

# 1 Introduction

## 1.1 WHAT ARE ATTITUDE REPORTS?

---

Humans lead rich mental lives, and the languages that we speak afford us rich vocabularies for describing them. A sampling of some of that vocabulary as found in English (sorted into syntactically and semantically relevant groups) is shown in (1).

- (1) a. think, believe, know, conclude, doubt, guess, understand, ...;  
 b. dream, imagine, pretend, fantasize, ...;  
 c. (be) happy, (be) sad, (be) mad, (be) surprised, ...;  
 d. want, wish, hope, like, love, hate, fear, ...;  
 e. intend, plan, decide, aim, try, ...

All of these verbs and adjectives share a syntactic behavior that has attracted a huge amount of attention from linguists and philosophers alike: they can embed sentences, or, in some cases, sentence-like constituents (in particular, infinitives, also known as nonfinite clauses). For example, all of the words in (1-a–c) and some of the words in (1-d–e) can embed the sentence in (2), yielding complex sentences like (3).

- (2) It's raining.  
 (3) Beatrix thinks [it's raining].

Similarly, some of the words in (1-a–b) and all of the words in (1c–e) are able to embed some species or another of nonfinite clause, as illustrated in (4-a–b) for *want* and *intend*, respectively.

- (4) a. Beatrix wants [it to rain].  
 b. Beatrix intends [to buy an umbrella].

An important idea in generative grammar is that form does not always follow function; that is, not all syntactic behavior is explainable by appeal to semantic considerations. But *some* syntactic behavior is

so explainable, and surely it is no accident that many of the verbs and adjectives we use for describing our mental lives have the ability to embed sentences or sentence-like constituents: it is emblematic of what philosophers of mind call INTENTIONALITY,<sup>1</sup> which is the capacity of the mind to represent mind-external objects.<sup>2</sup> Beliefs and desires, for example, have objects, and often those objects are of the sort that we can use sentences to name or describe. Perhaps not all of the mental states and actions described by the words in (1) work like this in every situation. Maybe, for example, I can have ‘undirected’ anger (ultimately, this is a question for psychologists or philosophers, not linguists). But I can also be angry *about something* or angry *that something is the case*.

In many theories of meaning, sentences denote PROPOSITIONS, which we might define, initially, as things that can be true or false. Accordingly, Russell (1940) coined the term PROPOSITIONAL ATTITUDE as a label for what we are talking about when we use sentences built around sentence-embedding psychological verbs like *believe*, *desire*, and *doubt*. In the meantime, it has become commonplace to use the term PROPOSITIONAL ATTITUDE REPORT as a label for the sentences themselves. For the sake of concision, I will often refer to these – as I do in the title of this book – simply as ‘attitude reports.’

Let me now mention a couple of phenomena that – given what’s just been said – one might be surprised to see included in this book, as well as one phenomenon that one might be surprised to see excluded.

First, I consider INDIRECT SPEECH REPORTS like (5) to be within the purview of this book.

(5) Beatrix says [it’s raining].

Although, strictly speaking, we would not want to consider sentences like this to be attitude reports in the narrow sense of naming a

<sup>1</sup> All small-capped terms in this book are listed alphabetically and defined in the Glossary at the end of the book, often with a cross-reference to the section of the book in which they are discussed. In general, I will use small caps for these terms only at their first occurrence in each chapter they appear in.

<sup>2</sup> Intentionality (with a ‘t’) is not to be confused with INTENSIONALITY (with an ‘s’). The latter stands in opposition to EXTENSIONALITY, and has to do with the semantic machinery (often modeled using possible worlds) needed for model-theoretic analysis of expressions involving possibility and necessity, including not only attitude reports but also modal expressions more generally. (More on this in Chapter 2.) To make matters more confusing, both intentionality and intensionality stand in contrast with *intention*, which, just as in ordinary usage, names a particular kind of mental attitude that involves a commitment to perform an action, often expressed in English with the verb *intend*.

## 1.2 *Why a Book about Attitude Reports?*

3

psychological state, indirect speech reports share so many semantic and syntactic properties with attitude reports in the narrow sense that it would be a mistake to ignore them entirely.

Second, I will also devote some discussion to sentences like (6), despite the fact that (at least superficially) they embed neither sentences nor sentence-like constituents; instead, they exemplify what are known as INTENSIONAL TRANSITIVE VERBS.

- (6) a. Beatrix wants [a frisbee].  
 b. Beatrix is looking for [a frisbee].

As we shall see, these sentences also share enough properties with overtly clause-embedding attitude reports that we would not want to exclude them on a superficial syntactic technicality. Not only that, we will see that foundational questions about the status of intensionality in natural language grammar turn in part on the analysis of sentences like (6).

As for what's not covered in this book: with the exception of some extremely brief comments in Section 2.6, I will have nothing to say about so-called PERCEPTUAL REPORTS like (7).

- (7) Beatrix saw/heard/felt [the frisbee fly by].

Perceptual reports are centered around sense verbs like *see*, *hear*, and *feel*, and are syntactically distinguishable from most attitude reports in that they embed a so-called bare or naked infinitive (an infinitive that lacks the infinitival marker *to* ordinarily found in nonfinite clauses in English). In spite of their obvious connection to mental states, there are good reasons for isolating perceptual reports as a class of sentences that are in some ways distinct from attitude reports; see Barwise (1981), Higginbotham (1983) for two relevant classics. That being said, I think that a comprehensive picture of attitude reports will ultimately need to elucidate their similarities and differences with respect to perceptual reports. But that will have to wait for another occasion.

## **1.2 WHY A BOOK ABOUT ATTITUDE REPORTS?**

---

So much for trying to define the object of study. Now on to a perhaps even more pressing question: why read (or for me, write) a book about attitude reports? In a nutshell, the answer is that attitude reports stand as one of the most central topics at the intersection between philosophy and linguistics. On the one hand, attitude reports bear on foundational questions about the nature of sentence meaning and

about the nature of proper names. These questions are typically studied by philosophers of language. On the other hand, attitude reports also interact in intricate ways with a host of semantically relevant grammatical phenomena. This makes attitude reports a very fertile area of study for linguists interested in natural language semantics and its interface with syntax and pragmatics. And yet, despite all this, there exists to date no book-length resource surveying the major findings and open questions and helping one navigate the enormous scholarly literature that attitude reports have inspired (though see Section 1.7 for a list of relevant survey articles and book chapters). While no single book could hope to do full justice to all the dimensions of a topic as rich as attitude reports, we have to start somewhere, and this book is my modest attempt at beginning to fill this gap. Let me now elaborate on some of these themes in more detail.

### 1.2.1 Attitude Reports and Sentence Meaning

One central guiding idea behind formal semantics as ordinarily practiced is that sentences of natural language have meanings that are individuated by and storable in terms of TRUTH CONDITIONS. If I tell you that it is raining, you may not know whether I've spoken truthfully, but as a competent speaker of English, you know that what I've said is in principle either true or false, and you also have some idea of what the sentence's truth or falsity turns on. The predominant approach in formal semantics for modeling this property of sentences is to say that the meaning of a sentence is, on some level, a set of possible worlds, namely those worlds in which the sentence in question is true.<sup>3</sup> This approach is well suited for many natural language phenomena. But it threatens to break down for attitude reports, which seem to be sensitive, at least sometimes, to distinctions that are more finely grained than truth conditions.

To be sure, we already know, quite independently of attitude reports, that there are aspects of natural language meaning that are beyond the reach of truth conditionality. This is, after all, a cornerstone of speech act theory as first developed by Austin (1962), and similar themes continue to be explored today, sometimes under the label

<sup>3</sup> A variant of this approach is to say that the meaning of a sentence is a function from worlds to truth values, namely that function which returns the value true if and only if the world it applies to is one in which the sentence in question is true. The choice between the set approach and the function approach is irrelevant to the discussion here.

## 1.2 Why a Book about Attitude Reports?

5

USE-CONDITIONAL MEANING (see e.g. Gutzmann 2015). To take some extreme examples, it makes little sense to assign truth conditions to expressions like *Hello!* or *Ouch!* But the threat attitude reports pose to truth conditionality is an even more serious one. The reason is that when it comes to attitude reports, truth-conditional approaches seem to get the facts wrong on precisely the kinds of phenomena that they ordinarily excel at: accounting for logical inference patterns. To take one example from the literature (Kamp et al. 2011: 344), the two sentences in (8-a) and (8-b) are truth-conditionally equivalent, although it may take some mathematical sophistication to see this.

- (8) a. There are twice as many women in Bill's class as men.  
 b. Any set containing the number of the men in Bill's class and closed under the operation of forming addition will contain the number of the women in his class. (Kamp et al. 2011: 344)

In spite of this truth-conditional equivalence, it is not difficult to imagine a scenario in which we would be prepared to accept the truth of (9-a) but not prepared to accept the truth of (9-b). This is the so-called PROBLEM OF LOGICAL EQUIVALENCE for attitude reports.

- (9) a. Bill believes that [there are twice as many women in his class as men].  
 b. Bill believes that [any set containing the number of the men in his class and closed under the operation of forming addition will contain the number of the women in his class].

Considerations of this sort have led many scholars to hypothesize that the semantics of an attitude report sensitive in some way to the *form* that the embedded sentence takes, a sensitivity that discriminates even between differences in form that ordinarily do not engender differences in truth conditions. But if that's the case, then we are led to ask: does *any* difference in the form of the embedded sentence change the meaning of an attitude report? Consider another pair of truth-conditionally equivalent sentences, shown in (10). In contrast with the previous case, it is very difficult to imagine a scenario in which we would be prepared to accept the truth of (11-a) but not prepared to accept the truth of (11-b).

- (10) a. Beatrix is chasing Maggie.  
 b. Maggie is being chased by Beatrix.  
 (11) a. Polly believes that [Beatrix is chasing Maggie].  
 b. Polly believes that [Maggie is being chased by Beatrix].

It should be intuitively clear why the two cases differ: recognizing the equivalence in (8) requires mathematical sophistication that Bill may not have, whereas the equivalence in (10) is something that we'd expect any competent speaker of English to acknowledge, even if only tacitly. But it is quite another matter to build a theory that draws a principled grammatical distinction between the two kinds of cases. Suppose no such distinction can be drawn. (Not that we should give up that easily – but suppose it, just so we can explore its consequences for a moment.) In that case, we are left with a choice. At one extreme, we might pursue a theory in which differences in the form of the embedded sentence always lead to differences in the meaning of the report, and explain away the perceived equivalence of (11) as something that is, strictly speaking, not a matter of semantics. At the other extreme, we might pursue a theory in which differences in the form of the embedded sentence that are not ordinarily truth-conditionally consequential never lead to differences in the meaning of the report, and explain away the perceived non-equivalence of (8) on pragmatic grounds: technically speaking, we would say, they are equivalent, but they give rise to different conversational implicatures that cloud this judgment.

In this way, attitude reports constitute a crucial testing ground for theories of sentence meaning. As things currently stand, there are a great many proposals on the market but nothing close to a consensus about which one is right. We will revisit this matter in Chapter 2.

### 1.2.2 Attitude Reports and Proper Names

Another important idea about natural language meaning – and, one could argue, a key explanandum of semantic theory – is that we can use language to refer to language-external objects. Proper names are a central example. I can utter the words *Stephen King* and thereby refer to a particular individual, namely the popular American horror writer who wrote *It* and other bestselling novels. And, according to the predominant view in formal semantics, the meaning of a proper name like *Stephen King* consists *solely* in its capacity to refer to the relevant individual. But if this is right, then two proper names that refer to the same individual should be semantically identical. Together with some other reasonable assumptions, this leads to the expectation that *Stephen King* and *Richard Bachman* (a pen name that Stephen King has occasionally used) should be interchangeable in all contexts without affecting the meaning of the sentence that they appear in. This expectation usually seems to hold up; for example, it seems intuitive that if Stephen King

## 1.2 Why a Book about Attitude Reports?

7

is Richard Bachman, then (12-a) is true if and only if (12-b) is true. But, in what has come to be known as FREGE'S PUZZLE (after Frege 1892), this expectation famously breaks down in attitude reports. Suppose that (13-a) is true and that Beatrix does not realize that Richard Bachman is Stephen King. Then, (13-b) might seem false.

- (12) a. **Richard Bachman** wrote *Thinner*.  
 b. **Stephen King** wrote *Thinner*.
- (13) a. Beatrix thinks that [**Richard Bachman** wrote *Thinner*].  
 b. Beatrix thinks that [**Stephen King** wrote *Thinner*].

If one accepts the thesis that proper names that refer to the same individual are semantically identical (although not everyone does), then Frege's puzzle is a special case of the problem of logical equivalence considered in the previous subsection, and some of the available solutions are similar: we could pursue a theory in which formal differences in the embedded sentence – even down to the difference between two co-referring expressions – lead to differences in the meaning of the report. Or we could pursue a theory in which, despite our intuitions, (13-a) and (13-b) really are equivalent; they only seem otherwise because our intuitions are clouded by pragmatic factors. Yet another kind of approach proceeds by complicating the semantics of attitude reports, in such a way that they are sensitive not just to the proposition encoded by the embedded sentence but also to contextual parameters that can be influenced by the choice of one proper name over another.

As is the case for the problem of logical equivalence, there are many proposals on the market for solving Frege's puzzle but no consensus about which one is right. This will be the focus of Chapter 3.

### 1.2.3 Attitude Reports and Grammar

Aside from bearing on foundational questions about meaning, attitude reports also interact nontrivially with many other independently interesting and important semantic phenomena. Often, they do so in ways that seem to be orthogonal to the foundational questions about sentence meaning and proper names. This is a good thing, because it means that one need not solve the problem of logical equivalence or Frege's puzzle in order to make progress on other puzzles. Some of the phenomena with which attitude reports interact include: scope and intensionality (see Chapter 4 on the DE DICTO/DE RE AMBIGUITY and Chapter 7 on intensional transitive verbs); INDEXICALITY, LOGOPHORICITY, and CONTROL (see Chapter 5 on DE SE ATTITUDE REPORTS);

MOOD, MODALITY, GRADABILITY, FOCUS, and PRESUPPOSITION PROJECTION (see Chapter 6 on the semantics of *want*); TENSE (see Chapter 7); and negation (see Chapter 7). These interactions, I would argue, constitute the primary attraction that attitude reports have to offer to linguists interested in semantics and the syntax–semantics interface.

### 1.3 THE APPROACH

---

Let me say something about how this book is organized and about the theoretical framework that it employs. As far as organization goes, the book takes a ‘puzzle-driven’ approach: a typical chapter begins by illustrating some phenomenon related to attitude reports that poses some puzzle or question for semantic theory. This then leads to a discussion of solutions that have been proposed in response and their theoretical implications. It also leads to further puzzles prompted by these solutions, which in turn spur refinements and alternative solutions. In recognition of the reality of the field, the point is never to come down firmly on any particular solution, but instead to illustrate as clearly as possible what is at stake in the choice between the various solutions. In this way, I hope to convey some of the richness of attitude reports as a topic of investigation.

Important work on attitude reports has been carried out within a number of different theoretical frameworks and intellectual traditions, and this of course poses a challenge for a book aiming to synthesize the important findings of these disparate sources. To the extent possible, I will cleave to the framework with which I am most familiar and in which much of the relevant work reviewed here has been carried out, namely that introduced and summarized by Heim and Kratzer (1998). More specifically, I assume that natural language grammar has two components relevant to the study of attitude reports: a generative (syntactic) component that assembles structures out of units drawn from a lexicon, and an interpretive (semantic) component that assigns denotations to those structures. I take the generative component to have a (broadly construed) Principles and Parameters architecture (see e.g. Chomsky and Lasnik 1993), but nothing in the book relies heavily on the details of any particular version of this theory. As for the interpretive component of the grammar, I assume as a working hypothesis that sentence meanings define truth conditions, derived compositionally via a small inventory of type-sensitive compositional rules (including at least Functional Application, Predicate



1.4 *A Note on Readership and Topical Emphasis*

9

Modification, and Predicate Abstraction) that operate locally on sister constituents in the structure. I assume at the outset an ontology of semantic types that includes individuals and truth values as atomic types, and over the course of the book we will entertain a number of other atomic types that may prove useful, namely worlds, times, eventualities, and degrees.

---

**1.4 A NOTE ON READERSHIP AND TOPICAL EMPHASIS**

---

The primary target audience for this book is students or researchers of linguistics who have had at least one or two graduate-level courses in formal semantics. As already mentioned, the text that most closely matches the background theory of this book is Heim and Kratzer 1998, with augmentations developed by von Stechow and Heim 2011 for dealing with intensional phenomena. Some of the material we will be considering requires technology that goes beyond what is introduced in either of these texts, and I will do my best to explain that technology as lucidly as I can without derailing the discussion, with references to other relevant sources where I think that would be helpful. To get the most out of this book, it is therefore recommended that readers be at least somewhat familiar with a Heim and Kratzer 1998-style framework for investigating semantics. Also highly recommended is the first chapter of von Stechow and Heim 2011, which extends Heim and Kratzer's framework, in a very accessible and lucid way, to intensional semantics. Some of the relevant background from both of these works will be briefly covered in Section 2.2 of this book.

One of the challenges in surveying attitude reports is that it is a topic investigated both by philosophers and by linguists, but often with different emphases. As already touched on above, philosophers tend to be more concerned with how attitude reports bear on foundational issues such as the nature of sentence meaning and how proper names refer, whereas linguists tend to be more concerned with how the meanings of attitude reports are grammatically encoded and how they interact with semantically relevant grammatical phenomena such as scope, binding, tense, and presupposition. This book makes no secret about being written by a linguist for linguists, but I nonetheless include some material that is more traditionally in the domain of philosophy, especially some of Chapter 2's discussion of HYPERINTENSIONALITY, as well as pretty much all of Chapter 3 on attitude reports and proper names. These are topics that linguists working on attitude reports ought to know something about, even if they're not going to be

engaging with them directly in their own work. The rest of the book, by contrast, including most of the material after Chapter 3, is more thoroughly grounded in the linguistics literature, even if much of it is ultimately traceable to important philosophical forebears, including especially Jaakko Hintikka, David Kaplan, Saul Kripke, David Lewis, W. V. O. Quine, and Robert Stalnaker.

Let me also say something about why so much of the book emphasizes belief reports as opposed to other kinds of attitude reports such as desire or intention reports. There are a couple of reasons for this. One is that belief reports are historically the variety that is best studied, and therefore the arena in which a lot of the core issues targeted by this book play out. Another is that it so happens that, for many of these core issues that we will be focusing on, the issue works the same way no matter whether we are looking at belief reports or some other kind of attitude report (though I will do my best to flag exceptions to this). But see Chapter 6 for a dedicated look at issues that come up when we turn our attention to other kinds of attitude reports. When we move beyond core issues, there is a great deal of richness to be found in exploring variation between different kinds of attitude reports, and there is still much work to be done in this area.

### 1.5 GUIDE TO LOGICAL SYMBOLS AND NOTATIONAL CONVENTIONS

---

In this section, I provide an informal key to the main logical symbols and related notational conventions that will come up over the course of this book. Let me begin by warning the reader that this guide will probably not be particularly helpful if this is your first exposure to these symbols and the concepts behind them; instead, this is intended as a refresher and quick reference guide. For a thorough, linguistically oriented introduction to these and other concepts from logic, see Partee et al. 1990.

First, from propositional logic, we borrow the concepts and corresponding symbols illustrated and informally defined in (14). Let  $p$  and  $q$  stand for arbitrary propositions – each is either true (1) or false (0) – and ‘iff’ abbreviates ‘if and only if.’ Note also that the *ors* in the definitions of (14-b) and (14-c) are to be understood as *inclusive*: true even if *both* of the disjuncts are true.

- (14) *propositional logic*
- |    |  |                    |
|----|--|--------------------|
| a. | $\neg p = 1$ iff $p = 0$                 | <i>negation</i>    |
| b. | $p \wedge q = 1$ iff $p = 1$ and $q = 1$ | <i>conjunction</i> |