

THE SEMANTICS OF COPY RAISING

A SEMANTIC ANALYSIS OF ENGLISH COPY RAISING
CONSTRUCTIONS

BY
DIANE DORAN, B.A. (Hons)

A THESIS
SUBMITTED TO THE DEPARTMENT OF LINGUISTICS & LANGUAGES
AND THE SCHOOL OF GRADUATE STUDIES
OF MCMASTER UNIVERSITY
IN PARTIAL FULFILMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
MASTER OF SCIENCE

© Copyright by Diane Doran, August 2015
All Rights Reserved

Master of Science (2015)
(Linguistics & Languages)

McMaster University
Hamilton, Ontario, Canada

TITLE: A Semantic Analysis of English Copy Raising
Constructions

AUTHOR: Diane Doran
B.A. (Hons), (University of Western Ontario)

SUPERVISOR: Dr. Ivona Kučerová

NUMBER OF PAGES: vi, 74

Lay Abstract

This thesis investigates the linguistic meaning associated with the “copy raising” sentence construction, e.g. “Your cat looks like she wants to go outside.” I argue that the interpretation of these sentences is dependent on establishing the individual whose perspective is conveyed in the sentence, which does not need to be the speaker. After examining the range of contexts in which various different copy raising constructions can be used, I propose an analysis of their core meaning that draws on the philosophical idea of possible worlds, and the psychological notion of stereotypicality. I also address the question of whether these constructions are related to the phenomenon of evidentiality, a property of certain languages which allows the speaker to linguistically mark the source of evidence for their claim.

Abstract

This thesis is an investigation of the structural and formal semantic properties of copy raising constructions in English, as well as their expletive counterparts. The first main claim is that contrary to what has been previously assumed, the perceiver of the event (i.e. the PGOAL in Asudeh & Toivonen’s 2012 terms) is an obligatory syntactic and semantic argument of the matrix verb. I argue that the identification of the PGOAL is not left to pragmatics, but rather that is represented as a silent pronoun in the structure: one that picks up a logophoric antecedent. The result of this is that the material in the embedded clause is semantically interpreted with respect to the PGOAL’s perspective. The second major claim of the thesis is that this perspective-sensitivity is most appropriately captured using a modal semantic framework (Kratzer, 1977, 1981 von Stechow & Heim, 2002). Specifically, I argue that each of the different copy raising verbs encodes a different accessibility relation between possible worlds or situations, while the PGOAL’s information state provides the relevant domain of worlds. Using these insights, I propose truth conditions for these constructions, which ultimately are sensitive to a kind of stereotypical ordering, and account for inter-speaker variability. Finally, I discuss the anomalous class of copy raising constructions with non-thematic subjects, and argue that overlapping discourse functions may have resulted in a shift away from modal semantics in these cases.

Acknowledgements

First and foremost, I would like to thank my supervisor Dr. Ivona Kučerová, for all of her guidance in the past two years. When I began this degree, I did not know what a modal was. And while this admission may not instill the greatest confidence in my readers, I mention it nonetheless, because it is a testament to just how far I have come under her supervision. Without her keen insight, vast knowledge, and demand for precision and clarity, this thesis would not be what it is. She devotes a great deal of her own time and energy to her students' development, and I'm extremely grateful to her. I also have to thank all of the exceptional people in the Department of Linguistics and Languages, whom it has been a pleasure to work alongside. I've had the benefit of being a part of the lovely group of ladies who make up the Syntax Lab: Jitka Bartošová, Cassandra Chapman, Heather Stephens, Rachael Hardy, Malaree Baraniuk, and Olena Kit, who I thank for all their feedback and support.

My life would have been a whole lot more difficult this year if the Social Sciences and Humanities Research Council of Canada (SSHRC) had not stepped in, so they definitely deserve a thank-you. I am also lucky to have received plenty of moral, as well as financial support: thanks to Mary Lynn Tobiasz for frequent and much needed commiseration, Amanda Ho for being top-notch office company, Richard Mah for helping me get to conferences one way or another, and Daniel Schmidtke for always looking after my sanity. Thanks as well to the Social Science and Humanities Research and Writing Group at McMaster - the final months of thesis-writing would have been a lot more miserable without it.

And last but not least, I'm forever grateful to my marvelous mom, who, technically speaking, is to thank for all of this.

List of symbols and abbreviations

$\exists$	Existential quantifier (possibility)
λ	Lambda abstractor
$\forall$	Universal quantifier (necessity)
A&T	Asudeh and Toivonen
CRC	Copy raising construction
LOG	Logophor
NOC	Non-obligatory control
PC	Perceptual construction
PCC	Perceptual Characterization Condition
PCG	P-source-Copy Generalization
PGOAL	Perceptual Goal
POPT	Predicate of personal taste
PP	Prepositional phrase
PSOURCE	Perceptual Source
s*	Topic situation

Contents

Lay Abstract	iii
Abstract	iv
Acknowledgements	v
List of symbols and abbreviations	vi
1 Introduction to copy raising	1
1.1 The syntax of copy raising	2
1.1.1 The matrix subject	3
1.1.2 The verb of perception	4
1.1.3 The complement clause	6
1.1.4 The “copy”	6
1.2 The P-SOURCE interpretation	7
1.2.1 Asudeh and Toivonen 2012	8
1.2.2 Landau 2011	9
1.2.3 Kim 2014	12
1.3 Summary	13
2 The perceiver role	15
2.1 Pragmatic accounts	16
2.1.1 Entailment	16
2.1.2 Presupposition	17
2.1.3 Conventional implicature	18

2.2	The judge parameter	19
2.2.1	Background	19
2.2.2	Characteristics of POPTs and PCs	21
2.3	Summary	24
3	P-GOAL's place in the grammar	25
3.1	P-GOAL as a silent element	25
3.1.1	Syntactic properties of P-GOAL	26
3.1.2	P-GOAL as a logophoric pronoun	32
3.2	Situational restriction	36
3.2.1	World variables	37
3.3	Summary	40
4	A modal semantic account	41
4.1	False truth scenarios	41
4.2	Proposal	43
4.2.1	Accessibility relations	46
4.2.2	Domain of quantification	48
4.2.3	P-SOURCE restriction	51
4.3	Summary	55
5	Modality, evidentiality, and speaker certainty	56
5.1	Are PCs evidential?	56
5.1.1	P-SOURCE as evidential	57
5.1.2	The verb as evidential	58
5.2	Tests for epistemic modality	59
5.2.1	Challengeability	59
5.2.2	Infelicity if the proposition is known to be false	61
5.3	Speaker certainty and semantic shift	63
5.4	Summary	65
6	Conclusions and remaining questions	67

Chapter 1

Introduction to copy raising

The English construction known as *copy raising* (Rogers, 1971, 1972, Potsdam & Runner, 2001, Asudeh, 2002, Landau, 2011) is illustrated in (1).

- (1) Your cat looks like she is hungry.

Copy raising constructions are characterized by a matrix subject and verb of perception, followed by a *like* or *as though* complement containing a pronoun which is coindexed with the matrix subject. In (1), the embedded pronoun *she* refers to the matrix subject *your cat*.

The same set of perception verbs that feature in copy raising constructions (henceforth CRCs) can also take an expletive *it* as a matrix subject, while expressing the same content in the embedded clause. This is what I will refer to as an expletive variant, as in (2). These two related constructions will be the focus of this thesis.

- (2) It looks like your cat is hungry.

Syntactically, CRCs resemble Subject-to-Subject raising (*Your cat seems to be hungry*) and small clauses (*Your cat looks hungry*), while semantically they also share properties with predicates of attitude (*I think your cat is hungry*) and epistemic modals (*Your cat must/might be hungry*). Like S-to-S raising, the property attributed to the matrix subject can be expressed with a full clause. Like small clauses, the verb specifies the mode of perception used in making the claim. Like predicates of attitude and epistemic modals, an individual's impressions of reality are conveyed in the utterance. The structural and truth-conditional properties of these similar constructions have been well-investigated, while CRCs have received far less attention, and their exact nature remains unclear.

One main issue is that while predicates of attitude specify an overt attitude-holder, and epistemic modals represent the knowledge-base of the speaker, the person whose impressions are expressed is often left implicit in a CRC or expletive variant. Is it necessarily the speaker who expresses their own views, or can they express the views of others, or do they use these constructions to simply make some generic statement? And what is the nature of the views that they express? If CRCs are assertions, as they certainly seem to be, then what exactly does the speaker assert, and how do we judge whether the assertion is true? In other words, *what precise meaning do these constructions encode, and how is this accomplished?* This is a question that to my knowledge has not yet been sufficiently answered.

Using the issue of perspective as my point of entry, in this thesis I will explore the syntactic and semantic properties of CRCs in order to arrive at a formal answer to the above question. After investigating the role of the individual whose perspective is conveyed in the sentence, I will conclude that the speaker is not obligatorily tied to this role in any systematic way. Instead I will propose that CRCs incorporate a silent element referring to the perceiver; one which picks up a contextually-determined linguistic antecedent. I will then argue that the structural presence of the perceiver element is what provides the restriction on the worlds of evaluation for the embedded content; specifically, I claim that aspects of the perceiver's mental state determine the worlds that make up a *modal base*, using a modal framework in the style of Kratzer (1977, 1981) and von Stechow & Heim (2002). I show that the various copy raising verbs can be thought of as representing different accessibility functions, and that the set of worlds returned by these functions are those that are ultimately quantified over to yield a truth value.

I further propose that the final step in determining the set of evaluation worlds involves a stereotypical or normative ordering, where the embedded proposition is essentially ranked with respect to likelihood, given the accessible facts. This ordering explains the variable and vague nature of truth judgements. Finally, I suggest that the semantic notion of likelihood may be responsible for mediating a shift away from a modal interpretation to something more like an evidential interpretation. This is reflected in the observable differences between two alleged types of CRCs (thematic and non-thematic subjects (Landau, 2011)), and relates to the apparent "direct perception" requirement on the matrix subject that has been the focus of much of the work on CRCs to date.

1.1 The syntax of copy raising

I begin with an overview of the current state of knowledge surrounding the syntactic properties of CRCs, to provide a basis for the arguments that are to follow. While a complete consensus on the correct analysis of copy raising has not yet been reached,

I will review the prominent theories from recent years and state which assumptions I will be adopting, where applicable.

1.1.1 The matrix subject

Copy raising is so named due to the fact that early analyses (Rogers, 1971, 1972) assumed that its derivation was similar to canonical raising, in that the matrix subject was believed to have originated within the complement clause. For movement analyses like Rogers', it was argued that when the embedded subject moves to the matrix clause, it leaves behind a pronominal copy in its place, which results in the surface word order that is characteristic of CRCs.

However more recent investigations of these constructions have posed a number of problems for this kind of treatment. The first, raised in Potsdam & Runner 2001, is that movement analyses involve raising the embedded subject of a finite clause, which is in a Case position. This is problematic because movement from a Case position violates the Tensed S Condition (Chomsky, 1982), which states that A-movement from a tensed clause is prohibited, or more currently, the Phase Impenetrability Condition (Chomsky, 2001). It also remains to be seen how exactly movement would result in an overt pronoun being generated¹. Furthermore, there are many examples in which no copy at all is present in the embedded clause. Kim (2014) provides the following example, taken from the Corpus of Contemporary American English:

- (3) [The flight attendants] just seemed like it was any other day.
(Kim 2014: ex. (81b.))

For these reasons, Potsdam & Runner (2001) offer an alternative analysis, in which the matrix subject does not move from the embedded clause, but is instead base-generated in its matrix position. Support for a non-movement analysis comes from the differences in scope possibilities seen with canonical raising versus CRCs.

- (4) a. Two people seem to have won the lottery.
(two>seem, seem>two)
b. Two people seem like they have won the lottery.
(two>seem, *seem>two)
(Adapted from Potsdam and Runner 2001: ex. (34, 35))

The canonical raising sentence in (4-a) has two possible interpretations: one in which the quantifier scopes over the verb *seem*, and one in which it has the inverse

¹Unlike resumptive pronouns, which can sometimes occur in the context of illicit movement, these pronouns can be bound by a quantificational antecedent (Heycock, 1994, Potsdam & Runner, 2001).

reading, with *seem* scoping above the quantifier. The CRC in (4-b) however, only has the first reading. Canonical raising uncontroversially involves movement, and so the two readings can be explained if the quantifier is interpretable both at its surface position and its trace position.² If CRCs were also the result of movement, we would expect the same scope possibilities to arise - however the quantifier is only interpretable in its matrix position, strongly suggesting that it was not moved.

With the exception of Asudeh & Toivonen (2012), whose analysis will be discussed in more depth in subsequent sections, most recent work on CRCs follows Potsdam and Runner's (2001) reasoning, and assumes that the matrix subject is generated in situ (Landau, 2011, Kim, 2014). I will also adopt this assumption. A final point of note is that the subject is always either definite or generic; no true indefinite subjects are found with CRCs (Kim, 2014).

1.1.2 The verb of perception

There are only a small number of verbs which can appear in English CRCs. These are verbs which all share a core meaning of some kind of perception: *seem*, *appear*, *look*, *sound*, *smell*, *taste*, and *feel*. Of these verbs, only *seem*, *appear*, and sometimes *look* participate in canonical raising. Asudeh (2002) and Asudeh & Toivonen (2012) claim that there is an important distinction between *seem* and *appear*, which they call *copy raising verbs*, and all the rest, which they call *perceptual resemblance verbs*. They argue that true copy raising verbs always require a co-referential pronoun in the embedded clause, while perceptual resemblance verbs can be used without them.³ Asudeh (2002) provides the following examples:

- (5) a. *Richard seems like Gonzo has been baking.
 b. Richard smells like Gonzo has been baking.
 (Asudeh 2002: ex. (9.20, 9.21))

However this distinction has been challenged by other researchers, who present counterexamples from internet searches and corpus data. Landau (2011) argues that copy pronouns are not always necessary with copy raising verbs, citing a number of examples from the internet:

- (6) Maybe she didn't think it would be that hard, but when I talked to her she

²While the phenomenon of reconstruction in A-chains is not straightforward, and Chomsky (1993, 1995) and Lasnik & Saito (1999) have argued that A-movement does not reconstruct, more recent work suggests that A-reconstruction is possible and even necessary in certain circumstances (von Stechow & Heim, 2002, von Stechow & Iatridou, 2003).

³However they do acknowledge that certain speakers (who they label Dialect D) allow copy raising verbs without a copy pronoun.

seemed like there would be no issues.

(Landau 2011: ex. (21b.))

- (7) In fact, even the sky appeared as though the clouds themselves had been stripped of life.

(Landau 2011: ex. (21e.))

He suggests that for a number of speakers, the presence of a copy pronoun with copy raising verbs is merely a preference, and that the factors that determine whether a copy is necessary are related to pragmatic, rather than syntactic constraints; i.e., whether the embedded event is plausibly inferable from the matrix perceptual event. For Landau, the only important differences between the different verbs has to do with how permissive/restrictive they are in the different types of sensory stimulus they are compatible with. *Seem*, for example, is much more permissive than *taste*, in that it is compatible with stimuli that are perceived through virtually any modality, as well as mental representations of them (*Her voice seems strained*, *The punishment seems harsh*, etc.). *Taste* on the other hand essentially only allows gustatory stimuli.

Kim (2014) brings further empirical evidence to the discussion, using a wealth of corpus data from Corpus of Contemporary American English (COCA) and Corpus of Historical American English (COHA), which indicate that copy raising verbs can indeed be used without a copy pronoun:

- (8) For me, studying Yiddish seemed as though I were traveling, instead, through the streets of a long-forgotten hometown.

(Kim 2014: ex. (35a.))

- (9) In spite of that, or just for that reason, she appeared as if everything were finally in its place.

(Kim 2014: ex. (35b.))

He also finds no difference between the two purported verb types in terms of where the copy pronoun, if present, is likely to appear. I too will suppose for the remainder of this paper that there is no reason to distinguish between copy raising verbs and perceptual resemblance verbs in this regard, though I will discuss other types of differences.⁴

⁴While Asudeh & Toivonen (2012) maintain that copy raising verbs *seem* and *appear* select a single propositional argument, independent of which constructions they appear in, Landau (2011), Kim (2014) and others posit lexical ambiguity for all of the verbs, in that their argument structures vary depending on the thematic status of the subject. I will revisit the issue of argument structure later in this chapter.

1.1.3 The complement clause

The next element in a CRC is the embedded complement,⁵ headed by *like* or *as though/if*. This complement is necessarily finite, and usually (but as we have seen, not always) contains a pronoun referring to the matrix subject. There is some debate about the label of this complement. Asudeh (2002), Asudeh & Toivonen (2012), Heycock (1994), Potsdam & Runner (2001) all argue for treating the complement as a predicative PP, headed by *like* or *as though*, which itself contains a finite CP. Kim (2014) argues that *like* and *as though* in these cases are complementizers themselves, and thus the copy raising verbs directly select for a CP. Landau (2011) focuses on the question of whether and when the complement is predicative or propositional, and so presumably he takes the PP/CP status of the complement to be dependent on this distinction. As I will mostly follow Landau's syntactic account in what is to come, I will also assume that the complement can be either a PP or CP.

1.1.4 The “copy”

As we have already seen, the view that the embedded coreferential pronoun in the complement clause is anything like a “copy” is quite difficult to argue and has generally been abandoned. Even Asudeh (2002), Asudeh & Toivonen (2012), whose analysis does involve the matrix subject having raised to its surface position, do not assume that it originated in the position of the pronoun, but rather that it raised from the subject of the open PP complement. They argue that the relationship between the matrix subject and embedded pronoun is then one of simple anaphoric binding, and that there is no restriction on where in the complement it can appear.

This is consistent with Kim's (2014) corpus data, in which he finds many examples of CRCs whose complements feature coreferential pronouns in non-subject positions.

- (10) The girl seemed as if **her** mom was dying.
(Kim 2014: ex. (4a.))
- (11) The bed appeared as if someone had recently been dragged from **it**.
(Kim 2014: ex. (4b.))
- (12) The cabinets looked as though someone had thrown the pots into **them** from across the room.
(Kim 2014: ex. (43b.))

Though Asudeh & Toivonen (2012) allow the copy to occur in any position, recall

⁵There is abundant evidence that it is indeed a complement, cf. Kim 2014

that they state that copy raising verbs (i.e., *seem* and *appear*) always require one in their complements. The examples from Kim (2014) in section 1.1.2 illustrate that this is too strong, at least for some speakers. Landau (2011) and Kim (2014) have proposed an alternative: that the presence/location of the copy is related to whether or not the subject is thematic. I will discuss their claims in the next section, and ultimately side with Landau (2011).

1.2 The P-SOURCE interpretation

Having reviewed the claims that have been made regarding the correct syntactic analysis of CRCs, we now turn to an important problem that has been the focus of much of the recent literature on copy raising. Existing knowledge of the properties of CRCs does not predict that there should be any difference in the semantic interpretations of CRCs and their expletive variants, and yet in certain cases, the two are found to be inequivalent. Asudeh & Toivonen (2012) introduce what they call the Puzzle of the Absent Cook, which is reproduced in (13) and (14).

- (13) Context: A and B walk into Tom's kitchen. Tom is at the stove doing something, but exactly what is unclear.
- a. Tom seems to be cooking.
 - b. It seems like Tom is cooking.
 - c. Tom seems like he is cooking.
- (Asudeh and Toivonen 2012: ex. (23, 24, 25))

In the above scenario, all three sentence variants (canonical raising, expletive, and copy raising, respectively) are perfectly acceptable, and seemingly interchangeable. However the slightly different scenario in (14) yields different results:

- (14) Context: A and B walk into Tom's kitchen. There's no sign of Tom, but there are various things bubbling away on the stove and there are several ingredients on the counter, apparently waiting to be used.
- a. Tom seems to be cooking.
 - b. It seems like Tom is cooking.
 - c. #Tom seems like he is cooking.
- (Asudeh and Toivonen 2012: ex. (26, 24, 25))

Here we see that a CRC is not acceptable when the matrix subject cannot be directly perceived. The infelicity of the CRC in this scenario is unexpected, given that *seem* is typically analyzed as an unaccusative verb, meaning that it does not assign a

thematic role to its subject, and thus we should not expect there to be any semantic relation between the verb of perception and its subject.

Given the apparent requirement that the subject of a CRC must be directly perceived somehow, many authors have relied on the notion of a Perceptual Source (PSOURCE) (Asudeh & Toivonen, 2012, Landau, 2011, Rett et al., 2013, Camilleri et al., 2014, Kim, 2014), which was first hinted at in Rogers (1974). The core idea is that the subject of a CRC is somehow marked as the source of perception for the embedded claim. I will first discuss the original proposal of Asudeh and Toivonen (2012), who coined the term, and I will then examine two subsequent accounts which suggest that matters are somewhat more complicated.

1.2.1 Asudeh and Toivonen 2012

After providing a number of examples from both English and Swedish which show that the matrix subject of a CRC has to be interpreted as the origin of the impression expressed, Asudeh and Toivonen (henceforth A&T) attempt to provide an account of this phenomenon under an LFG framework. They argue that the matrix subject position is not a thematic one, and thus that the direct perception requirement could not be a component of meaning inherent to some theta role. As evidence that the matrix subject position is non-thematic, they present data from tests of thematicity, borrowed from Potsdam & Runner (2001), which include raising of expletives and idiom chunks:

- (15) %There seem like **there** are problems.
(Potsdam & Runner 2001: ex. (4b.))

- (16) %The cat seems like **it's** out of the bag.
(Asudeh & Toivonen: ex. (63a.))⁶

Certain speakers accept sentences like those in (15) and (16), where we see an expletive *there* and idiom chunk, respectively, in subject position. As Postal & Pullum (1988) convincingly argue, these tests demonstrate that the position in question is not assigned a thematic role by the verb. A&T thus assume that the matrix subject of a copy raising verb cannot be thematic.⁷ They instead propose that there is need for a more generalized kind of relation which they call a *semantic role*. Semantic roles are essentially a broader class of thematic roles, but include roles that do not

⁶The % symbol is used to indicate inter-speaker variation.

⁷Recall that for Asudeh and Toivonen, the only “copy raising” verbs are *seem* and *appear*. They acknowledge that the “perceptual resemblance verbs” (*look*, *sound*, etc.) may be optionally thematic, as argued by Potsdam & Runner (2001).

necessarily need to be tied to a semantic or syntactic argument. These are roles which represent inherent participants in a certain event, but which are often left unspecified. For example, most verbs do not require a location or temporal duration argument for saturation, and yet these aspects are logically and necessarily present. A&T propose two semantic roles which are inherent to CRCs and perceptual events in general: PSOURCE (Perceptual Source), the thing or person perceived, and PGOAL (Perceptual Goal), the perceiver.

They argue that the subject of a copy raising verb is interpreted as the PSOURCE, and the PGOAL may or may not be specified using a to-PP, e.g.:

(17) **Beth** seemed **to me** like she had won.

A&T argue that copy raising verbs include specifications in their lexical entries which require that there is a PSOURCE, and that the PSOURCE be its subject.⁸ The Puzzle of the Absent Cook can then be explained as a kind of presupposition failure occurring when the asserted PSOURCE in 17, Tom, does not match the actual PSOURCE: the state of the kitchen.

But if PSOURCE is not a thematic role, the next question A&T have to answer is how the matrix subject is semantically licensed. Recall from section 1.1.2 that A&T claim that copy raising verbs always require a copy pronoun in their embedded complements. They argue that copy raising verbs lexically contribute a “manager resource,” which effectively removes the copy pronoun from composition and allows the matrix subject to compose in its place, thus licensing it.

In the next section I will review the observations of Landau (2011), which suggest that this analysis may not adequately represent the actual language data.

1.2.2 Landau 2011

Landau explores a range of copy raising data and presents two new generalizations which undermine many of Asudeh and Toivonen’s claims: 1) copy raising verbs do not always require a copy in their embedded complement, and 2) the matrix subject is not always interpreted as the PSOURCE.

We saw a number of examples in Section 1.1 from both Landau (2011) and Kim (2014) where the verbs *seem* and *appear* are used without a copy of the matrix subject anywhere in the embedded clause (examples 3, 6-9). Landau argues that many of the copy-less CRCs which have previously been flagged as ungrammatical are in fact only seemingly infelicitous, due to an additional pragmatic constraint that the embedded event be plausibly inferable from the matrix perceptual event.

⁸Rett et al. (2013) even extend this idea to the category of evidentiality, and argue that copy raising subjects encode *direct evidence*.

- (18) a. *Tina seems like Chris has been baking sticky buns.
 (Asudeh & Toivonen 2012: ex. (16a.))
 b. *Bill appears as if Mary is intelligent.
 (Lappin 1984: ex. (16b.))

Landau argues that these examples are not truly ungrammatical, but rather lack sufficient contextual information to satisfy the plausible perceptual inference condition. However when the contextual cues are strong enough, these verbs can indeed be used without a copy:

- (19) [e.g., in a party context, where Mary is Bill's wife]
 Bill appeared as if Mary had just said the most embarrassing thing.
 (Landau 2011: ex. (36c.))

The observation that copy raising verbs do not always require copy pronouns poses a major problem for A&T's analysis, in that for them, "the licensing of the copy raising subject rests on its composition in the place of the copy pronoun," (2012, p. 358) as we saw above. And since they argue that the matrix subject is not thematic, they are left with no obvious way of explaining how it might be licensed.

Landau instead argues that the matrix subject *is* sometimes thematic, and that PSOURCE is in fact its accompanying *thematic* role. The arguments that Asudeh & Toivonen (2012) present in favour of the subject being non-thematic merely show that its status is ambiguous. Because A&T claim that the PSOURCE interpretation is not sufficient for semantic licensing (i.e. there must be a copy as well), they are forced assume that PSOURCE is not a theta-role. With numerous examples showing that copies are not always necessary, this step becomes unwarranted.

Furthermore, Landau is able to provide a number of examples in which the matrix subject of a CRC cannot be construed as a PSOURCE:

- (20) [Upon receiving the gradesheet]
 Oh dear, John looks like he has failed the exam.
 (Adapted from Landau 2011: ex. (27c.))
- (21) [Upon hearing a friend describe their car troubles over the phone]
 Your car sounds like it needs tuning very badly.
 (Adapted from Landau 2011: ex. (23), attributed to Heycock (1994))
- (22) The future seems like it's looming over my head.
 (Landau 2011: ex. (59b.))

In (20), John himself is not present to serve as the source of information - rather the information is gained by looking at the gradesheet. Similarly in (21), the conversation takes place over the phone, and thus the speaker has no direct perceptual contact with the car. The future (22) is of course inherently non-perceivable.

It may be tempting, in order to salvage the PSOURCE-subject generalization, to say that the subjects in Landau's examples have metonymic or ellipsis readings, such that John's name on the gradesheet in (20) represents an extension of John, that *your car* in (21) stands for "the description of your car," etc. However more conclusive examples can be constructed:

- (23) [You are a skilled musician with a highly trained ear. Through the thin walls of your apartment, you can hear your neighbour playing the guitar, but the chords sound slightly off, like the guitar is missing a particular string. You remark:]
The B string sounds like it's missing.

There is no conceivable way of construing the subject as a PSOURCE in this example, as it is not present even in a metonymic sense.

Importantly however, all of the sentences in which the subject is not a PSOURCE feature a copy pronoun in the embedded clause. Unlike cases where the subject can be interpreted as a PSOURCE, Landau argues that these sentences *do* require a copy pronoun in order to be grammatical. For instance, example (20) could not be rephrased as (24).

- (24) [Upon receiving the gradesheet]
#Oh dear, John looks like the exam was difficult.

However in a situation where John can be interpreted as the PSOURCE, for example upon seeing John's disheartened expression as he emerges from the exam room, the sentence in (24) is fine. Landau demonstrates this contrast with the minimal pair in (25).⁹

- (25) *Context:* I watch a news story about a famous bakery. They are now launching a new type of bread, which is presented in baskets along the counter. Dozens of people are crowded in there, savoring the smell of the loaves as they're waiting in line to buy the new bread. I say to my friend:
a. The bread/smell looks like it is irresistible.

⁹The data that Landau provides is in Hebrew, and he claims that the same interpretations arise in English as well, though he acknowledges that certain English speakers may find these degraded.

- b. The bread/*smell looks like nobody's on a diet anymore.
(Landau 2011: ex. (31))

When there is a copy in the embedded clause (25-a), the use of a non-visual subject like *the smell* is perfectly grammatical. However without a copy (25-b), *smell*, which is not a suitable PSOURCE here, cannot be used as the subject. These observations provide the basis for Landau's PCG, given in (26).

- (26) *The P-source-Copy Generalization (PCG)*
Given a sentence "DP_i V_{perc} (to DP_j) like CP",
A copy_i is necessary in CP iff DP_i is not a P-source.
(Landau 2011: ex. (26))

As Landau argues that PSOURCE is a thematic role, this can be rephrased as: a copy is only necessary in the embedded clause iff the matrix subject is not thematic. This is because Landau's analysis essentially works as follows: when the matrix subject is not thematic it cannot be licensed by the matrix verb, so it needs to be licensed by predication. A null operator merges with the embedded CP to yield a predicate, and this operator needs a variable to bind:¹⁰ the copy. The PCG is able to correctly predict the felicity of CRCs with remarkable consistency.

Contra A&T, Landau's work on CRCs strongly suggests that the subject of a CRC is not always interpreted as the PSOURCE - but that when it is, it is a thematic argument.

1.2.3 Kim 2014

Building on the work of Potsdam & Runner (2001), Landau (2011) and others, Kim (2014) also supports the view that the matrix subject of a CRC may be either thematic or non-thematic. To show that in some cases, the matrix subject must be in a theta-position, he repeats from data Potsdam & Runner (2001) (and overlooked by Asudeh & Toivonen (2012)), which shows that expletives and idiom chunks are not always permissible in subject position:

- (27) *There seems like John expects there to be an election.
(Kim 2014: ex. (14))
- (28) *The other foots appears like the shoe is on it.
(Kim 2014: ex. (15a.))

¹⁰This is in accordance with the Ban on Vacuous Quantification (Chomsky, 1982).

In these cases, Potsdam and Runner argue that the matrix subject position is thematic, and thus cannot tolerate expletives or idiom chunks. In fact, they propose that the subject is thematic when it does not bind a pronominal copy in the embedded clause's *subject* position, but non-thematic when it does. Kim also maintains this distinction, although he does not speculate as to what type of theta role the subject bears when it is thematic; unlike Landau (2011), he does not think of PSOURCE as a thematic role. Instead he develops the PSOURCE condition into a pragmatic constraint applying to all CRCs:

- (29) Perceptual Characterization Condition (PCC):
 The matrix subject of a CR construction, serving as the topic, is “perceptually characterized” by the rest of the utterance.
 (Kim 2014: ex. (42))

Kim does not provide a more precise definition of what it means to “perceptually characterize” something, other than to say the embedded clause must “say something characteristic” about the subject. But he argues that as long as the PCC is observed, there is no reason that a copy pronoun would be *required* in the embedded clause.¹¹ His work is the most recent contribution to the literature on CRCs.

1.3 Summary

Having examined the literature on copy raising constructions, it becomes clear that the English version could have been more aptly named. The “copy” is not actually a copy of the matrix subject: it can refer to other entities, appear in non-subject positions, or be absent altogether. Consequently the matrix subject is not a product of raising either: it is likely generated in situ.

We have looked at the puzzling properties of the matrix subject, both with respect to its interpretation as the perceptual source for the claim, as well as its optionally thematic status, which seems to be related to the presence or placement of the “copy” pronoun. While Landau (2011) and Kim (2014) disagree over precisely which environments require the matrix subject to be thematic, it is not immediately obvious how one would test the predictions of each account, and in fact it may be that they each are pursuing the analysis of a slightly different dialect. Yet what can be taken away from their work is that the matrix subject is almost certainly ambiguous: it can sometimes be non-thematic, but it can also be thematic. This fact casts doubt over the supposed necessity of Asudeh and Toivonen's (2012) *avante-garde* introduction

¹¹So long as the matrix subject is thematic, presumably. See Kim (2014) for his full account using an HPSG framework.

of “semantic roles” to the system. Depending on whether Landau or Kim’s analysis turns out to be correct, the PSOURCE interpretation is either inherent to the meaning of the theta-role PSOURCE, or it is the biproduct of a pragmatic constraint (the PCC). As Landau’s account both uses principled arguments that deal with established phenomenon, and gives explicit and reliable predictions for the acceptability of CRCs, I find it to be the more convincing and more testable of the two. I will adopt his central claims for the remainder of this thesis.

Having reviewed the evidence for promoting the PSOURCE subject to the status of a thematic argument, I now turn my attention to its partner “semantic role”, the PGOAL. In this thesis I will explore the proposed role of the PGOAL and its place in the linguistic system, ultimately arguing that it too has been wrongfully denied its place as an argument. I will show that its role in the grammar is in fact crucial for the interpretation of CRCs and their expletive variants, and that acknowledging this can lead us to a far better understanding of the many puzzling properties of these constructions. Chapter 2 examines the possibilities available to the role of the PGOAL, and proposes that it has representation in the syntactic structure. Chapter 3 investigates source of the PGOAL’s referent and argues that it can be analyzed as a logophoric pronoun, and Chapter 4 presents my account of these constructions using a modal semantics à la Kratzer (1977, 1981) and von Stechow & Heim (2002). In Chapter 5 I return to Landau’s non-thematic subject CRCs, and suggest that the relationship between epistemic modality and evidentiality can help to explain their properties. I conclude in Chapter 6 with a brief summary, and discussion of outstanding issues.

Chapter 2

The perceiver role

The individual that plays the role of the perceiver in the perceptual event described by a copy raising verb, i.e., the PGOAL in Asudeh and Toivonen's (2012) terms, has received much less attention in the copy raising literature than the role of the perceiver, or PSOURCE. As we saw in the previous chapter, Asudeh & Toivonen (2012) consider the PGOAL to be a *semantic* role, similar to the Experiencer thematic role, which is an entailed participant in an event that does not necessarily correspond to a semantic argument. They argue that because the PGOAL is optional in syntax, it cannot be an argument in the sense of the Theta Criterion.¹² However this conclusion does not necessarily follow; the same verb may have different underlying argument structures, as is the case with causative alternations (e.g. John opened the door vs. The door opened).¹³ Furthermore, many languages (the so-called *pro-drop* languages) allow the subject of a verb to be omitted in speech, but this does not mean that these subjects are not arguments, and we have no reason to exclude the possibility that something similar is true for the PGOAL.

Landau (2011) assumes that the Experiencer argument is somehow present at least at the semantic level, as the lexical entries that he provides for the copy raising verb *sound* contain a slot for an Experiencer variable, yet he does not elaborate on how this argument is represented or where it might get its value.

In this chapter, I aim to investigate the nature of the PGOAL more carefully. I will first examine the possibility that the identity of the PGOAL is determined through pragmatic means such as presupposition or implicature. I will ultimately reject this possibility, and in Section 2.2 I will instead explore the similarities between the role of the perceiver in CRCs, and the role of the judge in predicates of personal taste,

¹²Their other argument is that an overt PGOAL cannot be extracted from, while extraction from argument positions is usually permitted. However Experiencer arguments tend to be the exception to this rule crosslinguistically (Preminger, 2006).

¹³Thanks to Ivona Kučerová for mentioning this point.

which many people argue should be represented as a silent element present in the structure.

2.1 Pragmatic accounts

Asudeh and Toivonen's analysis of CRCs focuses mainly on the PSOURCE, and they acknowledge that an investigation of the PGOAL is still needed (2012, p. 364):

“The second issue for further research concerns the status of PGOAL. In the absence of other information (e.g., a to-PP in English), the speaker seems to fill the role of PGOAL. However, it is a legitimate question whether this is an entailment, with the PGOAL being equated in the semantics with the speaker index. Alternatively, the information that the speaker is the PGOAL could be a presupposition or conventional implicature, with the PGOAL being existentially closed in the semantics on a par with PSOURCE.”

In this section I will show that though these suggestions seem plausible at first glance, none of them are able to account for the range of interpretations that can arise in conversation. Because much of what I will discuss in this thesis applies to both CRCs *and* their expletive variants, I will use the term *perceptual construction*, henceforth PC, to refer to the superset containing both.

2.1.1 Entailment

If it were an entailment that the speaker is the PGOAL, we would expect, as Asudeh and Toivonen put it, that the speaker is equated in the semantics with the speaker index, and thus that the interpretation that comes with a PC is always that the speaker is expressing their own impressions, even when no overt PP is used. However it is fairly easy to imagine situations in which speakers may express the impressions of other individuals.

- (30) Jim: Why is the dog growling at the mirror?
Bob: It looks like another dog is staring at him.

In this example, Bob clearly is expressing what it must look like from the dog's perspective. Bob has the cognitive capacity to recognize a mirror image, but the dog does not. The dog must act as the PGOAL in this sentence, i.e. “It looks to the dog like another dog is staring at him.”

Many more examples are conceivable, however one is enough to rule out the possibility that the speaker is entailed to be the PGOAL. Entailments are not cancellable; an entailment relation is such that whenever s is true, s' is also true. Thus if this were an entailment, we should expect that whenever a perceptual event is expressed by a PC, the perceiver of that event is the speaker. The above example shows that this is false.

2.1.2 Presupposition

The defining characteristic of a linguistic presupposition is that it can be projected to or inherited by a higher clause, in environments where entailments cannot (Beaver, 2001). Beaver uses the examples in (31) to demonstrate the difference between an entailment and a presupposition.

- (31) a. Whoever discovered the elliptic form of the planetary orbits died in misery.
 b. Somebody died in misery.
 c. Somebody discovered the elliptic form of the planetary orbits.
 (Beaver 2001: ex. (E1, E3, E7))

The proposition in (31-b) is entailed by (31-a), but (31-c) is *presupposed* by (31-a). Presuppositions are known to survive when embedded under negation, in the antecedent of a conditional, and under modal possibility operators, among other environments. Consequently, the presupposition that somebody discovered the elliptic form of the planetary orbits survives in all of three environments:

- (32) Whoever discovered the elliptic form of the planetary orbits did not die in misery.
 (Beaver 2001: ex. (E4))
- (33) If whoever discovered the elliptic form of the planetary orbits died in misery, he should have kept his mouth shut.
 (Beaver 2001: ex. (E5))
- (34) Perhaps whoever discovered the elliptic form of the planetary orbits died in misery.
 (Beaver 2002: ex. (E6))

When these presupposition tests are applied to the hypothesis that the speaker is presupposed to be the PGOAL, we see contrasting results:

- (35) a. Josh looked like he was tired yesterday.
b. Presupposition?: I saw Josh yesterday.
- (36) Josh did not look like he was tired yesterday.
- (37) Maybe Josh looked like he was tired yesterday.
- (38) If Josh looked like he was tired yesterday, he should have hidden it from his boss.

The conditional and modal operator in (37) and (38) certainly do not allow the assumption that the speaker is the PGOAL to project into the main clause; the interpretation that arises is that the speaker cannot conclude what John looked like yesterday. It is less clear whether negation (36) changes the interpretation in this respect, though importantly it still allows for the possibility that the speaker learned about John's tired appearance from someone else. And in fact I will later argue that the interpretation where the PGOAL = speaker need not have been the initial interpretation at all. At the moment however, we may conclude that overall, the evidence seems to weigh against the idea that PGOAL is *presupposed* to be the speaker.

2.1.3 Conventional implicature

The last of Asudeh and Toivonen's suggestions is that the information that the speaker is the PGOAL could be a conventional implicature. A conventional implicature, as defined by Grice (1975), is an implication that is conveyed by the conventional meaning of a word or phrase, but does not affect the truth conditions of an utterance. Grice gives the following example:

- (39) He is an Englishman; he is, therefore, brave.
(Grice 1975: p. 44)

By using the word *therefore* the speaker implicates that it follows from the referent's being an Englishman that he is brave. But in order to satisfy the truth conditions, it must only be the case that he is an Englishman, and that he is brave. The additional meaning that there is some relation between being an Englishman and being brave is a conventional implicature. One property of conventional implicatures is that they cannot be overtly canceled. For example, the same speaker could not utter (40) without contradicting themselves:

(40) #He is an Englishman and therefore brave, but those two facts are unrelated.

If the information that the speaker is the PGOAL were a conventional implicature, it would not be cancellable. However we see in (41) and (42) that this information is easily cancellable without the same sense of contradiction:

(41) Josh looked like he was tired yesterday - though I wasn't the one who saw him.

(42) Tina sounded like she misses home - but I wasn't the one who talked to her.

The identity of the PGOAL in PCs cannot be explained away as a conventional implicature, nor an entailment, nor a presupposition. Gricean-type accounts of this phenomenon fail to capture the range of readings that are available to these constructions. Our theory needs some other way of accounting for how speakers interpret these types of sentences when no PGOAL is overtly specified; and I will later provide examples which indicate that the sentences that are utterable in certain situations could only be felicitous if we take the PGOAL to be represented grammatically. In the next section, I will discuss the parameter of the judge, which is argued to have representation in the syntax/semantics, and bears many important similarities to the PGOAL.

2.2 The judge parameter

Many of the properties of PCs that were observed in the previous section are reminiscent of the properties associated with predicates of personal taste (henceforth POPTs). To explain the truth conditions of POPTs, many authors argue that there is a contextual “judge” whose tastes are relevant for evaluating the proposition, and that this judge has a place in the semantic (and perhaps syntactic) representation of an utterance. I will ultimately propose something similar for PCs, and in this section I will expose the similarities between these two constructions to lend support to my proposal.

2.2.1 Background

A predicate of personal taste is a predicate that expresses properties which are inherently subjective. The classic examples from Lasnik (2005) are seen in (43) and (44).

(43) Roller coasters are fun.

(Lasersohn 2005: ex. (1))

- (44) This chili is tasty.
(Lasersohn 2005: ex. (2))

The fundamental problem posed by such predicates is in determining how to evaluate their truth. As Lasersohn observes, though there is really no “fact of the matter” - speakers can disagree using POPTs, in what he calls “faultless disagreements.” These are exchanges where speakers can contradict each other using statements that appear have the same content, and yet neither speaker can be said to be wrong. For example:

- (45) John: This is fun!
Mary: That’s not true! This isn’t fun at all!
(Lasersohn 2005: ex. (31))

Lasersohn (2005) was the first to suggest that the notion of a contextual judge be incorporated into the semantics of POPTs to deal with this problem. His innovation was to propose that in addition to worlds and times, POPTs are evaluated with respect to *judges*, and thus the evaluation index is equal to $\langle w, t, j \rangle$. This means that an utterance containing a POPT such as (44) can now be judged as true iff the chili in w tastes good to j at t , where the speaker typically acts as the judge.

Stephenson (2007), noting certain similarities between POPTs and epistemic modals, modifies Lasersohn’s analysis in order to arrive at an account that can successfully capture the properties of both. Unlike Lasersohn, who assumes that these predicates only take a single syntactic argument, Stephenson treats POPTs as two-place predicates, whose first argument is the individual whose tastes are relevant. She argues that there are three possible options for this argument. First, it may be filled overtly, as in sentences like (46).

- (46) This rollercoaster is fun for Sam.
(Stephenson 2007: ex. (14-a, 51-a))

But if not, the argument position is filled by a silent pronoun. Stephenson posits two different types of pronouns which may appear in this position. The first she labels PRO_J , which is not a pronoun in the traditional sense, but rather more like a world-denoting variable, in that it selects the world of evaluation, and takes its reference from an index. This is essentially a reformulation of Lasersohn’s (2005) judge-dependent semantics. Stephenson however, assumes that speakers always take an autocentric perspective when this judge pronoun is present, meaning that they

make assertions using themselves as the judge.¹⁴ A speaker uttering (47) would thus be asserting that the rollercoaster in w is fun for him at t .

- (47) This rollercoaster is fun PRO_J
(Adapted from Pearson 2007: ex. (49-a))

The last option is for the argument to be *pro*, a null referential pronoun referring to a contextually salient individual.

- (48) This rollercoaster is fun pro_{Sam}
(Adapted from Pearson 2007: ex. (50-a))

In this case the interpretation is exactly the same as in (46), and the sentence is judged as true iff the roller coaster in w is fun for Sam at t . Presumably, this would need to be uttered in a context where Sam's perspective is made salient, though Stephenson does not elaborate on what the contextual requirements might be.

The truth value of a POPT is thus dependent on the setting of this argument, which more often than not is represented by a silent element. In the next section I will show that the relationship between a PC and its PGOAL is much the same.

2.2.2 Characteristics of POPTs and PCs

Like POPTs, PCs most often appear without any overt prepositional phrase referring to the PGOAL:

- (49) a. It looks like it's going to rain.
b. Shawn seems like he's had a bad day.
c. Trudy sounds like she's plotting something.

And also like POPTs, the person whose perspective is relevant is typically understood to be the speaker. However as we saw section 2.1, this is not always the case. Example (30), repeated as (50) below, shows that when another individual is made contextually salient, this individual can be understood as the PGOAL.¹⁵

- (50) Jim: Why is the dog growling at the mirror?
Bob: It looks like another dog is staring at him.

¹⁴However she assumes that attitude operators shift the judge parameter to the attitude holder.

¹⁵In fact, I will later argue that the notion of contextual saliency is not quite enough in this case, and that the PGOAL actually requires a certain linguistic antecedent.

This is precisely what Stephenson (2007) argues to be the case for the silent pronoun of POPTs, as in example (48). Example (50) was modeled after a canonical example that she attributes to Kai von Fintel, seen in (51), where the individual whose taste is expressed is understood to be the cat:

- (51) (After watching the cat devour its food)
 “I think the new cat food is tasty.”

This is what Lasersohn (2005) would have referred to as an example of a speaker taking an “exocentric” perspective, and what Stephenson argues is an example of the first argument containing an silent referential pronoun referring to the cat.¹⁶

The puzzle of “faultless disagreements” posed by Lasersohn (2005) for POPTs, can also be seen with PCs. Consider:

- (52) Sam: The man on that bench looks like he’s just been dumped.
 Sue: Nuh-uh, he looks like he’s got bad stomach cramps.

Suppose that the man on the bench is a total stranger, and neither Sam nor Sue will ever get to know what sort of day he was having. For the purposes of this exchange then, there is no fact of the matter. Each speaker seems to be expressing her own impressions of the situation, much like two speakers expressing their own gustatory impressions of a particular batch of chili. It seems sensible to suppose that, like POPTs, the truth of a PC is relative to a certain individual. Sam’s experiences with the world up until this point will have led her to form a series of beliefs and associative links pertaining to observable things in the world and their potential underlying causes. Perhaps in her experience, people who have recently endured heartbreak typically exhibit whatever facial expression the man on the bench is currently making. But Sue’s experiences of the world may be different, and perhaps for her, the man’s expression is characteristic of abdominal pain.¹⁷ It is true that it looks to Sam like the man has just been dumped, and also true that it looks to Sue that he has stomach cramps; just as it is true that one speaker finds the chili tasty, and that another does not.

I follow Stephenson (2007) and assume that the reason that speakers are able to disagree over such individually-relativized statements is that the disagreement functions as a signal which indicates to the first speaker that his or her statement,

¹⁶For Stephenson’s arguments against Lasersohn’s “exocentric perspective” explanation, the reader is encouraged to consult her paper.

¹⁷In Chapter 4, I develop this idea into a modal semantic analysis in the style of Kratzer (1981), von Fintel & Heim (2002).

which Stephenson views as a proposal to update the common ground, has not been accepted. The shared property of faultless disagreements is further motivation for a similar analysis for POPTs and PCs.

One final similarity to be discussed has to do with the behaviour of these constructions when embedded under epistemic modals. Pearson (2013) observes that both bare POPTs and the verb *seem* apparently require that the speaker has had some relevant perceptual contact with the subject, but that this requirement seemingly disappears when they are used with an epistemic modal.

- (53) a. This cake is tasty.
 (Pearson 2013: ex. (1))
 b. John seemed like he was tired yesterday.¹⁸

Pearson claims that the sentences in (53) give rise to the assumption that the speakers have tasted the cake, and seen John, respectively. In contrast, these assumptions are ruled out with epistemic modals:

- (54) a. The cake must be tasty.
 (Pearson 2013: ex. (30))
 b. John must have seemed like he was tired yesterday.

Pearson (2013) in fact conjectures that both POPTs and the verb *seem* take implicit Experiencer arguments - which of course is the line of thinking that is being pursued here - and that there is “a presupposition of direct perceptual experience built into the meaning of the Experiencer theta role” (p. 118).

Pearson’s judge-free alternative analysis of POPTs appeals to a notion of first-person genericity, in which speakers project beyond their own experience, and claim that the proposition holds true for *all those individuals with whom they identify*. She implements this with a generic operator (GEN) that binds a covert Experiencer argument; and because Pearson’s version of genericity is first-person oriented, she takes this Experiencer to be the speaker. Thus the presupposition of direct perceptual experience would normally apply to the speaker, as the *identity-with* relation is reflexive. To account for problematic cases such as epistemic modals and cat food, she says that that speaker is included in the set of individuals that GEN ranges over *unless there is some reason to exclude her*. This forces Pearson to develop an elaborate account of how the contextual restriction of her generic operator is determined.

¹⁸Pearson’s example actually involves a small clause (John seemed tired yesterday), but the observation also holds true of CRCs. While an investigation of the properties of small clauses is beyond the scope of this thesis, the interested reader is referred to Basilico 2003, Moulton 2013.

While the idea of there being a presupposition of direct experience built into the Experiencer theta role seems reasonable, we have already seen in Section 2.1.2 that for PCs, it cannot be a presupposition that this Experiencer is necessarily the speaker. What seems more sensible, at least as far as PCs are concerned, is to assume that there is indeed a presupposition of direct experience associated with the PGOAL argument, but this argument can take many values other than the speaker. The example with the epistemic modal in (54-b) does indeed invoke the interpretation that someone has had perceptual contact with the subject, but it is someone other than the speaker: perhaps John's coworker has emailed the speaker to express concern for his sleeping habits, and the speaker infers that John seemed tired to this coworker yesterday. Rather than proposing, as Pearson does, that because epistemic *must* signals indirect evidence the speaker is consequently excluded from the contextual restriction of GEN, we can simply assume that the Experiencer/PGOAL in this sentence refers to another individual.¹⁹

While my intention is not to enter into the debate over the correct analysis of POPTs, the version of the analysis which features a silent pronoun is more easily transferable to an analysis of PCs, as it involves less machinery and seems to capture the interpretations more accurately.

2.3 Summary

In this chapter we have taken a closer look at Asudeh and Toivonen's (2012) proposed semantic role PGOAL: the perceiver of the event described by a PC. We have seen that straightforwardly accounting for the identification of its referent cannot be accomplished using pragmatic mechanisms such as entailment or implicature. Section 2.2 investigated the properties common to both PCs and POPTs, which have been analyzed (Stephenson, 2007, Keshet, 2005) as incorporating a silent element denoting the individual whose tastes are relevant. We saw that the PGOAL and the judge are analogous in many important respects, and I have suggested that a similar approach to the analysis of the PGOAL may be warranted. In the next chapter I will delve into the issue of how exactly the PGOAL is represented.

¹⁹Note that the idea of epistemic *must* signaling indirect evidence is perfectly compatible with this way of thinking as well. Matthewson (in press) has argued that *must* is an evidential signaling *not best Evidence Strength*, meaning it can only be used in cases where the speaker does not have the best possible evidence for their claim. First-hand sensory evidence is generally, and definitely in the cases of taste predicates and verbs of perception, the *best* possible evidence. Under this view, *must* would simply never occur in an instance of first-hand perception, and so in order to be felicitous, the PGOAL must have an available referent other than the speaker when embedded under an epistemic modal.

Chapter 3

P-GOAL's place in the grammar

Now that we have seen some preliminary evidence in favour of incorporating the PGOAL into the grammar, I turn to a more concrete investigation of the ways in which this element interacts with the syntactic structure, and the consequences it has on semantic interpretation. In Section 3.1 I will examine the possible syntactic relations between the PGOAL and other elements in the structure, and argue that it most closely resembles a logophoric pronoun (Clements, 1975, Reinhart & Reuland, 1993, Landau, 2001). Then in Section 3.2 I will illustrate that the PGOAL plays a crucial role in the semantics, in providing the restriction on the worlds of evaluation for the clause below it.

3.1 P-GOAL as a silent element

Before I argue that the PGOAL is present in the syntax as a silent element, I will first attempt to determine what type of silent element it might be. In light of the similarities between CRCs and POPTs, it might make sense for the PGOAL to resemble one of the silent elements that Stephenson (2007) posits to represent the judge: a null referential *pro*, which gets its value from the context (similar to the *pro* of discourse pro-drop languages like Chinese and Japanese), or a world-denoting $\text{PRO}_{J/G}$ which takes its value from an index. A priori, it is also conceivable that the silent PGOAL is related to the kind of silent PRO argument that is believed to be present in the complement of control predicates (Landau, 2001), and must be bound by a local controller. It could also be a version of the classic phonetically null *pro* which is deleted under conditions of recoverability (Chomsky, 1981, 1982), or it may even be something less specified, resembling the kind of null *phi*-Phrase or minimal null pronoun of Holmberg (2005) or Landau (2015) respectively, which lack certain features and therefore must acquire them from elsewhere.

There is a wide world of silent elements, which can often be difficult to tease apart from one another. In this section I will attempt to narrow down the possibilities by examining the properties of the silent element in question, before outlining my proposal.

3.1.1 Syntactic properties of P-GOAL

In at least some cases, the identity of the PGOAL is clearly supplied by a linguistic antecedent. If the PGOAL is the type of silent element that requires a linguistic antecedent, looking at the structural conditions on the relationship between antecedent and silent element will help to reveal its nature. I begin by examining these cases, before moving to cases where the role of an antecedent is less clear.

Non-local antecedents

A known characteristic of the silent PRO of Obligatory Control predicates is that PRO must be controlled locally (Landau, 2001):

- (55) Sarah tried PRO_S to make the bed.

The PRO_J of Stephenson (2007) is also bound locally, whether by the speaker or the closest attitude predicate, e.g.:

- (56) Sam thinks that Mary thinks the cake is tasty PRO_J

where the attitude operator *think* shifts the judge parameter to the subject of the nearest attitude verb, Mary.

The antecedent of PGOAL, however, need not be local, and in fact other potential antecedents may even intervene:

- (57) [Six-year old Jimmy's mother has arranged for his father to dress up in a Santa suit and pretend to arrive on Christmas Eve. But Jimmy is a terrible eavesdropper and has discovered the truth. To humor his parents, on Christmas Eve he pretends to watch in awe as Santa greets his mother.]
Jimmy thinks that his mom thinks it looks like she and Santa are old friends.

Although *his mom* is subject of the closest attitude predicate to the embedded clause, the interpretation is still undoubtedly that she thinks that it looks *to Jimmy* as though she and Santa are old friends. Given that (57) and other similar examples are acceptable, it does not seem plausible that the PGOAL is related to any of the existing versions of big PRO.

Quantificational antecedents

Importantly, the referent of PGOAL can have a distributive reading, as in (58), where the perspective represented in the complement (i.e. the PGOAL) covaries with each of the referents picked out by the quantifier (each mother).

- (58) [A group of mothers are having tea while through the window they watch their children playing outside. The children find a box of very convincing toy weapons under the playground, and they each take one and begin to play with it. The mothers glance up from their tea and begin to scream.]
 Naive onlooker A: Why is every mother screaming?
 Onlooker B: Because it looks like her child is playing with a weapon.

The understood reading of onlooker B's statement is that each individual mother's impression is that her own child is playing with a weapon, hence the screaming.²⁰ The fact that the PGOAL is able to take a quantificational antecedent indicates that it is, at least sometimes, a bound variable; its reference being partly fixed by another element. This suggests that it is different from the *pro* of null-subject languages, which make the phi-features of the referent of *pro* visible through rich verbal agreement, and thus do not require an antecedent to be mentioned in the previous discourse in order for the referent to be identifiable (Chomsky, 1981, Huang, 1984). Of course, because the PGOAL also happens to appear in the position of an indirect object, it does not interact with verbal agreement in any way. This not only means that it cannot be licensed through verbal agreement in the same way that the *pro* of consistent null-subject languages can, but it also means that we are unable to determine whether it possesses phi-features or not. Furthermore, it does not c-command the embedded clause, so there we have no straightforward way of establishing its binding properties.

Split antecedents

Like non-local antecedents, split antecedents are also known to be incompatible with Obligatory Control (Landau, 2003). Though it is difficult to conclusively determine the PGOAL's referent without the appealing to binding tests, sentences like (59) seem to show that the PGOAL can form a plurality from split antecedents:

- (59) Wendy went over the terms of the contract with Rob. It sounded like it was unreasonable, so they agreed not to sign it.

²⁰Admittedly, this example does not conclusively exclude the possibility that the PGOAL takes an arbitrary or generic value. I will provide arguments against this view in later chapters.

The most natural interpretation of (59) is that the contract sounded unreasonable to both Wendy and Rob.

Pragmatically inferable referents

While most null-subject languages permit the omission of subjects when the verbal agreement is strong enough to identify them, there is a great deal of variation within the so-called *pro*-drop languages. Languages like Japanese and Chinese, which feature no agreement morphology at all, tend to allow the deletion of elements in a number of different environments, so long as the referent is pragmatically inferable (Huang, 1984, Holmberg, 2005). And though Stephenson (2007) does not specify precisely what kind of pronoun her null referential *pro* might be, presumably this is the type of element she has in mind, as she claims that something like (48), repeated as (60), can be used referentially if Sam is somehow available in the surrounding context.

(60) This rollercoaster is fun *pro*_{Sam}

This silent element presumably requires no linguistic antecedent, but has its value resolved based on contextual and pragmatic factors. When attempting to construct CRCs with strictly inferred PGOALS, the situation becomes murky.²¹

(61) [Inside an illuminated tent, Jack is eating candy from a bowl on a low table. However between the tent wall and the table in a tall garbage can. Outside the tent, his wife Jill can only see his silhouette, and is standing on the side of the tent closest to the garbage can, and when Jack reaches his arm into the bowl of candy, Jill sees his arm plunging into the garbage can. Two of their friends, Ben and Ken, are sitting in lawn chairs parallel to the opening of the tent, and can see both Jack and Jill. Ben says to Ken:]

a.?#It looks like Jack is eating out of the garbage.

b.?#Jack looks like he is eating out of the garbage.

As both Ben and Ken can see that Jack is in fact eating out of a candy bowl, the intended PGOAL in these sentences is meant to be Jill. Without any previous mention of Jill in Ben and Ken's previous discourse, however, many informants find that this reading is unavailable for both the CRC and expletive variant.²² If this kind of scenario is not sufficient to license a silent PGOAL, it would mean that the PGOAL is not the type of element that can be freely omitted when it is inferable from the

²¹Though, to be fair, no more murky than with POPTs

²²Thanks to audiences at the 2015 CLA Meeting for helpful feedback on these types of examples.

wider context, and rather that it does have some additional (possibly structural) requirements. Let's try another:

- (62) [Mike and Martha's daughter Sally has a doll named Patches that she believes comes alive when nobody is looking. While Martha is cleaning the house one day, she finds Patches on the floor, so she picks her up and puts her on the windowsill, and in doing so she accidentally knocks over the potted plant that had also been on the windowsill. Mike and Martha watch as Sally comes into the room, and gazes upon the scene with a horrified look on her face. Martha whispers to her husband:]
It looks like Patches pushed the plant onto the floor.

Again, Martha and Mike, being adults, are aware that dolls do not come alive, and thus the intended PGOAL is Sally. While the corresponding CRC would be odd because the doll herself is not really the source of perception upon which the claim is based, (62) feels acceptable for many speakers. Why is the implicit PGOAL in (62) more easily inferable than in (61)?

Looking more closely at the contexts that allow an implicit third-party PGOAL, a common trait emerges. All these cases feature a prior discourse context in which the PGOAL's thoughts, impression, emotions, etc. have been made salient. In our first example of a third party PGOAL, repeated in (63), Jim's question highlights the dog's perspective:

- (63) Jim: Why is the dog growling at the mirror?
Bob: It looks like another dog is staring at him.

I turn for a moment to a related example discussed by Stephenson (2007) (originally raised by Egan et al. (2005)), in relation to the judge parameter that she claims is operative with epistemic modals.

- (64) [Context: Ann is planning a surprise party for Bill. Unfortunately, Chris has discovered the surprise and told Bill all about it. Now Bill and Chris are having fun watching Ann try to set up the party without being discovered. Currently Ann is walking past Chris's apartment carrying a large supply of party hats. She sees a bus on which Bill frequently rides home, so she jumps into some nearby bushes to avoid being spotted. Bill, watching from Chris's window, is quite amused, but Chris is puzzled and asks Bill why Ann is hiding in the bushes. Bill says:]
I might be on that bus.
(Stephenson 2007: ex. (59))

Stephenson claims that epistemic modals are always dependent on the judge parameter, i.e., that they are always interpreted from the perspective of the speaker or local attitude holder, and thus that they can never take a silent referential argument. Bill's statement in (64) seems to contradict this claim, as it expresses what *might* be the case according to Ann. However Stephenson argues that Bill's answer actually involves ellipsis,²³ such that his full statement has the form of (65):

- (65) Ann is hiding in the bushes because I might be on that bus.
(Stephenson 2007: ex. (60))

This means that Bill's answer is actually embedded under the *because*-clause. Stephenson suggests that *because*-clauses "express a person's conscious reasoning or rationale" (p. 506), and thus shift the judge parameter in the same way as an attitude predicate. The same reasoning can easily be applied to CRCs with implicit PGOALS. Bob's answer in (63) would also be understood as reflecting ellipsis:

- (66) Bob: The dog is growling at the mirror because it looks like another dog is staring at him.²⁴

If we assume that (66) is the underlying form of Bill's utterance, then the implicit PGOAL (the dog) not only has an antecedent in the higher clause, but its antecedent is an individual whose conscious reasoning is expressed. But in a similar dialogue where the form of the appropriate answer would not contain a *because*-clause, the exchange sounds somewhat odd:

- (67) Jim: What is the dog growling at?
Bob: ?#It looks like another dog is staring at him.

Note that the oddness does not originate solely from the fact that the question and answer are of different forms. (68) is also not a direct answer to the question, but is improved from (67).

- (68) Jim: What is the dog growling at?
Bob: He thinks that another dog is staring at him.²⁵

²³In order for it to be a proposition of the appropriate form; see Hamblin 1973, Groenendijk and Stokhof 1984.

²⁴Thanks to Bronwyn Bjorkman for this suggestion.

²⁵Thanks to Cassandra Chapman for suggesting this alternant.

Though both (67) and (68) serve as pragmatically reasonable responses to Jim's question, it seems that the lack of an appropriate antecedent for the PGOAL in (67) leaves Bob unable to refer to the dog implicitly.²⁶ The same logic can explain the felicity of example (58), dialogue repeated in (69):

- (69) Onlooker A: Why is every mother screaming?
 Onlooker B: Because it looks like her child is playing with a weapon.

Here the *because* is overt, strongly supporting that assumption that the rest of the matrix clause has been elided. Here as well, the implicit PGOAL has an available antecedent whose consciousness has been made salient in the higher clause.

We can now see why Ben's attempt to use Jill as an implicit PGOAL in (61) fails. Neither Jill nor her perspective has been made salient in the discourse between Ben and Ken, or even in the surrounding situation. Ben thus has no recourse to implicitly refer to her as the PGOAL. In (62) the situation is slightly different. Sally has shown an observable reaction to the scene at hand, via her horrified expression, and as her mother, Martha likely has the relevant knowledge of and experiences with Sally in order to be able to interpret her behaviour. Martha and Mike's attention is focused on Sally as she enters the room, and so Martha's utterance is naturally construed as her providing Mike with an explanation of Sally's behaviour. Thus it does not seem terribly far-fetched to assume that, given the shared knowledge and context between Mike and Martha, this utterance also involves ellipsis, something like:

- (70) Sally looks horrified because it looks (to her) like Patches pushed the plant onto the floor.

In this way, the implicit PGOAL in (70) has an available antecedent as well. Similarly, we would expect that the sentences in (61) would improve if Ken were to ask "Why is Jill scowling in disgust?"

We also see implicit PGOALS in many cases other than *because*-clauses, where the PGOAL's thoughts, speech, etc. have been referred to previously:

- (71) John said Mary looked (to him) like she'd been through an ordeal.
 (72) Jen thinks Greg seems (to her) like he's been under stress lately.
 (73) Carl imagined that Lisa had sounded (to him) like she was happy to see him.

²⁶However it may be that certain speakers are willing to accommodate the pragmatic intention of the question to the extent that they find (67) acceptable.

In this section, we have seen that the PGOAL can have a long-distance antecedent, act as a bound variable, and retrieve its reference from individuals who have been previously mentioned in the discourse. This sounds very much like the description of a standard pronoun, and yet there is a specific set of contexts which allow the use of a silent PGOAL. In the next section I will outline my proposal for the PGOAL's representation.

3.1.2 P-GOAL as a logophoric pronoun

The set of contexts that allow the PGOAL to be left unpronounced - contexts involving the individual's speech, thought, reasoning, etc. - almost perfectly resemble the set of contexts that license what are known as *logophoric* pronouns. Logophoric pronouns are found in some African languages, and they are particular pronouns whose reference in embedded environments is supplied by an antecedent "whose speech, thoughts, feelings or general state of consciousness is reported" (Clements, 1975, p. 141). Example (74) features an example of the logophoric pronoun *yè* from Ewe:

- (74) Kofi be *yè*-dzo
 Kofi say LOG-leave
 "Kofi said that he (Kofi) left."
 (Clements 1975: ex. (1))

This logophoric pronoun obligatorily refers to the author of the speech act, Kofi. Logophoricity has also been useful in accounting for instances of anaphors which do not obey binding Condition A (i.e., do not appear in the binding domain of their antecedent (Reinhart & Reuland, 1993)), such as (75).

- (75) **Lucie** thought that a picture of **herself** would be nice on that wall.
 (Reinhart and Reuland 1993: ex. (45b.), emphasis added)

Unlike Condition A anaphors, Reinhart and Reuland claim that these logophoric reflexives are not licensed by syntax, but rather by discourse factors such as perspective and center of communication. Thus, they are licit in contexts where an individual's consciousness is highlighted.

Landau (2001, 2015) also argues that logophoricity is responsible for identifying the controller in non-obligatory control (NOC) constructions. Interestingly, like the PGOAL, the PRO in NOC can take both long-distance and split antecedents. As Landau points out: "for a (non-deictic) pronoun to be felicitous, its antecedent merely needs to be familiar in the discourse. A logophor, however, can only refer to antecedents that are sources or targets of mental/communicative reports" (2001, p.

147). This is what we have seen to be the case with the PGOAL: the conditions on its antecedent are less strict than traditional binding conditions (the binder can be non-local), but stricter than the conditions on pronominal co-indexation (the antecedent must be logophoric). I thus propose that the silent element which I have been referring to as PGOAL is best characterized as a null logophoric pronoun.

As was demonstrated in Section 3.1.1, the PGOAL can refer either to the speaker or to an antecedent that is found in a logophoric context. Note that without the appropriate logophoric context, a salient antecedent is not sufficient to license a PGOAL interpretation. Consider the following minimal pair:

- (76) a. Vince is crying because it appears as though his girlfriend is kissing another guy.
 b. Vince is crying while it appears as though his girlfriend is kissing another guy.

Both (76-a) and (76-b) are compatible with the interpretation that the speaker holds the impression that Vince's girlfriend is kissing another guy, and are not inconsistent with a scenario where Vince also has this impression. Crucially, however, (76-b) is incompatible with a scenario in which the speaker does not hold this impression themselves. In other words, only (76-b) cannot support a reading where Vince is the unique PGOAL. If the speaker (reporting to a friend over the phone on how the party is going, for instance) can see that Vince's girlfriend is only pretending to kiss another guy as a joke, he or she cannot utter (76-b). The conjunction *while* merely expresses cotemporality, and thus does not reference Vince's mental state or reasoning in the same way that *because* is able to; Vince in (76-b) therefore is not a possible logophoric antecedent.

In a sense, this condition is entirely logical: in order for a speaker to convey another person's perspective without overt marking, they need the appropriate context to do so. And as a silent element, the PGOAL needs a method of recovery. In most cases, we assume that it is the speaker, and this is unsurprising if we understand PGOAL as a logophoric pronoun; speakers' utterances most commonly express their own thoughts, and refer to their own mental states. But when another antecedent is available in a logophoric context, this reading may overpower the default interpretation, though it does not exclude it. However when the speaker is ruled out as the PGOAL by semantic or pragmatic factors (e.g. the complement expresses a proposition that conflicts with the speaker's knowledge or perspective ((57), (61), (62), etc.), the hearer will search for another possible antecedent. If one is not available in a logophoric context in the preceding discourse, the sentence will be infelicitous. In other words, the implicit PGOAL/logophor is licensed if the speaker is expressing either their own views, or the

views of a logophoric antecedent.²⁷

Of course, these conditions should only apply when the PGOAL is left unpronounced. I assume that when there is an overt PP containing a PGOAL (e.g. *to me*), there are no antecedence conditions. I further assume that a null PGOAL also appears with a null preposition (*to*), from which it is assigned Case. Thus it appears in complementary distribution with its overt counterpart.

With regards to the internal composition of this silent pronoun, matters are still unclear. Holmberg et al. (2009) investigated a similar phenomenon in what he calls *partial null subject languages* like Finnish and Brazilian Portuguese, where the subject of an embedded clause may be omitted under certain conditions:

- (77) O João₁ disse que (ele₁) tinha comprado uma casa [BP]
 the João said that he have-PST.3SG bought a house
 “João said that he had bought a house.”
 (Holmberg 2009: ex. (12a.))

The embedded subject in (77) must obligatorily refer to the matrix subject, João, who is the author of the reported speech act. The embedded subject of such a construction can be omitted in complements of saying, thinking, or perceiving, when there is a matrix subject representing the sayer/thinker/perceiver.²⁸ In the partial null subject languages that Holmberg discusses, the null element must be in a particular structural relationship with its antecedent in order to be controlled (Holmberg, 2005). He thus assumes that the null element must be deficient in some way. Specifically, he proposes that the null subject of the embedded clause is represented as a *phi*-Phrase, which consists of a set of valued *phi* features, but has an unvalued [D] feature. Without a valued [D] feature, it cannot be referential, and so it must be controlled by a DP (usually) in the next clause up.²⁹

Though Holmberg does not consider this phenomenon to be related to logophoricity, the contexts which allow null embedded subjects in these languages bear a certain similarity to the contexts which allow a null PGOAL in PCs. Yet the structural conditions on the relation between antecedent and null element are much more clearly and strictly defined in Holmberg’s cases: the null subject must be c-commanded by the argument in the higher clause. Nevertheless, we might adopt the assumption that the

²⁷Frascarelli (2007) has proposed that every phrase has features representing the speaker in the C-domain, so that the speaker is always available as an antecedent. Something of this nature might be relevant here.

²⁸As well as in a few other environments, see Holmberg et al. 2009.

²⁹The exception being Finnish, which allows connectivity under specificational predication (Holmberg et al., 2009).

PGOAL is somehow deficient, as it requires an antecedent in order to be interpreted.³⁰ It may be the case then that the PGOAL is also a bundle of *phi*-features, lacking a [D] feature value. The valuation might occur through some kind of chain formation or feature matching; however the precise mechanisms through which PGOAL is linked to its antecedent is a matter I must leave for future research.

I have proposed that when no overt PGOAL is specified in a PC, there is a null pronoun in its place which is anaphoric to an individual whose consciousness is expressed. Recall from Chapter 2 that in order to account for similar problems of determining the evaluator of a POPT, Stephenson (2007) posits the existence of two different types of silent pronouns: judge-dependent PRO_J , and referential *pro*. However Pearson (2013) has shown that this analysis sometimes makes incorrect predictions. The judge parameter is not necessarily bound locally, as Stephenson claims:

- (78) The cat thinks that John thinks that the cat food is tasty (since he keeps buying it for her).
(Pearson 2013: ex. (17))

This POPT example is equivalent to (57); the local attitude holder (John, Jimmy's mom) is not understood as the judge or PGOAL, respectively. Furthermore, Pearson points out that the lack of restrictions on where *pro* can appear in Stephenson's system fails to rule out cases where a referential interpretation is not in fact available (cf. Pearson 2013). If we were to apply Stephenson's analysis to the PGOAL, the same problems would arise: the local attitude holder does not necessarily represent the PGOAL (example (57)), and the availability of a referential *pro* would permit a third-party PGOAL reading in (61) and (76-b), where it does not seem to be available. While I do not claim that my proposal should extend to POPTs, it appears that for PCs at least, positing these two different silent elements would not accurately capture the facts. My proposal instead capitalizes on the observation that the environments in which these elements are found all necessarily involve focus on an individual's perspective. The most common case is for a speaker to express their own perspective, but under the right conditions, speakers can covertly refer to the perspectives of others. Thus under my proposal, the PGOAL need only be represented by a single element, and its referent is determined by a single mechanism: logophoricity.

³⁰This is not conclusive, as certain speakers may accept exchanges like those in (61). It may be the case that PGOAL can be used deictically (perhaps only for certain speakers), and thus it might in fact have a [D] feature. More naturalistic corpus data will be required to resolve this issue.

3.2 Situational restriction

If the PGOAL is indeed an obligatory component of PCs, how does it affect the formation and interpretation of these sentences? I have already hinted at the fact that the content of the utterance must be compatible with an available PGOAL's knowledge or impressions, otherwise the sentence will be infelicitous. Consider the following example, which takes place in the same campsite context as (61):

- (79) [Jill has an irrational fear of Big Bird. Jack is inside the tent making shadow puppets. He's remarkably good at it, and eventually makes one in the shape of Big Bird. Outside the tent, Jill screams.]

Ken: Why is Jill screaming?

Ben: Because it looks (to Jill) like Big Bird is in the tent.

Even though Ben and Ken are reasonable people who realize that Big Bird is fictitious, Ben's utterance is acceptable when Jill is specified as the explicit or implicit PGOAL. However the corresponding CRC would be strange:

- (80) #Big Bird looks (to Jill) like he's in the tent.

Though the content of the utterance is compatible with Jill's impressions, Ben's statement in (80) is infelicitous. In other words, it cannot be interpreted as a paraphrase of "Jill thinks that Big Bird is in the tent." We might understand this strangeness as being due to the PSOURCE condition that was discussed in Chapter 1: the matrix subject is interpreted as the source of evidence for the claim. As Big Bird is not actually present, he cannot be perceived. But the oddness remains in (81):

- (81) [Jack has been out late drinking, and happens to be wearing all black. When he comes home he is too drunk to find the right key, so he spends a lot of time fumbling around at the keyhole and trying to pull the door open. His wife Jill has stayed up late watching a news story about a burglar who has been active in their neighbourhood, and is sitting in the living room. Through the curtains, she sees a man in black busy at her door handle and begins to search for weapons. Their next-door neighbours Ben and Ken are watching all of this through the window. Ken asks Ben why Jill is so frightened.]

a. Ben: It looks (to Jill) like the burglar is trying to break in.

b. Ben: # The burglar looks (to Jill) like he is trying to break in.

Though the speaker (Ben) and the PGOAL (Jill) have different beliefs about his identity, (81) features an individual who is perceivable. Yet as (81-b) shows, perceiving this individual is not enough: when in matrix subject position, the definite description *the burglar* must be interpreted as *the burglar according to the speaker*, and not *the burglar according to Jill*. As such, (82) is fine:

(82) Ben: Jack looks (to Jill) like he is trying to break in.

Though both (81-a) and (81-b) are intended to express Jill's perspective, the individual Jack can only be referred to as "the burglar" when the DP appears in the lower clause. I thus argue that the PGOAL is the individual whose perspective restricts the worlds/situations of evaluation for the *embedded* clause. In this section, I will use the framework of Percus (2000) and Keshet (2008), which incorporates world-denoting variables into the syntactic structure, in order to briefly demonstrate the logic behind this claim. This framework will be useful for illustrative purposes, however it should be noted that I ultimately will opt for a different implementation in my final analysis.

3.2.1 World variables


Percus (2000) advocates the view that linguistic structures contain variables over possible worlds or situations,³¹ which attempts to account for certain types of semantic ambiguity. Without going into undue detail, his system makes the following assumptions:

1. Clauses take lambda abstractors in the Specifier of IP, which create functions from situations to truth values.
2. Certain lexical items select for what Percus calls *situation pronouns*: unpronounced variables which take situations as their value.
3. Each situation pronoun in the structure is coindexed with one of the lambda abstractors.
4. The coindexation of situation pronouns determines the world/situation in which the element that it combines with is interpreted.

Each VP and strong DP (Keshet, 2008) selects for a situation pronoun. Given the propositional attitude sentence "Mary thinks my brother is Canadian," Percus represents the positions of the situation variables with "S" and "T" in Figure 3.1:

³¹Percus acknowledges that sometimes the domain of quantification is smaller than entire worlds, and thus he uses the term *situation*. For the moment, the terms *world* and *situation* can be thought of as equivalent.

Figure 3.1: Percus 2000: ex. (26)



An element that selects for a situation pronoun has the possibility of being interpreted with respect to one among (possibly) many situations; in the above sentence, elements could either be evaluated with respect to the actual world, i.e. the matrix situation, or they could be evaluated with respect to Mary's belief worlds. These multiple coindexation possibilities can often give rise to semantic ambiguity, as we will see below.

After demonstrating that the distribution of these situation elements closely resembles the distribution of ordinary pronouns, Percus proposes a binding theory that can account for this. His most notable contribution is the binding condition which he calls Generalization X:


- (83) Generalization X:
 The situation pronoun that a verb selects for must be coindexed with the nearest λ above it.
 (Percus 2000: ex. (34))

This means that the s_1 in Figure 3.1 can only be coindexed with λ_1 , as is shown, but also that the situation pronoun selected by the embedded verb (**S**) must be coindexed with λ_2 . This is what gives us the interpretation that the brother *is Canadian* in Mary's belief worlds. The situation pronoun on the DP *my brother* (**T**) however, is not selected by a verb, and thus Generalization X does not apply, and it can be

coindexed with either λ_1 or λ_2 , the utterance situation or Mary's thought situations,³² respectively. Thus, the sentence can be true both in situations where the speaker does in fact have a brother, whom Mary believes to be Canadian (s_1), and also in situations where there is a person who Mary simply *believes* to be the speaker's brother (even if he is not), whom she also believes to be Canadian (s_2). The DP *my brother* is ambiguous in this way; it can be coindexed with either of the lambdas that dominate it.

Applying this to the sentences in (81), we can see why the DP *the burglar* is only felicitous in the expletive variant:


(84) Expletive variant:



Just as the attitude holder Mary in Figure 3.1 is the individual whose belief worlds are represented in the embedded clause, the PGOAL Jill in (84) is the individual whose impression worlds are represented in the embedded clause of a PC. In its position in the complement clause, the DP *the burglar* is dominated by both λ_1 and λ_2 , and so it can in principle receive either index. When it receives the index of λ_2 , as in (84), it has opaque use, i.e., it is interpreted with respect to the situations consistent with Jill's impressions. However the corresponding CRC would have the following structure (85):

(85) Copy raising construction:

³²To use Percus's terminology, the DP has either transparent use (actual world) or opaque use (possible worlds.)



Because *the burglar* appears in the matrix clause, it can only be coindexed with λ_1 , the utterance situation. This would entail that this individual is a burglar in the actual world. But in our scenario the speaker is well aware that it is in fact Jack who is at the door, and thus the sentence is infelicitous. The (in)felicity of (84) and (85) tells us that the PGOAL is the individual whose perspective constrains the worlds of evaluation for the embedded clause, akin to the subject of an attitude predicate.

3.3 Summary

In this Chapter, I reviewed the observable properties of the implicit PGOAL element, and concluded that it can be understood as a silent logophoric pronoun: it is anaphoric to an individual whose perspective has been made salient in the discourse. Though the factors that determine which of the possible antecedents represents the PGOAL are contextual in nature and sometimes involve world-knowledge, once the referent of the PGOAL is fixed, it becomes an important component of the semantics of these constructions. The PGOAL, much like the attitude holder of a propositional attitude, is the individual whose impressions of beliefs provide the domain of evaluation for the proposition expressed by the embedded clause. Contrary to what has previously been assumed, this makes the PGOAL a necessary element in the structure; an argument, rather than an adjunct. In the next Chapter, I will expand the idea of restricted quantification over possible worlds into a modal analysis in the sense of Kratzer (1977, 1981) and von Stechow & Heim (2002).

Chapter 4

A modal semantic account

In this chapter I will argue that preceptual constructions are *modalized assertions*, and after reviewing the basis framework, I will outline my full analysis.

4.1 False truth scenarios

To set the stage for my forthcoming arguments, I begin with another puzzle, which I will call the *puzzle of false truths*. This refers to the fact that the utterer of a PC may express an embedded proposition that she believes to be false, and yet we judge the utterance to be true. First consider a standard use of PCs, like in (86).

- (86) [Word has gotten around the neighbourhood that Jim's classmate Timmy's dog has just died. At school the next day, Jim notices that Timmy's eyes are red and puffy and his lip keeps quivering. Jim says to his friend:]
- a. It looks like Timmy's really upset.
 - b. Timmy looks like he's really upset.

In this commonplace example, the embedded proposition *Timmy is really upset* is surely a proposition that Jim believes to be likely true, or at least considers possibly true.³³ What is puzzling is that the same utterances can be used truthfully in a scenario where the speaker is not under the impression that the embedded proposition might be true.

³³I assume that because Jim does not have direct access to Timmy's mental state and cannot know for sure how he feels, pragmatic principles require that he make a weaker assertion than a plain declarative, and thus the use of a PC signals that the speaker does not have complete knowledge of *p*.

- (87) [Jim and his brother are at their friend Timmy's house for dinner. Timmy excuses himself to go into the kitchen to cut up some onions, and when he comes back, tears are streaming down his face. Jim says to his brother:]
- a. (Haha,) it looks like Timmy's really upset.
 - b. (Haha,) Timmy looks like he's really upset.

In this scenario, Jim and his brother are both perfectly aware that Timmy has just been cutting onions, and that he's not actually upset. We have to assume that Jim, the speaker, places himself as the PGOAL of these utterances, as there is no other logophoric antecedent available in context. It is odd then, that the speaker seems to be able to express that he himself holds the impression that Timmy is upset, whilst knowing that he's really reacting to the onions. And notice that the speaker is not using any rhetorical devices such as sarcasm or hyperbole, nor is he violating any conversational maxims; thus cases like these cannot be written off as reflecting some kind of pragmatic phenomenon. His utterance is literally true, just as true as would be in (86); and this is puzzling.

A further puzzle is that not all copy raising verbs have this property. The following responses to the scenario in (87) are infelicitous.

- (88)
- a. #Haha, it sounds like Timmy's really upset.
 - b. #Haha, Timmy sounds like he's really upset.
 - c. #Haha, it seems like Timmy's really upset.
 - d. #Haha, Timmy seems like he's really upset.

While the infelicity of (88-a) and (88-b) is undoubtedly due to the fact that there is no auditory information in this context which could generate the impression, the infelicity of (88-c) and (88-d) is somewhat surprising, given that *seem* is generally regarded as simply being a very semantically bleached perception verb, meaning that it vaguely refers to perception in general, but does not specify a sensory modality (Landau, 2011, Rett et al., 2013, Kim, 2014).

If we construct a similar context which features auditory information, we see that the verb *sound* can be used to express these false truths as well, but *seem* remains odd:

- (89) [Your brother Ryan has just had his wisdom teeth removed, and your mother asks you to call him to see if the surgery went alright. Ryan answers, and his mouth is still full of cotton so he has trouble speaking. You report to your family:]
- a. Haha, it sounds like Ryan has a swollen tongue!

- b. Haha, Ryan sounds like he has a swollen tongue!
- c. #Haha, it seems like Ryan has a swollen tongue!
- d. #Haha, Ryan seems like he has a swollen tongue!

When the context is auditory in nature, utterances like (89-a) and (89-b) are acceptable and true, even if Ryan had begun the conversation by saying “Sorry, my mouth is full of cotton so I can’t talk clearly.” But the utterances with *seem* in (89-c) and (89-d) are strange here as well. Rett et al. (2013) conducted an online rating study in which they tested the acceptability of CRCs and their expletive variants in response to contexts which contained different types of sensory information. Their results show that the verb *seem* is acceptable in response to a wider range of contexts than *look* or *sound*, suggesting that it is, in Rett et al.’s terms, the least restrictive copy raising verb. It is interesting, and deserving of explanation, that *seem* should be the least restrictive in terms of the sensory sources it allows, but that it appears to be the most restrictive in these false-truth cases.

As the speaker acts as the PGOAL in these sentences, a successful semantic analysis of PCs will have to account for how it is that they can be used truthfully with complements that the PGOAL believes to be false, and why it is that not all copy raising verbs allow this. And the only possibility is indeed for the PGOAL to be the speaker; appealing to the notion of an arbitrary or generic PGOAL to account for cases like (87-a), (87-b), (89-a) and (89-b) would not be satisfactory, as it would fail to explain the unacceptability of the examples with *seem*. It is not at all clear why this resource would not be available to rescue the *seem* examples. One might argue that the word *like* contributes a merely comparative meaning, and thus that it is not surprising that false-truth utterances like (87) are possible. But again, the fact that CRCs with *seem* do not allow this kind of innocent comparison demonstrates that the semantics of *like* alone cannot account for this property.

In the next section, I will propose an analysis of PCs which treats them as modalized assertions, using the possible world framework proposed by Kratzer (1977, 1981) and developed by von Stechow & Heim (2002). Under this framework, I will argue that the differences between copy raising verbs observed above can be understood as reflecting the different accessibility relations denoted by the verbs; i.e. the different types of propositions which are allowed into the *modal base*.

4.2 Proposal

The conventional way of thinking about the semantics of modals is that they are quantifiers over possible worlds. Angelika Kratzer (1997, 1981) was the first to implement this idea into a viable semantic framework. Kratzer proposed that a

modal (e.g. *must*, *can*, *might*, etc.) lexically supplies only a quantificational *force*, either universal ($\forall$) or existential ($\exists$), and that the set of worlds that is quantified over is determined by context, or what she calls a *conversational background*. A conversational background consists of two parts: a *modal base* and an *ordering source*.

The modal base is, loosely speaking, the part of the conversational background that determines the set of *accessible* worlds. The modal base essentially tells us what kinds of worlds we consider when we interpret a modal statement. For example, the conversational background could tell us that the relevant worlds to be considered are those that are compatible with *what is known* in the actual world, i.e., an epistemic modal base, or it could be that we are to consider only those worlds that are consistent with *some relevant set of circumstances*, i.e., a circumstantial modal base. To give an example, the English modal *must* carries universal force, quantifying over a complete set, but depending on the conversational background, the set of worlds that are quantified over will differ.

(90) Frank must be at home today.

Minimally, this sentence expresses that all possible worlds are such that Frank is at home today. But its interpretation will be relative to whatever conversational background is active in the discourse. For instance imagine that someone at Frank's office has just inquired as to why Frank's desk is vacant. The speaker knows that Frank has arranged it with the boss that he can work from home on days when it rains, and it is currently raining. In this case a reply of (90) is understood to have an epistemic modal base: of the worlds consistent with what the speaker knows, all are worlds in which Frank is at home that day. A simpler paraphrase might be "Given all the things I know, it is necessary that Frank is at home today." Alternatively, if Frank is a teenager who has been grounded by his parents, his mother's utterance of (90) would be understood to have a circumstantial modal base: of the worlds consistent with the current set of circumstances (i.e. Frank being grounded), all are worlds in which Frank stays at home. (As will clarify below, though this sentence receives a *deontic* interpretation, deontic modality involves a circumstantial modal base (von Stechow, 2006).)

More specifically, what the conversational background actually provides to the modal base is a set of propositions: the set of all propositions that express young Frank's circumstances, or the set of all propositions that express what adult Frank's coworker knows. Any given possible world is deemed an *accessible* world iff it is consistent with the propositions in the modal base.

Without the modal base, any number of worlds could be compatible with the scenario where Frank's desk is vacant. Frank might be on holiday. He might be running late. He might have fled the country, or died. We would not be able to

universally quantify over possible worlds, because “Frank is at home today” is not true in *all* those worlds. Instead we quantify over the worlds that are consistent with the propositions in the modal base; in this case, the propositions that make up Frank’s coworker’s knowledge. This will include things like “Frank is not at work today,” “The boss permits Frank to stay home on rainy days,” and “It is raining today.” So while it is not true that *all* conceivable possible worlds are worlds in which Frank is at home today - after all, anything is possible - it is true that all the worlds consistent with Frank’s coworker’s knowledge are worlds in which Frank is at home.

But in fact the modal base alone is not quite enough to arrive at the desired meaning. The second component of Kratzer’s modal system, the ordering source, is the (contextually determined) set of propositions that ranks the worlds in the modal base. Because even in the set of possible worlds that are consistent with what Frank’s coworker knows, or with what teenage Frank’s circumstances are, there are still a great number of worlds in which strange and unexpected things happen. Frank’s coworker may know about Frank’s arrangement to stay home on rainy days, but he will surely also know that people sometimes spontaneously drop dead. So though they are far less plausible, worlds in which Frank is absent from work because he is dead are not technically inconsistent with his coworker’s knowledge. This means that not universally *all* the coworker’s epistemically accessible worlds are worlds in which Frank stayed home today.³⁴ This is where the ordering source comes in; in this case, where we have an epistemic modal base, the conversational background provides a stereotypical or normative ordering source, which contains the set of all propositions expressing the way that things typically are. The worlds in the modal base which are consistent with the highest number of propositions in the ordering source are designated as the “best” worlds. It is this set of best worlds which the modal ultimately quantifies over; in all of the most stereotypical worlds consistent with what Frank’s coworker knows, Frank is at home today.

Like the modal base, the ordering source can come in a number of different types, depending on context. In the scenario where Frank is a grounded teenager, the ordering source is deontic; his mother is expressing something like a command. The worlds consistent with Frank’s circumstances make up the modal base, as mentioned above, and it is this set of worlds that are then ranked according to the set of all propositions that express the rules of Frank’s punishment. Though in the actual world Frank finds himself grounded, there exist many possible worlds where Frank disregards his parents orders, and goes to the movies for instance.³⁵ The ordering source ensures that we only quantify over worlds in which the terms of Frank’s punishment are

³⁴At least not by choice.

³⁵Notice that the modal base could not be deontic, one reason being that if we only considered worlds where Frank obeyed his parents rules, he would not have been grounded in the first place.

obeyed.

That, in a nutshell, is the possible world theory of modal semantics. Since Kratzer’s initial work on modal verbs, this framework has been adapted to provide semantic accounts of many other linguistic phenomena, such as conditionals (Kratzer, 1986, von Fintel & Iatridou, 2002), propositional attitudes (Hughes & Cresswell, 1996, Stalnaker, 1999), the progressive (Portner, 1998), the future (Copley, 2009, Giannakidou & Mari, 2012), and many more.

Building on the seminal work of Hintikka (1969), von Fintel & Heim (2002) begin to develop a possible world semantics for propositional attitudes. The idea is largely the same as the semantics of modals, however rather than having the set of accessible worlds be determined by a conversational background, the function that determines the set of relevant worlds is given by the lexical content of the attitude verb. For example, the sentence “George believes that Henry is a spy” (von Fintel & Heim, 2002, p.19) is true if in all the worlds that are compatible with George’s beliefs in *w*, Henry is a spy. The same applies to verbs like *hope*, *think*, *regret*, etc. Different propositional attitude verbs thus each correspond to different *accessibility relations*, which essentially are functions from individuals and worlds to sets of propositions, which gives us the set of all worlds where the proposition is true.

I will now illustrate how this type of machinery can be applied to PCs.

4.2.1 Accessibility relations

In Chapter 3, I argued that the PGOAL is both a syntactic and semantic argument of the copy raising verb, whether it is pronounced or not. The PGOAL is the individual who experiences the (visual, auditory, olfactory, etc.) sensations that give rise to an impression. In this way the PGOAL is very closely related to, if not equivalent to, an Experiencer or attitude holder. I also argued in Chapter 3 that the PGOAL has a role in restricting the worlds of evaluation for the embedded clause. Hopefully it is clear at this point that these facts suggest that the kind of possible world semantics discussed above might be useful here.

I wish to argue that each of the different copy raising verbs, like propositional attitude verbs, represents a different accessibility relation. The verb *look* represents the relation *compatible with what PGOAL can see*, the verb *sound* refers to the propositions *compatible with what PGOAL can hear*, etc. Importing Kratzer’s system, the modal base that accompanies the copy raising verb *look*, for example, will contain all those propositions which express what the PGOAL in *w* can see at *t*. This will take us to only those worlds which are compatible with what the PGOAL can see. Recall the strangeness of (61), where there is no logophoric antecedent available to act as the PGOAL, and so we are forced to interpret Ben’s utterances as reflecting his

own impressions. The visual information available to Ben at utterance time includes, among other things, the fact that Jack is reaching his hand into what Ben recognizes to be a candy bowl. The infelicity of his utterance (“Jack looks like he is eating out of the garbage”) can be explained: worlds in which Jack is eating out of the garbage are not consistent with Ben’s visual information, thus they are not accessible, and not included in the modal base.

The verb *seem* however, is not a purely perceptual verb; I argue that its accessibility relation is an *epistemic* one. It assigns to any world and individual the set of propositions that express what the PGOAL *knows* (or believes) in that world at utterance time. What PGOAL knows will include what he or she can perceive through various sensory sources, but it will also include knowledge of past events or facts. As I will demonstrate in the next section, this is the precise reason for the infelicity of *seem* in false-truth utterances.

I am not the first to suggest that the verb *seem* has epistemic qualities; this idea can be found in Iatridou 1990, Matushansky 2002, and Aijmer 2009, for instance. In addition to the fact that *seem* is infelicitous in false-truth utterances, evidence for the epistemic use of *seem* in PCs comes from the fact that in general, the use of *seem* with a third-party PGOAL is degraded, whether it is overt or covert:

(91) It seems to Jess like Mark is hungry.

(92) Vince is crying because it seems like his girlfriend is kissing another guy.

(91) is acceptable if Jess has previously expressed her impressions about Mark’s hunger to the speaker. Otherwise, in some sort of context where the speaker would have to infer what Jess’s thoughts are, this utterance is slightly odd. Likewise, in the scenario where the speaker of (92) can see that Vince’s girlfriend is not really kissing another guy, (92) is somewhat strange. This is because we do not have direct access to the epistemic states of others. (92) is odd because we only have direct access to the epistemic state of the speaker, who knows that Vince’s girlfriend is only pretending. We can be fairly certain, using our own visual faculties and our implicit knowledge of visual angles, of what another person can *see*, but we can rarely be certain of what they *know*, unless it has been reported to us. However in rare scenarios where we do have access to another individual’s knowledge, this line of reasoning predicts that PCs using *seem* would improve. This might be the case when a parent makes a statement about the epistemic state of their young child. Just like in context (62), a parent would know roughly what stage of development their child is in, what concepts he or she is aware of, and any peculiar beliefs that they hold. The following exchange thus seems reasonable:

- (93) [Mom and dad have a toddler son. One day they are all in the living room; dad is sitting with the toddler, and mom is pottering around cleaning while listening to music on her headphones. The toddler want his mom's attention, and calls to her. She can't hear him over the music, and the toddler begins to shout louder and louder, kicking his legs and banging his fists. When mom finally takes off her headphones, dad says:]
He was getting really mad because it seemed like you were ignoring him.

Dad knows that mom wasn't ignoring her son: she couldn't hear him. So we know that the PGOAL cannot be dad, as this proposition is inconsistent with dad's knowledge. But mom and dad recognize that their toddler has neither the world knowledge nor the cognitive capacity to understand that his mother was listening to music on headphones, and thus could not hear his cries. We would probably be inclined to grant the parents the authority to speak for their child's epistemic state, so that dad's utterance using the toddler as the PGOAL in (93) is felicitous. Assuming this reasoning about the different types of accessibility relations is on the right track, I now move on to the specifics of how the worlds to be quantified over are selected.

4.2.2 Domain of quantification

Returning to the false-truth scenario (87) from the beginning of this chapter, we will see how a modal semantic analysis can account for both puzzles surrounding these types of sentences. When the speaker, Jim, utters (87-a) or (87-b), the verb *looks* returns the set of all propositions that express what Jim, the PGOAL, can see at utterance time. This will include propositions such as:

- (94) 1. Timmy is standing in front of me.
2. Timmy has blonde hair.
3. Timmy has clear liquid coming from his eyes.

The modal base would then contain all possible worlds in which these propositions are true. In the case of *seem* however, (88-c) and (88-d), the modal base³⁶ will contain the set of all propositions that express what Jim knows/believes at utterance time, and this is a much larger set. It will include relevant propositions such as:

- (95) 1. Timmy has clear liquid coming from his eyes.
2. Timmy went into the kitchen to cut onions.

³⁶For simplicity, I use the term "modal base" to refer both to the set of propositions provided by the accessibility relation, and the set of worlds picked out by those propositions.

3. Timmy has been in a good mood all day.

If we were to simply apply universal quantification over the worlds picked out by these modal bases - so that our truth conditions would say that the utterance of (87-a) or (87-b) is felicitous iff the proposition “Timmy is really upset” is true in all of those worlds - we would have a condition that is too strong. As it stands, the sets of worlds picked out by *look* are not *incompatible* with worlds in which Timmy is not in fact upset. Many other worlds are compatible with the propositions in (94): worlds where Timmy has recently undergone laser eye surgery, worlds where Timmy is overcome with joy, and also worlds in which Timmy has been cutting onions. Yet on the other hand, existential quantification does not seem to be strong enough a condition; among the worlds picked out by *seem* (and compatible with the propositions in (95)), there is at least one world in which Timmy happens to be upset, yet uttering (88-c) or (88-d) is not felicitous.


To arrive at the appropriate truth conditions for PCs, the worlds in the modal base must be further restricted. Like modals under Kratzer’s (1977, 1981) system, the modal bases of PCs must be subject to ordering by an ordering source. I propose that the ordering source operative here is one that is stereotypical or normative, i.e., containing the set of propositions that express the way things typically are, or how things typically happen. Propositions applicable here might include:

- (96)
1. Liquid coming from the eyes is usually tears.
 2. Tears are usually caused by excess emotion.
 3. The emotion that causes tears is usually sadness.
 4. Cutting onions usually results in tears.

It is important to note that the stereotypical ordering is not absolute: it is relativized to the speaker. The propositions that feature in the stereotypical ordering source may vary slightly depending on the speaker or PGOAL; their experiences with the world to date, their beliefs about causation, their implicit biases, etc. It is this variation which allows for the “faultless disagreements” of the kind seen in (52). To one speaker the visual landscape may be characteristic of one state of affairs, but to a different speaker, it may be characteristic of another. To illustrate how this type of ordering gets us to the appropriate truth conditions, consider Figure 4.2., which is a visual representation of what the speaker’s (radically simplified) stereotypical ordering source might look like.

According to this ordering source, the majority of occurrences of tears are occurrences of tears caused by emotion, and the majority of tears caused by emotion are tears caused by sadness/being upset. And overall, tears caused by sadness make

Figure 4.2: Ordering source for context (87)



up a greater proportion of the total occurrences of tears than any other kind. The worlds included in the modal base for the sentence in (87-a) (“It looks like Timmy is really upset”) are all worlds in which Timmy has tears on his face, but the propositions expressing what PGOAL can see do not provide much more relevant information than that. The ordering source will use the relevant propositions about stereotypicality (the solid circles) to select the set of worlds from the modal base that are most typical or likely. In this case, given that Timmy has tears on his face, the most likely kind of world for us to be in is a world in which Timmy has tears on his face due to being upset, as this type of world is more common than any other. Using this first-past-the-post system of stereotypical ordering, we can say that “It looks like Timmy is really upset” is true iff in all the most stereotypical worlds compatible with what PGOAL can see, Timmy is really upset.

If we try to use a PC involving *seem* however, the modal base now contains the proposition that Timmy has just been cutting onions. This narrows the set of accessible worlds, so that we are now focusing on only those worlds in the intersection of occurrences of tears and occurrences of onion-cutting. Given this set of worlds, the most typical worlds are not ones in which Timmy’s tears were caused by sadness (though it is still a slim possibility), but rather ones in which his tears are due to onion-cutting. Thus it is false that in all the most stereotypical worlds compatible with what PGOAL knows, Timmy is really upset, and as a result the statements in

(88-c) and (88-d) are infelicitous.³⁷

This semantics not only provides reasonable truth conditions for perceptual constructions, but it also explains our intuition that the speaker/PGOAL commits themselves the embedded proposition when they use the verb *seem*, but not necessarily when they use any of the other perceptual verbs. Essentially, to use a PC containing *look*, *sound*, etc. is to make the claim: “*p* is the most likely state of affairs based on my visual/auditory/etc. information” - and nothing more. There is nothing additional that would entail that the speaker actually believes that *p* is the case. However there is an inherent contradiction in making the claim: “*p* is the most likely state of affairs base on my *knowledge*” when one knows (or at least believes) that *p* is not the case. This is the result of using epistemic *seem* in false-truth utterances.

4.2.3 P-SOURCE restriction

We are now much closer to understanding the complete truth conditions for PCs, however we are not quite all the way there. Still yet to be accounted for is the puzzle of the absent cook which was introduced in Chapter One. Currently, this semantics does not differentiate between CRCs and their expletive variants. It predicts that the sentence “Tom seems like he is cooking” should be felicitous simply iff in all the worlds compatible with what PGOAL knows, Tom is cooking. This would surely be the case in the absent cook scenario, where, though Tom is nowhere to be seen, there is ample evidence that he is cooking - but the utterance is not in fact felicitous. As we have seen, there is an apparent requirement that the matrix subject be perceived (through the relevant sensory mode) by the speaker, at least in the case of a thematic matrix subject (Landau, 2011). It is these cases that I will attempt to provide truth conditions for in what follows, ignoring, for the moment, the rarer non-thematic subject cases discussed by Landau (2011).

Some further restriction on the truth conditions of CRCs specifically is needed in order to arrive at the correct analysis. Before this can be done however, I will need to refine some of the assumptions and terminology that have been useful in illustrating the patterns so far, but which now need to be made more precise. Until now we have been dealing with possible *worlds*. Possible worlds are, by nature, complete worlds, extending to all the relations and entities within them. However there is a large body of literature that argues that in many cases, quantification over entire worlds is too strong. Kratzer (2011), along with many others, advocates the view that more often than not we are actually quantifying over parts of worlds, or *situations*. Situations can

³⁷However this does not mean that stating “It seems like Timmy has been cutting onions” would be felicitous in this scenario either. Epistemic modals (of which I argue *seem* is one) are incompatible with the expression of statements that already make up part of the speaker’s knowledge.

be as big as entire worlds, or as small as a teapot. A classic example of quantification over situations comes from quantifiers such as “everyone,” which rarely quantify over every existing person in the world, but rather every person in a certain situation:

- (97) [Context: Research assistants are conducting an experiment where they have to monitor sleeping people. They must wait for all the participants to fall asleep before they begin their tests.]
 Everyone is asleep, let’s begin.
 (adapted from Barwise & Perry 1983: p. 160)

Here it is obvious that *everyone* is only quantifying over the participants in the study, and not the research assistants or anyone else in the world, for that matter. What is quantified over is a small part of the world: a situation containing only the participants. While the exact ontology of situations is still a matter of contention, their implementation has led to breakthroughs in the analysis of a wide range of phenomena, including implicit quantifier domain restrictions, donkey pronouns, and exhaustive interpretations (Kratzer, 2011).

I will assume (with good reason) that situations, rather than worlds, are the relevant entities involved in the semantics of PCs. The idea of situational restriction goes back to Austin (1950) who argued that propositions are true of certain contextually salient situations, which he called *topic situations*. I adopt this terminology, and assume that the situation (loosely construed) in which a PC is uttered corresponds to its topic situation, and thus the truth of a PC is evaluated with respect to it.³⁸

Following Landau (2011), I assume that the matrix subject of a CRC is a thematic argument which bears the thematic role PSOURCE, and this is what gives rise to the interpretation that the subject is the source of information upon which the embedded claim is based. I make the further claim that this meaning component carries with it the entailment that the PSOURCE is the type of entity that can be perceived through the modality indicated by the verb, and that the PSOURCE subject is a part of the topic situation (which I will signify with: s^*).³⁹

- (98) $\lambda s^*. \lambda e. \exists e. [\text{PSOURCE}(e)] \wedge [e \leq s^*]$

We can immediately see how this accounts for the infelicity of examples like (80), where the use of the DP “Big Bird” as the matrix subject is infelicitous. Though the entity which Jill imagines to be Big Bird (the shadow) is the perceptual source for

³⁸The situation pronoun of Percus (2000) which is found in the matrix clause (s_1) can perhaps be thought of as a technical implementation of the notion of a topic situation.

³⁹The $\leq$ symbol represents a proper part relation.

her claim, the CRC as a whole is evaluated with respect to the topic situation, not the situations of Jill's impressions. The topic situation for the context in (79) is the situation in the actual world involving Jill, the tent, Ben and Ken, etc. Big Bird is not a part of this situation, and thus the entailment does not hold.

This is the additional restriction that was needed: a CRC can only be true of a situation that, at a minimum, contains the PSOURCE. The propositions that then yield the set of worlds in the modal base are the propositions expressing what PGOAL sees/hears/knows *in s**. Speakers may choose to make claims about topic situations that contain the PSOURCE and nothing else, or they may refer to larger topic situations, i.e. the entire scene that includes the PSOURCE. For example, in a slight variation on scenario (87), the speaker can choose a topic situation that includes only Timmy.⁴⁰

- (99) [Timmy and his friends Jim and Kim are in Timmy's kitchen, chatting while Timmy prepares dinner. Timmy begins to cut onions, and as he does so his eyes begin to water. Jim says to Kim:]
Haha, Timmy looks like he's really upset.

Here, if we stick to the early version of the analysis, the sentence should probably come out false. Both Jim and Kim can see that Timmy is cutting onions; the most stereotypical worlds compatible with their visual information would then be worlds in which Timmy's eyes are watering due to onion-cutting, not to being upset. But if the topic situation is restricted to include only Timmy, then (99) will be true, as the only (relevant) visual information that feeds into the modal base is the proposition that Timmy has liquid coming from his eyes, and this is stereotypically indicative of sadness.

However in certain other cases, it is necessary for the topic situation to be larger than a situation containing only the PSOURCE. For example consider the scenario from (61), where Jack is inside the tent eating candy from a low bowl. Ben says to Ken:

- (100) Jill is scowling in disgust because Jack looks (to her) like he is eating out of the garbage.

The topic situation for this utterance cannot be a situation that contains Jack (and the PGOAL) and nothing else. To evaluate this utterance we need to consider many other things as well: the position of the garbage can, the translucent tent wall, the position of Jill with respect to both of those things, etc. If the topic situation were

⁴⁰However I assume that the PGOAL, as an existentially closed argument of the matrix verb, is also necessarily included in the topic situation.

restricted to Jack alone, in the same way that we can restrict the topic situation in (99) to Timmy alone, then we would not have access to enough information to make (100) true. The only information available would be propositions about Jack's appearance, and the movements that he makes with his arms. We might be able to infer that (it looks to Jill like) he is eating, but we would certainly have no reason to believe that he is eating out of the garbage.

So the topic situation that a speaker intends to refer to can be variable in size, so long as the condition that the PSOURCE is contained within it is met. This ensures that the PSOURCE is available to be interpreted as the source of information. The final version of the truth conditions says that a CRC is true of a topic situation iff the individual specified by the matrix subject is a part of s^* , and in all the worlds compatible with what PGOAL can see/hear/know in s^* (that are ranked most stereotypical by the ordering source), the embedded proposition is true. This is partially formalized in (101).⁴¹

$$(101) \quad \lambda s^*. \lambda e. \exists e. [\text{PSOURCE}(e) \wedge e \leq s^*] \ \& \ [\lambda p. \lambda x. \forall s' : s^* R_x^V s' \rightarrow p(s') = 1.]$$

The modal accessibility relation $s^* R_x^V s'$ can be defined as (102), where x stands for the PGOAL and V represents the lexical content of the verb, whether it be vision, audition, etc.

$$(102) \quad s^* R_x^V s' \text{ iff } s' \text{ is compatible with what } x \text{ sees/hears/knows in } s^*.$$

So just as the PGOAL plays the role of partially restricting the worlds (or rather situations) in which the *embedded proposition* is evaluated, as we saw in Section 4.2.2, it's partner PSOURCE has the role of restricting the situation in which the *entire assertion* is evaluated, though the restriction is accomplished through different means.

In the cases where the matrix subject of a CRC is non-thematic, it does not bear the PSOURCE theta-role, and thus there is nothing that would trigger the requirement that the subject be present in the topic situation. However, why there should be this distinction between thematic and non-thematic subjects at all is a perplexing question. I will provide some speculation on this issue in the following chapter.

⁴¹Formalization of the ordering source component is omitted, however this notion should be made more precise through formalization in future work.

4.3 Summary

In this chapter, I have used the framework of situation semantics in order to account for the (in)felicity of PCs in a wide range of scenarios, and to provide suitable truth conditions for them. I have argued that each copy raising verb represents a different accessibility relation between situations, and this is able to account for two puzzles relating to CRCs in false-truth scenarios: why speakers are sometimes able to express impressions they themselves do not hold, and why they cannot do this with *seem*. The purely sensory accessibility relations of perceptual verbs like *look* and *sound* allow speakers to quantify over situations that are merely consistent with their sensory information, and thus do not necessarily require that the speaker believes the embedded proposition to be true. The epistemic accessibility relation associated with *seem* ensures that speakers quantify over situations which are consistent with their knowledge, and thus any propositions that the speaker believes to be false are excluded. This modal-style analysis of PCs can account for the (in)felicity of both false-truth scenarios and ordinary ones, using exactly the same system.

For thematic-subject CRCs, I have incorporated the PSOURCE interpretation into the truth conditions as an entailment which requires the matrix subject to be a part of whatever topic situation the utterance is evaluated with respect to. All of this demonstrates that we can think of PSOURCE and PGOAL not just as abstract semantic roles (Asudeh & Toivonen, 2012), but as crucial reference points which establish domain of quantification for the matrix and embedded propositions, respectively.

Chapter 5

Modality, evidentiality, and speaker certainty

I have proposed an analysis of perceptual constructions that essentially assimilates the semantics of copy raising verbs to the semantics of modal verbs: they are quantifiers over possible worlds/situations, and the domain of quantification is restricted by means of a modal base and ordering source. In the case of *seem*, I have argued that the lexically supplied accessibility relation is an epistemic one; which means that *seem*, at least in its use in PCs, is essentially an epistemic modal.

There are a handful of recent proposals in the semantic literature which claim that epistemic modals can be thought of as evidentials (von Stechow & Gillies, 2007), and vice versa (Izvorski, 1997, Matthewson et al., 2007, McCready & Ogata, 2007), and some have even claimed that evidentiality *is* epistemic modality (Matthewson, 2014). It is perhaps not surprising then, in light of my analysis, that some authors have claimed that CRCs are evidentials (Rett et al., 2013). Yet many others maintain that the phenomena of epistemic modality and evidentiality are fundamentally distinct, though they may often overlap (De Haan, 1999, Lazard, 2001, Aikhenvald, 2004, Cornillie, 2009). Among the proponents of the latter view, it has also sometimes been argued that epistemic modality, but not evidentiality, encodes the level of speaker certainty (De Haan, 1999). In this chapter I will weigh in on this debate insofar as it pertains to PCs, and suggest that the close relationship between these three categories can provide insight into the peculiar cases of non-thematic subject CRCs.

5.1 Are PCs evidential?

The study of evidential phenomena is a relatively young enterprise, only gaining considerable attention in the mid-1980s. As such, it is still not very well understood, and there is little consensus within the literature on what exactly it means for

something to be *evidential*. At the very least, evidentiality necessarily involves marking of the source of information for a statement. In its strictest interpretation, the term evidentiality only applies if the meaning in question is grammatically encoded and its marking is obligatory (Aikhenvald, 2004). But the marking of evidentiality can be accomplished in various ways, and it is often hard to know where to draw the line between the grammatical category of evidentiality, and the simple linguistic expression of evidential concepts. The common view of English is that it is not an evidential language, because we are not required to mark the source of our information in the same way that we are required to mark tense on all of our utterances. Hence the claim that evidentiality may be present in specific constructions in English must be investigated with caution, and with careful consideration as to the implications this would have for our understanding of evidentiality more generally. With caution then, I now approach the issue, to which I will also return briefly in Section 5.3.

5.1.1 P-SOURCE as evidential

Asudeh & Toivonen (2012) hint that their PSOURCE may be related to evidential phenomena, and Rett et al. (2013), building off of the work of Asudeh & Toivonen, explicitly claim that CRCs encode *direct* evidentiality, while their expletive variants are unmarked for evidentiality. Their claim that direct evidentiality is encoded is motivated by the fact that the matrix subject of a CRC must be directly perceived in order for the utterance to be felicitous. This is problematic for two reasons; one of them empirical, and the other typological. The first is that this claim fails to account for the examples raised in Landau 2011, where the matrix subject clearly cannot be interpreted as the perceptual source. These are the cases that Landau argues involve a non-thematic subject (e.g. “The smell looks like it’s delicious”). And in fact there is no indication that either Asudeh & Toivonen or Rett et al. are aware of these cases.

The second issue is that the claim that there is a requirement of direct perception specifically on the subject is quite unorthodox with respect to the existing literature on evidentials. As far as I am aware, all evidential languages, no matter what grammatical means are used to encode evidentiality, mark the source of information that supports the *entire proposition expressed*. The various types of information that are encoded may differ between languages - some will distinguish between firsthand and non-firsthand evidence only, some will distinguish between visual, other sensory, and reported evidence, etc. (Aikhenvald, 2004) - but the types of evidence that languages encode generally seem to fall into predictable categories. If, as Rett et al. argue, CRCs encode evidentiality, then they encode some hitherto unattested evidential category which amounts to something like “direct perception of sentential subject.” To my knowledge, there are no purported evidentials which

make explicit reference to the grammatical structure. Alternatively, if CRCs merely encode “direct perceptual” evidentiality, then it is a similarly unorthodox kind which applies not to the whole proposition expressed, but to a particular constituent, i.e. the matrix subject. This raises suspicion about whether such a thing really constitutes evidentiality per se, or whether it might be indicative of something more like a theta-role, which is precisely what Landau (2011) argues.

As I follow Landau in treating PSOURCE as a theta-role, under my analysis the evidential flavour of CRCs arises as a consequence of fact that any perceptual event logically involves a perceptual source. As discussed in Chapter 4, I assume that this means that the individual marked as the PSOURCE is entailed to be present in the topic situation, i.e. the situation that corresponds to the perceptual event denoted by the matrix verb. The evidential interpretation is thus achieved indirectly via this entailment. The question of whether or not such incidental types of evidential meaning should be included under that heading of evidentiality can be debated, however opening up the definition to include such cases would leave the term with little interesting content. I will henceforth only use the term *evidential* to refer only to linguistic items whose primary function is to mark the source of evidence for the entire assertion.

5.1.2 The verb as evidential

While I consider it unlikely that the PSOURCE criterion is an evidential phenomenon, there is another possibility which I do not believe has yet been considered, and that is that the copy raising verb itself acts as a kind of evidential marker. Essentially, this would amount to the verb marking the type of sensory evidence that the speaker has to support the embedded claim. This would seem to be, at the very least, consistent with the examples that have been discussed in this thesis, both for CRCs and their expletive variants. It would also be far more consistent with evidentiality as we know it; different sensory sources are frequently encoded in evidential languages, and evidential markers often evolve from verbs of perception (Aikhenvald, 2004).

However this is where things get hairy; I have argued that copy raising verbs each represent a different accessibility relation in a modal semantic system, and thus that they determine what type of information is fed into the modal base. The verb *sounds* will select for a modal base that contains only those worlds that are consistent with the speaker/PGOAL’s auditory information, and so the interpretation that the speaker has auditory evidence for their claim is inevitable. Does this mean that modality and evidentiality are one and the same? Possibly, though not necessarily.

There may be no reliable tests which can tell us what is or is not an evidential, mainly due to the lack of consensus as to what exactly an evidential is. But there are a number of ways of diagnosing epistemic modality. In the next section, I will apply a

pair of tests from Matthewson et al. 2007 (some of which are taken from Faller 2002), which the authors use to argue that the evidential markers found in St'at'imcets are also epistemic modals. I apply these tests to both thematic and non-thematic CRCs, and show that each type behaves differently, suggesting that only the thematic cases involve modality. Implications will be discussed in Section 5.3.

5.2 Tests for epistemic modality

I will be highlighting two relevant tests from Matthewson et al. 2007: Challengeability (aka Assent/Dissent), and Infelicity if the proposition is known to be false. These tests were intended to distinguish between an epistemic modal analysis and illocutionary operator analysis for St'at'imcets evidentials; thus they provide support for an epistemic modal analysis (when passed) but do not exclusively indicate that something is an epistemic modal. I use them here to further validate my modal analysis, and to reveal crucial differences between CRCs with thematic and non-thematic subjects.

Following Landau's (2011) syntactic account of CRCs, we know that the majority of CRCs are indistinguishable between having a thematic or non-thematic subject. As non-thematic subjects are licensed by predication and thus require a copy in the embedded clause, the only time we can be certain that a subject is thematic is when the embedded clause contains no copy pronoun. And as thematic subjects bear the theta-role PSOURCE, the only time we can be certain that a subject is non-thematic is when it is non-perceivable in the situation (i.e. not a possible PSOURCE). I will use these cases in the tests wherever possible.

5.2.1 Challengeability

This test is designed to determine whether or not a linguistic item contributes to the truth of the proposition expressed, and is originally due to Faller 2002. If "the meaning of the element in question can be questioned, doubted, rejected or (dis)agreed with" (p. 110), then it contributes to the truth conditions of the proposition expressed. There is some debate about whether epistemic modals are challengeable, however my own intuitions suggest very strongly that they are, and many other authors share the view that they can be challenged at least sometimes (Faller, 2002, von Stechow & Gillies, 2007, Matthewson et al., 2007). The following example from Matthewson et al. (2007) illustrates this.

- (103) A: Jo must be the thief.
 B: That's not true. There are some other plausible suspects. Jo may be

entirely innocent.
 (Matthewson et al. 2007: ex. (47))

It is not the truth of the embedded proposition that B is challenging here, but the modal claim. Challenging a modal claim usually involves challenging the premises that restrict the modal base or ordering source (Faller, 2002), and this is what B does in (103): he adds the information that there are other plausible suspects to the modal base, meaning that not *all* accessible worlds are ones in which Jo is the thief.

We have already seen an example similar to this with CRCs in Chapter 1, where two observers argue about what sort of day the man on a nearby bench is having (example (52)). In that example, the disagreement cannot be about the fact of the matter; neither of them have or will ever know the facts about this stranger. They are arguing about the modal claim, about the propositions that make up the ordering source. This example certainly feels thematic, but to be certain we must look at copyless examples:

- (104) [Upon watching John emerge from the exam room with an unhappy expression on his face, your friend says “John looks like the exam was difficult.” You reply:]
- a. That’s not true, that’s the face he makes when he’s pleased with himself.
 - b. ?That’s not true, I know for a fact that the exam was easy.

If (b.) is acceptable for some speakers, at the very least it seems uncooperative or pedantic to challenge the embedded claim, because this was not the main point of the assertion. Challenging with (a.) is more fitting, where it is the modal claim that is contested. The speakers negotiate the set of propositions in the ordering source, and not the embedded proposition itself. This indicates that the modal claim does contribute to the truth conditions of the utterance, and in this way, thematic-subject CRCs pattern with epistemic modals.⁴²

As to non-thematic CRCs, consider another one of Landau’s (2011) examples, repeated in (105):

- (105) Mark Twain seems like he was a good man.
 (Landau 2011: ex. (59-a.))

Mark Twain died long ago, and thus he is not available as a potential PSOURCE, so the subject must be non-thematic. Challenging the modal claim here is much more

⁴²This is not to say that the embedded proposition does not contribute to the overall truth conditions. It can be challenged, so it does contribute, however doing so is pragmatically odd, indicating that it is not the central focus of the assertion. I will revisit this point in Section 5.3

odd than in (104).

- (106) a. #That's not true, you don't know anything about Mark Twain!
 b. That's not true, everyone knows he was a jerk.

The same pattern is observed when "John looks like he has failed the exam" is uttered in the scenario where the speaker is looking at the grade sheet, and so the subject *John* is non-thematic:

- (107) a. #That's not true, you don't have your glasses on!
 b. That's not true, I know for a fact that he passed. .

If challenging the matrix claim of a non-thematic CRC is at all possible (the (a.) sentences), it comes across as highly unexpected and strange, while challenging the embedded content in (the (b.) sentences) is much more natural. In (107), the (a.) sentence attempts to challenge the premises upon which the speaker's claim is based, and such a challenge appears to be unwarranted in this scenario. With respect to this test as well, thematic CRCs pattern with epistemic modals, while non-thematic CRCs display different behaviour.

5.2.2 Infelicity if the proposition is known to be false

The second test that Matthewson et al. use as a diagnostic of epistemic modality is that the utterance should be infelicitous if the embedded proposition is known to be false. The logic of this test is that the modal base selected by an epistemic modal consists entirely of worlds that are consistent with the speaker's knowledge, and thus a proposition that the speaker believes to be false will not be true in any of those worlds. This test is clearly related to the "false-truth" scenarios that I discussed in Chapter 4. Recall that the verb *seem* is infelicitous in such cases, whereas *look* and *sound* (and presumably *smell* and *feel* as well) are potentially felicitous. Given that I have argued that *seem* acts as an epistemic modal, its infelicity in this test is predicted, however the behaviour of other perceptual verbs warrants some discussion.

The traditional view of epistemic modals is that their quantificational force corresponds to the speaker's degree of certainty that the proposition expressed is true; universal quantification (e.g. English *must*) over epistemically-accessible worlds indicates very strong certainty, whereas existential quantification (e.g. English *might*) expresses only weak certainty, or possibility. I have argued that perceptual verbs like *look* and *sound* involve modal semantics, though they are not fully epistemic. As with epistemic modals, establishing the worlds in the modal base of a perceptual verb requires orientation to a certain individual (unlike deontic or circumstantial modals),

however the modal base does not have access to all of that individual's knowledge, but rather a more restricted picture of the individual's sensory information. After stereotypical ordering, the modal universally quantifies over the accessible worlds, but crucially in this case universal quantification is not tied to speaker certainty. Thus the fact that the perceptual verbs can be used felicitously when the proposition is known to be false (as in false-truth scenarios like (87)) does not mean that they are not *modal*, it merely means that they are not *epistemic* modals. The nature of their accessibility relations allows for a larger set of worlds to be quantified over, including worlds in which things that they believe to be false are true.

Interestingly however, non-thematic subject CRCs seem to fail this test no matter which verb is used. Using the non-thematic version of "John looks like he has failed the exam" from above, we can construct a scenario in which the speaker believes the proposition to be false:

- (108) [Professor Plum has asked his secretary to compile a list of all his student's exam grades and print it out for him. When he goes to pick it up from her, she tells him that she's printed it, but that she just noticed a mistake: she entered John's grade as 7 instead of 75. He tells her that he needs the sheet right away so he will take it as is, and remember the error. John is a very arrogant student and Professor Plum doesn't like him very much. When he looks at the gradesheet later, he smugly thinks to himself:]
- a. ?Heh, John looks like he has failed the exam.
 - b. Heh, It looks like John has failed the exam.

The difference is subtle, but (a.) seems strange if we understand that Professor Plum knows that John actually passed the exam, and that he cannot see John himself. The same subtle difference can be observed in a scenario using *sounds*:

- (109) [You and your friend are sitting with Dan, who is playing the guitar. You know that the particular song that Dan is plucking does not feature any notes on the B string, and you jokingly remark to your friend:]
- a. ?The B-string sounds like it's missing.

Unlike the related scenario in (23) from Chapter 1, here the speaker can clearly see that the guitar has all its strings, and knows that the song just so happens to lack any notes played on the B-string. And unlike in (23), this sentence is now odd. Non-thematic subject CRCs seem to always be infelicitous when the embedded proposition is known to be false.

However, just as failing this test does not necessarily indicate that something is *not* modal, passing it does not necessarily mean that something *is* either. There are other types of constructions that are infelicitous if the embedded proposition is known to be false: factives, first-person verbs of attitude, etc. Ordinary unmodalized assertions are also of course infelicitous if they are known to be false; they are called lies. Thus we cannot conclude on the basis of this test that non-thematic subject CRCs are epistemic modals. What is curious however, and what can be taken from this section, is that the two types of CRCs differ in this respect. In the next section, I will consolidate the results of the tests and propose a possible explanation for the behaviour of these two types of CRCs.

5.3 Speaker certainty and semantic shift

The tests performed in Section 5.2 provide additional support for the proposed analysis of thematic-subject CRCs as involving modal semantics. However the non-thematic subject CRCs behave curiously, indicating that they may be distinct from thematic CRCs in their semantic as well as syntactic properties.

There is no clear way of accounting for the behaviour of non-thematic CRCs under the present modal system. Using the verb *looks*, for example, should mean that in all (stereotypical) worlds compatible with what PGOAL can see, the embedded proposition is true. As we saw above, this alone should not commit the speaker to any level of belief in the embedded proposition, and thus false-truth utterances should be possible. But non-thematic CRCs were shown to be infelicitous when the embedded proposition is known to be false, indicating that they *do* commit the speaker to the embedded proposition to some degree. The differences between thematic and non-thematic subject CRCs thus extend beyond the status of the matrix subject.

I wish to argue that these differences are a symptom of the fact that the primary communicative function of CRCs has undergone, or is in the process of undergoing, reanalysis from modal (or evidential-modal) to evidential. And I believe that the secondary implication of speaker certainty that is often associated with modal claims may have mediated this shift. This is because the majority of instances of CRCs in everyday usage are likely not the false-truth kind seen in (87); but rather the genuine kind: the speaker is trying to express what is most likely to be the case based on the information that is available to them. And in these instances, the speaker would in fact be under the impression that the embedded proposition they are expressing is the most reasonable one given the situation; essentially, they would appear committed to belief in *p*. It is likely that over time, speakers may have come to understand this to be the communicative function of the utterance, i.e. that the utterer of a CRC wishes to express that *p* is most likely true.

I suggest that this frequent overlap in meaning has led speakers to reinterpret the

main assertion of CRC: where previously the modal claim served as the main assertion (as is the case with thematic CRCs), the embedded clause is now sometimes being used as the main assertion, resulting in non-thematic CRCs. Yet the matrix verb still must have a role to play, and so it is reanalyzed as a (basically parenthetical) comment indicating that the speaker's assertion of *most likely, p* is supported by evidence. In these non-thematic CRCs then, the matrix verb no longer acts as a modal operator, but rather as a kind of evidential marker which specifies the nature of the evidence that supports the assertion.

One of the most commonly-cited traits of evidential markers is that they are *not-at-issue* (Faller, 2002, Matthewson et al., 2007, Murray, 2009), i.e. they are not the main point of the utterance, and as such they are usually not negotiable. This is precisely what the challengeability test in Section 5.2.2 showed us: with non-thematic CRCs, the premises supporting the embedded claim cannot be contested, though claim itself can. I argue that in non-thematic CRCs, the matrix verb semantically acts as a not-at-issue evidential marker which applies to the embedded proposition. The speaker asserts *most likely, p*, and merely adds the information that *p* is supported by evidence.

This type of evidentially-marked assertion is very similar to what Faller (2002) finds for the “direct” evidential *-mi* from Quechua. Faller analyzes Quechua evidentials as illocutionary operators, which modify the sincerity conditions of a speech act. The marker *-mi* is analyzed as encoding that the speaker has the *best possible grounds* for his or her assertion. Faller provides the following example, and indicates the propositional and evidential meanings.

- (110) Irqi-kuna chakra-pi-n puklla-sha-n-ku.
 child-pl field-loc-**mi** play-prog-3-pl
 p= The children are playing in the field.
 ev= speaker sees the children playing in the field
 (Faller 2002: ex. (69))

She explains:

“It seems odd to say that the speaker considers the proposition *The children are playing in the field* an epistemic necessity or possibility because they saw them play there. Rather, the speaker of (69) is presenting this proposition as an unmodalized fact, and at the same time indicates that they can present it as a fact because they have direct evidence” (p. 92).

Something similar seems to be in place for non-thematic CRCs; though the speaker does not need to have the best possible grounds for their assertion, but rather they

need to believe that their assertion is likely to be true. Like the direct evidential *-mi*, the evidential verb of a non-thematic CRC acts as an explicit justification of the assertion.

The reasoning presented here would also be in line with Landau's (2011) account of how the verbs in each of the types of CRCs compose with their arguments. When a copy raising verb takes a thematic subject, Landau argues that it composes directly with the PSOURCE subject and a propositional complement. However when the matrix subject is non-thematic, this subject must be semantically licensed in some other way: Landau proposes that the verb's complement in these cases forms a predicate, which combines with the matrix subject to yield a proposition. It is this resulting proposition that then combines with the matrix verb. For the non-thematic CRC "John looks like he has failed the exam" this might give us something roughly like the following at LF:

(111) [looks] [John has failed the exam]

If indeed the embedded claim forms a complete proposition before it composes with the matrix verb, there are a number of conceivable ways that the verb's evidential meaning could then be applied to it, in the same way that attested evidential markers apply to entire propositions. It may be that the alternative licensing strategy that Landau argues must exist (i.e. predication via Op-merger) developed in response to the shift toward a primarily evidential, rather than modal meaning. The specific mechanism by which the evidential verb then combines with the proposition remains an open question. Copy raising verbs could act as illocutionary operators, similar to those argued for by Faller (2002), or perhaps they carry presuppositions (Izvorski, 1997), or conventional implicatures in the sense of Potts 2005. They may even be related to the "spatio-temporal operators" found in Korean (Chung, 2007). Unfortunately a full investigation of this phenomenon is beyond the scope of this thesis.

5.4 Summary

The meanings of thematic and non-thematic CRCs often appear to amount to the same thing: thematic CRCs claim that based on a certain body of evidence, *p* is the most likely case, while non-thematic CRCs claim that *p* is most likely the case, based on evidence. Yet I argue that it is precisely this difference in the directionality of meaning that makes the two distinct. Thematic CRCs are modal assertions: they begin by establishing a body of facts, and then project a conclusion from those facts. But non-thematic CRCs are assertions modified by an evidential marker: they state

the conclusion that has been reached, adding that this conclusion is based on a certain kind of evidence. This explains the differing semantic properties that we have seen for each type of CRC; the main assertion of a thematic CRC is the entire matrix clause, while the main assertion of a non-thematic CRC is the embedded clause. As a result, the source of evidence denoted by the matrix verb of non-thematic CRC cannot be challenged; it is not at-issue. Furthermore, the utterer of a non-thematic CRC cannot believe the embedded proposition to be false, as the main function of their utterance is to assert that it is most likely the case. I propose that this contrast in the overall point of the assertion is the essential difference between thematic and non-thematic CRCs.

I have argued that thematic CRCs are primarily modal in nature, however in theory this does not exclude the possibility that they are also evidential. Those who support the view that evidentials can be modals, such as Matthewson et al. (2007), might argue that this is simply one such case. The matrix verb that restricts the modal base can be viewed as both a modal operator and an evidential marker, as both functions perform essentially the same job here. However the stronger claim, the claim that evidentiality *is* epistemic modality (Matthewson, 2014), does not appear to be tenable; the existence of non-thematic CRCs demonstrates that one can exist without the other.

Chapter 6

Conclusions and remaining questions

In this thesis, I have attempted to reveal the core semantic properties of perceptual constructions by testing the limits of their acceptability. Beginning with the often-overlooked role of the individual whose perceptual experience is expressed in the sentence, the PGOAL, I showed that the process which identifies its referent cannot be reduced to pragmatic phenomena such as entailment, presupposition, or conventional implicature. Instead, I argued that the PGOAL has a place in the syntactic representation, as a null pronominal argument of the matrix verb. After examining the various contexts which do and do not allow an individual other than the speaker to fill this role, I concluded that the PGOAL's referent is recovered through logophoricity: it must have an antecedent whose thoughts, perspective, or emotions have been made prominent in the preceding discourse.

Once the role of the PGOAL was recognized, the truth conditional properties of PCs became observable. In Section 3.2.1, I showed that the material inside the embedded clause, but not necessarily the matrix clause, is interpreted with respect to the perspective of the PGOAL. I then argued that a modal semantic analysis, akin to the kind used to account for epistemic modals and verbs of attitude, is well-suited to deal with these constructions, where the PGOAL acts as the individual whose consciousness provides the modal base with a set of accessible worlds. The differences in felicity observed in Section 4.1 between the different copy raising verbs were argued to stem from the types of accessibility relations that the verbs lexically encode; the set of worlds that are accessible from an individual's knowledge state (*seem*) will be much smaller than the set of worlds accessible from an individual's modality-specific perceptual information (all other perceptual verbs).

I then proposed that the (relevant set of) worlds in the modal base are ranked with respect to stereotypicality by an ordering source before universal quantification

applies. The stereotypical/normative nature of this ordering source explains why speakers can “faultlessly” disagree (Lasersohn, 2005) over what something looks or sounds like: different people may have different ideas about what is most typical. This also explains why our truth-value judgements of CRCs can sometimes be fuzzy. Most people would probably be inclined to say that “Timmy looks like he’s really upset” is technically true, but what about “Timmy looks like he’s just had laser eye surgery,” or “Timmy looks like water from the roof has leaked onto his face?” A roof-leakage explanation is theoretically compatible with the visual scene (i.e., Timmy with clear liquid on his face), but surely most of us would hesitate to grant this sentence the status of truth. My analysis is supported by this intuition; there is a distinct sense that our hesitation is due to the fact that water-has-leaked-from-the-roof is not the inference that most people would make in this scenario, i.e. it is not the most normative explanation. If I am right in thinking that it is possible to define formal truth conditions for these kinds of utterances, then these conditions are reliant on on shared, conventional notions of typicality.

Importantly, the underlying semantics for thematic-subject CRCs is preserved throughout; independent of the verb that is used, and whether it is used in false-truth scenarios, honest approximations, third-person PGOALS, and perhaps even their expletive variants.⁴³ The reason that Asudeh & Toivonen’s cook from Chapter 1 cannot be absent when using a CRC is due to a requirement that the topic situation include the entity selected by the PSOURCE theta role. In this way, the PSOURCE plays the part of partially restricting the situation with respect to which the CRC as a whole is evaluated, while the PGOAL partially restricts the situation with respect to which the embedded clause is evaluated. The entirely sensible conclusion that emerges from these facts then, is that the semantics of perceptual events is vitally dependent on establishing both the perceiver and the perceivee.

Finally, I attempted to place these constructions within the context of modality and evidentiality more generally, and argued that the distinction between thematic and non-thematic subject CRCs may be reflective of an ongoing shift from a primarily modal to a primarily evidential meaning. The contrast between the two types of CRCs with respect to their observable properties and proposed compositional behaviour (Landau, 2011) can be summarized in the table below.

These categorical distinctions indicate that we are dealing with two separate constructions, however a crucial question that remains is how exactly these different CRCs are represented syntactically, and how the hearer decides between one or

⁴³It is my suspicion that the *it* in certain “expletive variants” is not an expletive at all, but rather a pronoun ranging over eventualities, as suggested by Rett et al. (2013). The eventuality denoted by *it* could then serve as the PSOURCE, thus assimilating the semantics of thematic CRCs with their expletive variants. Unfortunately it is not clear how one could empirically distinguish between the two hypotheses.

Thematic CRCs	Non-thematic CRCs
subject interpreted as PSOURCE	subject not necessarily interpreted as PSOURCE
subject is semantically licensed by theta role	subject is semantically licensed via predication
subject composes directly with matrix verb	subject composes with embedded predicate, which then composes with matrix verb
may or may not have a copy pronoun in the embedded clause	must have a copy pronoun in the embedded clause
contribution of matrix verb is at-issue	contribution of matrix verb is not-at-issue
do not commit speaker to belief in <i>most likely</i> , <i>p</i> when used with perceptual verbs	commit speaker to belief in <i>most likely</i> , <i>p</i> even when used with perceptual verbs

the other interpretation. The thematic subject may be the head of an additional functional projection optionally selected by the verb, or perhaps there is just a single head which selects for a feature that determines whether the subject is interpreted as the PSOURCE or not. These and many other conceivable explanations are possible, and the appropriate analysis of these facts will need to be pursued further. Until these issues are resolved, we will not be able to provide a complete solution to the puzzle of the absent cook that launched this investigation; we do not know what factors dictate whether a thematic interpretation is necessary, and thus we do not know why Tom *must* be present in the kitchen in order for the CRC to be felicitous.

An additional query raised by this work is to what extent the patterns and properties observed here extend to other similar constructions. For instance, do canonical subject-to-subject raising constructions (e.g. “Tom seems to be cooking”) involve something like a PGOAL? Does the A-movement of the subject have an important effect on the semantic interpretation? Are small clauses (e.g. “Tom seems hungry”) also evaluated with respect to a kind of stereotypical ordering? Do sentences where the verb *seem* takes a *that*-complement (e.g. “It seems that Tom is cooking”) somehow express a higher degree of speaker conviction? All of these are questions which I hope this work will be helpful in addressing in future research.

References

- Aijmer, K. (2009). Seem and evidentiality. *Functions of language*, 16(1), 63–88.
- Aikhenvald, A. Y. (2004). *Evidentiality*. Oxford University Press.
- Asudeh, A. (2002). Richard III. *CLS*, 38, 31–46.
- Asudeh, A., & Toivonen, I. (2012). Copy raising and perception. *Natural Language & Linguistic Theory*, 30(2), 321–380.
- Austin, J. (1950). Truth. In *Proceedings of the aristotelian society*.
- Barwise, J., & Perry, J. (1983). *Situations and Attitudes*. MIT Press.
- Basilico, D. (2003). The topic of small clauses. *Linguistic Inquiry*, 34(1), 1–35.
- Beaver, D. I. (2001). *Presupposition and assertion in dynamic semantics* (Vol. 29). CSLI publications Stanford.
- Camilleri, M., ElSadek, S., & Sadler, L. (2014). Perceptual reports in (dialects of) Arabic. In M. Butt & T. H. King (Eds.), *Proceedings of the LFG14 Conference*. CSLI Publications.
- Chomsky, N. (1981). Lectures on Government and Binding. *Dordrecht: Foris*.
- Chomsky, N. (1982). *Some concepts and consequences of the theory of Government and Binding* (Vol. 6). MIT press.
- Chomsky, N. (1993). *Lectures on Government and Binding: The Pisa lectures* (No. 9). Walter de Gruyter.
- Chomsky, N. (1995). *The minimalist program* (Vol. 1765). Cambridge, MA: MIT Press.
- Chomsky, N. (2001). Beyond explanatory adequacy, volume 20 of MIT Occasional Papers in Linguistics. *Cambridge, MA: MITWPL*.

- Chung, K.-S. (2007). Spatial deictic tense and evidentials in Korean. *Natural Language Semantics*, 15(3), 187–219.
- Clements, G. N. (1975). *Logophors in Ewe* (Tech. Rep.). Mimeo.
- Copley, B. (2009). *The semantics of the future*. Routledge.
- Cornillie, B. (2009). Evidentiality and epistemic modality: On the close relationship between two different categories. *Functions of language*, 16(1), 44–62.
- De Haan, F. (1999). Evidentiality and epistemic modality: Setting boundaries. *Southwest journal of linguistics*, 18(1), 83–101.
- Egan, A., Hawthorne, J., & Weatherson, B. (2005). Epistemic modals in context. *Contextualism in philosophy*, 131–168.
- Faller, M. T. (2002). *Semantics and pragmatics of evidentials in Cuzco Quechua* (Unpublished doctoral dissertation). Stanford University.
- Frascarelli, M. (2007). Subjects, topics and the interpretation of referential pro. *Natural Language & Linguistic Theory*, 25(4), 691–734.
- Giannakidou, A., & Mari, A. (2012). The future of Greek and Italian: An epistemic analysis. In *Proceedings of Sinn und Bedeutung* (Vol. 17, pp. 255–270).
- Grice, H. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Syntax and semantics* (pp. 41–58). Academic Press, New York.
- Groenendijk, J., & Stokhof, M. (1984). On the semantics of questions and the pragmatics of answers. *Varieties of formal semantics*, 3, 143–170.
- Hamblin, C. L. (1973). Questions in Montague English. *Foundations of language*, 41–53.
- Heycock, C. (1994). Layers of predication. In *Outstanding dissertations in linguistics*.
- Hintikka, J. (1969). Semantics for Propositional Attitudes. In L. Linsky (Ed.), *Models for modalities* (pp. 87–112). D. Reidel Publishing Co.
- Holmberg, A. (2005). Is there a little pro? Evidence from Finnish. *Linguistic Inquiry*, 36(4), 533–564.
- Holmberg, A., Nayudu, A., & Sheehan, M. (2009). Three partial null-subject languages: a comparison of Brazilian Portuguese, Finnish and Marathi*. *Studia Linguistica*, 63(1), 59–97.

- Huang, C.-T. J. (1984). On the distribution and reference of empty pronouns. *Linguistic Inquiry*, 531–574.
- Hughes, G. E., & Cresswell, M. J. (1996). *A new introduction to modal logic*. Psychology Press.
- Iatridou, S. (1990). The past, the possible, and the evident. *Linguistic Inquiry*, 123–129.
- Izvorski, R. (1997). The present perfect as an epistemic modal. In *Semantics and linguistic theory* (pp. 222–239).
- Keshet, E. (2005). A matter of taste. *Manuscript, MIT*.
- Keshet, E. (2008). Only the strong: Restricting situation variables. In *Proceedings of SALT* (Vol. 18, pp. 483–495).
- Kim, J.-B. (2014). English copy raising constructions: Argument realization and characterization condition. *Linguistics*, 52(1), 167–203.
- Kratzer, A. (1977). What ‘must’ and ‘can’ must and can mean. *Linguistics and philosophy*, 1(3), 337–355.
- Kratzer, A. (1981). The notional category of modality. *Words, worlds, and contexts*, 38–74.
- Kratzer, A. (1986). Conditionals. *Chicago Linguistics Society*, 22(2), 1–15.
- Kratzer, A. (2011). Situations in natural language semantics. *Stanford Encyclopedia of Philosophy*. Retrieved from <http://plato.stanford.edu/>
- Landau, I. (2001). Control and extraposition: The case of super-equi. *Natural Language & Linguistic Theory*, 19(1), 109–152.
- Landau, I. (2003). Movement out of control. *Linguistic Inquiry*, 34(3), 471–498.
- Landau, I. (2011). Predication vs. aboutness in copy raising. *Natural Language & Linguistic Theory*, 29(3), 779–813.
- Landau, I. (2015). *A two-tiered theory of control* (Vol. 71). MIT Press.
- Lappin, S. (1984). Predication and raising. In *Proceedings of NELS* (Vol. 14, pp. 236–252).
- Lasnik, P. (2005). Context dependence, disagreement, and predicates of personal taste*. *Linguistics and philosophy*, 28(6), 643–686.

- Lasnik, H., & Saito, M. (1999). *Minimalist analysis*. Blackwell Publishing.
- Lazard, G. (2001). On the grammaticalization of evidentiality. *Journal of pragmatics*, 33(3), 359–367.
- Matthewson, L. (2014, March). Semantic uniformity and diversity: Evidence from the Pacific Northwest. In *Jackman Lecture Series, University of Toronto*.
- Matthewson, L. (in press). Evidence type, evidence location, evidence strength. *Evidentials and Modals, CRiSPI, Emerald*.
- Matthewson, L., Davis, H., & Rullmann, H. (2007). Evidentials as epistemic modals: Evidence from St’át’imcets. *Linguistic variation yearbook*, 7(1), 201–254.
- Matushansky, O. (2002). Tipping the scales: The syntax of scalarity in the complement of seem. *Syntax*, 5(3), 219–276.
- McCready, E., & Ogata, N. (2007). Evidentiality, modality and probability. *Linguistics and Philosophy*, 30(2), 147–206.
- Moulton, K. (2013). Raising from the dead. *Linguistic inquiry*, 44(1), 157–167.
- Murray, S. E. (2009). A Hamblin semantics for evidentials. In *Semantics and linguistic theory* (pp. 324–341).
- Pearson, H. (2013). A judge-free semantics for predicates of personal taste. *Journal of Semantics*, 30(1), 103–154.
- Percus, O. (2000). Constraints on some other variables in syntax. *Natural Language Semantics*, 8(3), 173–229.
- Portner, P. (1998). The progressive in modal semantics. *Language*, 760–787.
- Postal, P., & Pullum, G. (1988). Expletive noun phrases in subcategorized positions. *Linguistic Inquiry*, 635–670.
- Potsdam, E., & Runner, J. (2001). Richard returns: Copy raising and its implications. *Proceedings of CLS*, 37, 453–468.
- Potts, C. (2005). *The logic of conventional implicatures*. Oxford University Press Oxford.
- Preminger, O. (2006). *Argument-mapping and extraction* (Unpublished master’s thesis). Tel Aviv University.
- Reinhart, T., & Reuland, E. (1993). Reflexivity. *Linguistic inquiry*, 657–720.

- Rett, J., Hyams, N., & Winans, L. (2013). The effects of syntax on the acquisition of evidentiality. In *Proceedings from the 32nd annual Boston University Conference on Language Development (BUCLD 37)* (pp. 345–357).
- Rogers, A. (1971). Three kinds of physical perception verbs. In *Seventh Regional Meeting of the Chicago Linguistic Society* (pp. 206–222).
- Rogers, A. (1972). Another look at flip perception verbs. In *Eighth Regional Meeting of the Chicago Linguistic Society* (pp. 303–315).
- Rogers, A. (1974). *Physical perception verbs in english: A study in lexical relatedness* (Unpublished doctoral dissertation). University of California, Los Angeles.
- Stalnaker, R. (1999). *Context and Content: Essays on Intentionality in Speech and Thought*. Oxford University Press.
- Stephenson, T. (2007). Judge dependence, epistemic modals, and predicates of personal taste. *Linguistics and Philosophy*, 30(4), 487–525.
- von Fintel, K. (2006). Modality and Language. In D. Borchert (Ed.), *Encyclopedia of Philosophy* (Vol. 2, pp. 20–27). Macmillan Reference.
- von Fintel, K., & Gillies, A. S. (2007). An opinionated guide to epistemic modality. *Oxford studies in epistemology*, 2, 32–62.
- von Fintel, K., & Heim, I. (2002). Lecture notes on intensional semantics. *ms.*, *Massachusetts Institute of Technology*.
- von Fintel, K., & Iatridou, S. (2002). If and when if-clauses can restrict quantifiers. In *Unpublished ms, MIT, contribution to the workshop in linguistics & philosophy at the University of Michigan*.
- von Fintel, K., & Iatridou, S. (2003). Epistemic containment. *Linguistic Inquiry*, 34(2), 173–198.