *Fundamentals of Functional Brain Imaging : A Guide to the Methods and Their Applications to Psychology and Behavioral Neuroscience*. Psychology Press.

Partee, B. (1986). Noun phrase interpretation and type-shifting principles. In Portner and Partee (2002), chapter 15, pages 357–381. Originally published in *Studies of Discourse Representation Theory and the Theory of Generalized Quantifiers*, Groenendijk *et al* (eds.) pp. 115 – 143, 1986, Foris Publications.

Partee, B. (1995). Lexical semantics and compositionality. In Gleitman, L. R. and Liberman, M., editors, *An Invitation to Cognitive Science*, volume 1, pages 311–360. MIT Press.

Partee, B. and Rooth, M. (1983). Generalized conjunction and type ambiguity. In Portner and Partee (2002), chapter 14, pages 334 – 356. Originally published in *Meaning, Use and the Interpretation of Language*, Bäuerle *et al*. pp. 361 – 393, 1983, Walter de Gruyter & Co.

Piñango, M. M. and Zurif, E. (2001). Semantic operations in aphasic comprehension: Implications for the cortical organization of language. *Brain and Language*, 79:297 – 308.

Piñango, M. M., Zurif, E., and Jackendoff, R. (1999). Real-time processing implications of enriched composition at the syntax-semantics interface. *Journal of Psycholinguistic Research*, (28):395 – 414.

Pickering, M., McElree, B., Frisson, S., Chin, L., and Traxler, M. (2006). Aspectual coercion and underspecification. *Discourse Processes*, 42:131–155.

Pickering, M. J., McElree, B., and Traxler, M. (2005). The difficulty of coerion: A response to de Almedia. *Brain and Language*, 93:1–9.

Poeppel, D. and Hickok, G. (2004). Towards a new functional anatomy of language. *Cognition*, 92:1 – 12.

Portner, P. and Partee, B., editors (2002). *Formal Semantics: The Essential Readings*. Blackwell Publishers.

Prince, A. and Smolensky, P. (1993). Optimality theory: Constraint interaction in generative grammar. Ms. Rutgers University, New Brunswick, N.J. and University of Colorado, Boulder.

Pustejovsky, J. (1995). *The Generative Lexicon*. MIT Press, Cambridge, MA.

Pustejovsky, J. and Bouillon, P. (1995). Aspectual coercion and logical metonymy. *Journal of Semantics*, 12:133–162.

Pylkkänen, L., Feintuch, S., Hopkins, E., and Marantz, A. (2004). Neural correlates of the effects of morphological family frequency and family size: an MEG study. *Cognition*, 91:B35–B45.

Pylkkänen, L., Llinàs, R., and Murphy, G. (2006). Representation of polysemy: MEG evidence. *Journal of Cognitive Neuroscience*, (18):1–13.

Pylkkänen, L. and Marantz, A. (2003). Tracking the time course of word recognition with MEG. *Trends in Cognitive Sciences*, (7):187 – 189.

Pylkkänen, L. and McElree, B. (2006). The syntax-semantics interface: On-line composition of sentence meaning. In Traxler, M. and Gernsbacher, M., editors, *Handbook of Psycholinguistics*. Elsevier, 2nd edition. In press.

Pylkkänen, L. and McElree, B. (2007). MEG study of silent meaning. submitted to Journal of Cognitive Neuroscience.

Pylkkänen, L., Stringfellow, A., , and Marantz, A. (2002). Neuromagnetic evidence for the timing of lexical activation: An MEG component sensitive to phonotactic probability but not to neighborhood density. *Brain and Language*, (81):666 – 678.

Pylkkänen, L., Stringfellow, A., Flagg, E., and Marantz, A. (2001). A neural response sensitive to repetition and phonotactic probability: MEG investigations of lexical access. In *Proceedings of Biomag 2000, 12th International Conference on Biomagnetism*, pages 363–367, Espoo, Finland. Helsinki University of Technology.

Quine, W. V. (1956). Quantifiers and propositional attitudes. *Journal of Philosophy*, 53:101–111.

Rayner, K. (1998). Eye movements in reading and information processing: 20 years of research. *Psychological Bulletin*, 124(3):372–422.

Rayner, K. and Duffy, S. A. (1986). Lexical complexity and fixation times in reading: Effects of word frequency, verb complexity, and lexical ambiguity. *Memory and Cognition*, 14:191–201.

Recanati, F. (2000). *Oratio Obliqua, Oratio Recta*. MIT Press, Cambridge, MA.

Rescher, N. (2005). *Epistemic Logic*. University of Pittsburgh Press, Pittsburgh, PA.

Reuland, E. and Everaert, M. (2003). Deconstructing binding. In Baltin, M. and Collins, C., editors, *The Handbook of Contemporary Syntactic Theory*, chapter 20, pages 634–669. Blackwell Publishing, Oxford, UK.

Richardson, D. C. and Spivey, M. J. (2004a). Eye-tracking: Characteristics and methods. In Wnek, G. and Bowlin, G., editors, *Encyclopedia of Biomaterials and Biomedical Engineering*. Marcel Dekker, Inc.

Richardson, D. C. and Spivey, M. J. (2004b). Eye-tracking: Research areas and applications. In Wnek, G. and Bowlin, G., editors, *Encyclopedia of Biomaterials and Biomedical Engineering*. Marcel Dekker, Inc.

Riggs, L. A., Merton, P. A., and Morton, H. B. (1974). Surpression of visual phosphenene during saccadic eye movements. *Vision research*, 14:997–1010.

Romero, M. (2004). Intensional noun phrases with *know* and *be*. *Catalan Journal of Linguistics*, 3:147 – 178.

Romero, M. (2005). Concealed questions and specificational subjects. *Lingusitics and Philosophy*, 28:687–737.

Romero, M. (2007a). Connectivity in a unified analysis of specificational subjects and concealed questions. In Barker and Jacobson (2007), pages 264–305.

Romero, M. (2007b). On concealed questions. In *Semantics and Linguistic Theory XVI*.

Ross, J. R. (1969). Guess who? In Binnick, R., Davison, A., Green, G., and Morgan, J., editors, *Proceedings of CLS*, volume 5, pages 252–286, Chicago, IL. Chicago Linguistic Society.

Sag, I. (1976). *Deletion and Logical Form*. PhD thesis, MIT, Cambridge, MA.

Sauerland, U. (1998). *The Meaning of Chains*. PhD thesis, MIT, Cambridge, MA.

Schönfinkel, M. (1924). Über die Bausteine der mathematischen Logik. *Mathematische Annalen*, 92:305–316.

Soames, S. (2005). *Reference and Description*. Princeton University Press, Princeton, NJ.

Staub, A. and Rayner, K. (2007). Eye movements and on-line comprehension processes. In Gaskell, M. G., editor, *The Oxford Handbook of Psycholinguistics*. Oxford University Press, Oxford, England.

Stockall, L., Stringfellow, A., and A, M. (2004). The precise time course of lexical activation: MEG measurements of the effects of frequency, probability, and density in lexical decision. *Brain and Language*, (90):88–94.

Stowe, L., Haverkort, M., and Zwarts, F. (2005). Rethinking the neurological basis of language. *Lingua*, (115):997 – 1042.

Swaab, T., Brown, C., and Hagoort, P. (1997). Spoken sentence comprehension in aphasia: Event-related potential evidence for a lexical integration deficit. *Journal of Cognitive Neuroscience*, 9:39–66.

Takahashi, S. (2006). *Decompositionality and Identity*. PhD thesis, MIT, Cambridge, MA.

Tarkiainen, A., Helenius, P., Hansen, P.C and, C. P., and Salmelin, R. (1999). Dynamics of letter string perception in the human occipitotemporal cortex. *Brain*, (122):2119–2132.

Tichý, P. (1982). Foundations of partial type theory. *Reports on Mathematical Logic*, (14):59–72.

Todorova, M., Straub, K., Badecker, W., and Frank, R. (2000). Aspectual coercion and the online computation of sentential aspect. In *Proceedings of the Twenty-second Annual Conference of the Cognitive Science Society*, pages 3–8.

Traxler, M.and Pickering, M. and McElree, B. (2002). Coercion in sentence processing: evidence from eye-movements and self-paced reading. *Journal of Memory and Language*, 47:530–547.

Uttal, W. R. (2001). *The New Phrenology*. MIT Press, Cambridge, MA.

van Bentham, J. and ter Meulen, A., editors (1997). *Handbook of Logic and Language*. Elsevier Press, Amsterdam, The Netherlands.

van Petten, C. and Luka, B. (2006). Neural bases of semantic context effects in electromagnetic and hemodynamic studies. *Brain and Language*, 97:279–293.

Zimmerman, M. (2000). Pluractional quantifiers: The occasional-construction in English and German. In Jackson, B. and Matthews, T., editors, *Proceedings of the 10th Conference on Semantics and Lingusitic Theory*, volume 10, pages 290–306, Ithaca, NY: Cornell University.

A
Eyetracking Materials

For each item, the verb before the "/" symbol is the CQ verb, the verb following this symbol is the control verb.

1. The name of the animal was learned / written by the students.
2. The title of the book was disclosed / listed by the author.
3. The flag of the nation was determined / carried by the king.
4. The age of the infant was guessed / confirmed by the doctor.
5. The alphabet of the language was predicted / printed by the linguist.
6. The contract of the company was decided / signed by the president.
7. The size of the herd was estimated / regulated by the inspector.
8. The variable in the equation was explained / omitted by the physicist.
9. The length of the operation was queried / recorded by the patient.
10. The anthem of the team was learned / shouted by the player.
11. The script of the show was disclosed / destroyed by the writer.
12. The costume of the actor was determined / delivered by the stylist.
13. The alias of the criminal was guessed / reported by the detective.
14. The paint of the portrait was predicted / stirred by the apprentice.
15. The number of the restaurant was asked / seen by the student.
16. The label of the shirt was decided / removed by the designer.
17. The volume of the music was estimated / noted by the engineer.
18. The design of the set was explained / sketched by the artist.
19. The distance of the flight was queried / stated by the pilot.
20. The speech of the senator was learned / published by the aide.
21. The recording of the band was disclosed / preserved by the singer.

22. The price of the phone was determined / reduced by the salesman.
23. The proof of the theorem was guessed / erased by the mathematician.
24. The peak of the stock was predicted / plotted by the broker.
25. The glaze of the cake was decided / eaten by the cook.
26. The value of the diamond was estimated / appraised by the jeweler.
27. The blueprint of the house was explained / modified by the architect.
28. The speed of the rocket was queried / computed by the observer.

B
MEG Materials

For each item, the verb before the "/" separates the CQ verb, from the Control verb, from the Anomalous verb.

1. The name of the furry animal had been learned / written / climbed by the elementary students.
2. The title of the popular book had been disclosed / listed / poured by the popular author.
3. The flag of the young nation had been determined / carried / married by the despotic king.
4. The age of the sick infant had been guessed / confirmed / strangled by the family doctor.
5. The alphabet of the exotic language had been predicted / printed / dressed by the field linguist.
6. The contract of the bankrupt company had been decided / signed / parked by the sickly president.
7. The size of the wild herd had been estimated / regulated / delighted by the wildlife inspector.
8. The variable in the stochastic equation had been explained / omitted / accused by the nuclear physicist.
9. The length of the risky operation had been queried / recorded / bothered by the nervous patient.
10. The anthem of the local team should be learned / shouted / advised by the baseball player.
11. The script of the primetime show should be disclosed / destroyed / convinced by the television writer.
12. The costume of the famous actor should be determined / delivered / surprised by the fashion stylist.
13. The alias of the illusive criminal should be guessed / reported / interested by the private detective.
14. The paint of the infamous portrait should be predicted / stirred / wounded by the young apprentice.
15. The logo on the company shirt should be decided / removed / trained by the new designer.
16. The volume of the rock music should be estimated / noted / drunk by the sound engineer.

17. The design of the movie set should be explained / sketched / silenced by the movie artist.
18. The distance of the international flight should be queried / stated / ground by the tired pilot.
19. The speech of the senior senator must be learned / published / satisfied by the faithful aide.
20. The recording of the jazz band must be disclosed / preserved / persuaded by the star singer.
21. The price of the digital phone must be determined / reduced / pleased by the electronics salesman.
22. The proof of the classic theorem must be guessed / erased / docked by the brilliant mathematician.
23. The peak of the volatile stock must be predicted / plotted / spilled by the stock broker.
24. The glaze of the wedding cake must be decided / eaten / wired by the pastry cook.
25. The value of the genuine diamond must be estimated / appraised / smothered by the discount jeweler.
26. The blueprint of the modern house must be explained / modified / reminded by the successful architect.
27. The speed of the distant rocket must be queried / computed / astute by the astute observer.