EXPLORING MEANING COMPOSITION: THE EPISTEMOLOGY OF MODELING PRACTICES IN FORMAL SEMANTICS

by

Marcia Andrea Gonzales Llanos

A Thesis Submitted in

Partial Fulfillment of the

Requirements for the Degree of

Master of Arts

in Philosophy

at

The University of Wisconsin-Milwaukee

August 2023

ABSTRACT

EXPLORING MEANING COMPOSITION: THE EPISTEMOLOGY OF MODELING PRACTICES IN FORMAL SEMANTICS

by

Marcia Andrea Gonzales Llanos

The University of Wisconsin-Milwaukee, 2023
Under the Supervision of Professors Joshua T. Spencer and William A. Penn

The present thesis attempts to offer an epistemological account of the modeling practices involved in linguistic model-theoretic truth-conditional compositional semantics. That is, it tries to answer the question of how formal semanticists can draw inferences when building and using scientific models, so knowledge or a deeper understanding of natural language can be gained. The methodological approach followed in this paper is to first examine case studies and then characterize the modeling practices in question by applying relevant insights from the contemporary philosophy of science literature which is focused in scientific practice and scientific modeling. Accordingly, an assessment of the inferences that semanticists can draw within models, between models, and from model to theory or model to data is provided. Finally, the following two claims are defended: that the modeling practices involved in formal semantics are exploratory in nature, and that the construction and application of semantic models license hypothetical inferences which provide modal knowledge about the syntax-semantics interface.

TABLE OF CONTENTS

ACKNOWLEDGMENTS

First and foremost, I want to thank the members of my thesis committee for their patience and unwavering support before and during the writing process, and for always going above and beyond in taking my ideas very seriously. Without them, earlier versions of this work would have never seen the light of day. Whereas both of my co-advisors, Joshua Spencer and Willy Penn, played important roles in helping me uncover and shape my deepest intuitions about language, metaphysics, science, and most importantly, philosophical theorizing, Nick Fleisher's insights and vast knowledge of the semantics literature were crucial for cementing the building blocks of my views and arguments. At the risk of increasing the number of pages of this very long document, I still want to thank each of them individually and acknowledge their contributions to my education, as well as my academic and personal development.

Special thanks go to Joshua for being the person who has influenced my life the most during my MA studies journey. Not only is Joshua an enthusiastic cat lover, a great professor, and a very insightful philosopher, but also an amazing human being. I especially appreciate that he has been incredibly generous with his time to discuss my philosophical and academic interests broadly (despite our *many* disagreements), as well as very unrelated matters. In retrospect, I have to say that whatever expectations I had about having an advisor by the time that I arrived at Milwaukee, Joshua did more than meet them by providing much needed guidance and moral support at every step of the way.

Special thanks go to Willy for looking after me since my (and his) very first weeks here, for his friendship, for engaging in many incredibly fun and interesting philosophical conversations, for sharing my complaints about back pain and other "old-age" related problems, for insisting in the idea that philosophy of science would make me happy (he was right!), and for pushing me to take what seemed to be imprudent risks at the time. Moreover, I am greatly indebted to him for stimulating the development of the ideas presented here and for enthusiastically listening to every single one of my rants about how amazing linguistics is.

Special thanks go to Nick for teaching me *everything* that I know about formal semantics, and for being such a great mentor, role model, and an endless source of witty comments and fun historical facts about linguistics. It is no exaggeration to say that I arrived at my first syntax-semantics seminar without knowing a single thing about linguistic semantics. But, as any reader will notice, this thesis is filled with details about some proposals in the semantics literature. If I did a good job at understanding these accounts and explaining them, it is due to the thoroughness and dedication of Nick's teaching. Furthermore, without his gentle disposition to engage in philosophical conversations, it wouldn't have been possible for me to pursue this project at all.

Other people at UWM also deserve special mention. Hamid Ouali has been immensely supportive of my interest in linguistics, and his excitement about his students' questions, ideas and projects never fail to be extremely motivating and contagious. Aidyn Cooper and Julia Pelletier have been the most loving and supportive friends, and the best office-mates. I especially appreciate that they have always tried to make me feel welcome, even before we became friends. What's more, none of my achievements would have been possible if they hadn't been caring and rooting for me for the last two years (in fact, I wouldn't have managed to complete the first full draft of this thesis if Julia hadn't introduced me to Midnights 3am edition). Lin Zheng has been the best work buddy I've ever had and a good friend. I have particularly enjoyed engaging with him in all sorts of never-ending conversations whenever I would stop by to say hi. Steven Alvarez has been the friend who made me laugh the most and who would never stop offering encouragement. Moreover, I have to acknowledge here his impressive success at manifesting into my future some very positive developments. Kenny Galbraith always offered reassurance, great advice and good vibes. Furthermore, he provided invaluable feedback on my incredibly long first completed draft. Juan Murillo Vargas is the first friend I made at UWM, and the person who taught me how to think about linguistic semantics in a philosophical way, and who inspired me to be passionate about language. Tessa Culleton has been the first and best linguistics friend ever. I will always be incredibly grateful to have had the opportunity to share so many mornings, second lunches, coffee/bubble tea shop visits and car rides (as well as laughs and anxieties) with her during my last semester here—I really couldn't have asked for more. I am also profoundly indebted to the other participants of the Fall 2022 Syntax and Semantics seminar—Aisha Fuddah, Daniel Quigley, Reem Alshammeri and Yunhui Kim—for making the class' sessions a welcoming and safe space for me at times of despair.

My Peruvian friends Alessandra Oshiro, Carlos Schoof, Denisse Deza, Erik Alvarado, Jorge Yakushi, Karl Palomino, Mario Sheing, Marlon Rivas, Pía Balmaceda and Ricardo Rojas, have contributed in one way or another to my mental sanity these past two years. Some adorable pets also did, so shout-out to Bianca, Bowie, Callie, Chispa, Lana, Nata, Pinot and Wesley for being so cute. And it goes without saying that my family excelled at being supportive throughout this journey. So, thanks a lot, everyone!

This book stems from a belief that linguistic semantics is a beautiful field, that the tools used to study formal semantics have yielded a rich body of results about fascinating and subtle data, that the field continues to produce exciting new insights at an impressive rate, and that there are simple and elegant tools to model how the syntax and semantics of a natural language work together. (Jacobson, 2014, p. 3)

Chapter 1

Introduction

Linguistic semantics is the field which studies natural language meaning, with formal semantics being the current dominant approach. This approach is usually characterized as the attempt to apply formal devices from mathematics, logic and computer science to the analysis of meaning composition. Hence, formal semanticists offer proposals of how the meanings of complex expressions like sentences can be determined by the meanings of the sentence's parts and the way in which these parts are combined by the syntax.

The present thesis will be focused on what is known as *compositional model-theoretic truth-conditional semantics*. That is, it will be concerned with linguistic theorizing which aims to provide a systematic account of how the truth-conditions of natural language sentences can be compositionally derived, by mainly drawing on resources from mathematical model-theory. But why does this form of theorizing aim to pair sentences with truth-conditions if semantics is concerned with *meanings*? The intuition behind this methodology is captured by the familiar slogan "to know the meaning of a sentence is to know its truth conditions" (Heim and Kratzer, 1998, p. 1). The idea here is that even if the notion of meaning hasn't been properly defined yet, semanticists can still theorize about it in virtue of the fact that semantic competence involves knowing what it would take for a sentence to be true—*i.e.*, knowing the sentence's truth-conditions. Thus, truth-conditions become a proxy for meaning, under this approach.

This methodology for the study of natural language meaning has led to many recent developments in the field, which have certainly not gone unnoticed by philosophers working in philosophy of language and its intersections. Still, philosophers haven't had much to say about the way that this theorizing is being carried out[1]. If the findings of linguists are being routinely used as evidence for philosophical arguments, shouldn't philosophers worry about the epistemic warrant that semantic claims can provide? Moreover, regardless of whether linguistic theorizing is instrumentally valuable for philosophers of language, paying attention to the practices of semanticists could be beneficial both for semanticists themselves and for philosophers of science. Philosophical work which aims to describe, contextualize and explain the methodology and aims of semantic theorizing could help semanticists to gain a better understanding of their own work and findings, and of their discipline as a whole. For philosophers of science, being aware of the practices involved in disciplines that go beyond traditional scientific fields would certainly provide them with a more informed understanding of scientific practice and its epistemology.

The present project hopes to contribute to the filling of this lacuna in the literature by addressing issues which are mainly epistemological and methodological in nature. For that reason, some tentative answers will be offered to the following questions (but not necessarily in that order):

(1) What is the object of study of formal semantics?

(2) What are the roles that the tools used in formal semantics play in the study of its subject-matter?

(3) How do semanticists derive knowledge from their theorizing practices?

The method that will be followed here is to study the theorizing practices of semanticists through a philosophy of science lens. Within philosophy, and to some extent, linguistics, the tools and products involved in the theorizing done by semanticists (*e.g.*, semantic frameworks and derivations of meaning composition) are referred to as "theories". It will be assumed here, without offering any explicit arguments, that it is better to think of them as scientific models. To this end, the examples that will be considered in Chapter 4 should elicit the intuition that most of the practices of semanticists involve modeling and not theory construction.

Accordingly, an assumption that will underlie the present approach is that thinking of semantic models[2] as *scientific* models can be a fruitful enterprise for gaining a deeper understanding of linguistic theorizing. Scientific models will be broadly conceived of here as the tools which are the

[1] There are some exceptions, of course—for a few examples, I refer the reader to the collections of articles edited by Burgess and Sherman (2014) and Ball and Rabern (2018).

[2] Throughout the paper, the term "(linguistic/semantic) models" will be used to talk about models in the scientific sense to be detailed below. This should not be confused with the use associated with model theory.

product of the model-building practices of scientists[3]. These tools have a variety of uses, which include (non-exhaustively) representing, explaining, intervening on, or learning about some relevant aspect of the world—that is, about the models' "target systems" (*cf*. Morgan & Morrison, 1999). As a consequence, one of the aims of this thesis will be to offer some reasons to start seriously thinking about linguistic semantics as a science, though this claim will not be argued for directly. Hence, what is being suggested here is that we can draw some interesting lessons from the epistemology of the modeling practices of semanticists regardless of whether one believes that linguistics is a cognitive science or a branch of psychology, or that semantic competence can be an object of scientific study.

The proposal to be offered here will be that the modeling practices involved in formal semantics can be thought of as *exploratory* in nature (*cf*. Gelfert, 2016), and that semantic models are reminiscent of the class of models that Massimi (2019) has labeled as *hypothetical perspectival models*. In line with this, it will be suggested that the function of semantic models is to deliver modal knowledge about a hypothetical target system—the syntax-semantics interface. That is, these models serve to reduce the space of objective possibilities in a manner that is constrained both by working assumptions and by linguistic data. If this is correct, their exploratory nature entails that these models work with possible entities, which are not necessarily known to be actual. But this should not be thought of as a disadvantage. Instead, the ability to make hypothetical inferences is what gives semanticists the necessary knowledge to choose which commitments to include as part of their future theories and models.

The structure of the thesis is as follows. Chapter 2 briefly reviews some of the insights from the philosophy of science literature about scientific modeling and scientific practice which will guide the methodology of this project. Chapter 3 offers an overview of some key concepts and methodological commitments that are assumed in the formal semantics literature, such as truth-conditions, semantic values and compositionality. Chapter 4 presents and assesses six case studies of modeling practices. Chapter 5 identifies certain features of these modeling practices, and offers an epistemological account which supports the claim that these practices involve the introduction of exploratory posits meant to drive the development of both future theory and future empirical study. Chapter 6 concludes and highlights some further issues for exploration and future research.

[3] This way of understanding scientific models is meant to stand in opposition to more traditional views which characterize them as interpretations or mathematical structures which satisfy scientific theories. In this paper, a stance on the relation between models and theories which is inspired by Cartwright's instrumentalist view of theories (*cf*. Cartwright et al., 1995; Suárez & Cartwright, 2008), and Morgan and Morrison (1999)'s autonomy approach to scientific modeling will be assumed. More will be said about this issue in Chapter 2.

Chapter 2

Scientific Modeling

Within the philosophy of science literature, it has been acknowledged that important aspects of the epistemology of science cannot be accounted for if we only take into consideration the roles that the "final" products of scientific theorizing play. That is, it has been claimed that in order to explain how many inferences that scientists make in the context of ordinary scientific inquiry are licensed, we need to take a closer look at the actual practices of scientists. A product of applying this methodology is the insight that scientific models play a crucial role at enabling users to draw inferences and learn about the world (amongst other things) by building the models or putting them to use (Morgan & Morrison, 1999, p. 11). As a consequence, some forms of contemporary research are driven by the metaphilosophical assumption that our understanding of science would be better informed by focusing our attention on the practice of scientific modeling (Frigg, 2022, p. 363).

But in what does this practice consist? Weisberg (2007) describes it as follows:

> Modeling, I will argue, is the indirect theoretical investigation of a real world phenomenon using a model. This happens in three stages. In the first stage, a theorist constructs a model. In the second, she analyzes, refines, and further articulates the properties and dynamics of the model. Finally, in the third stage, she assesses the relationship between the model and the world if such an assessment is appropriate. (2007, p. 209).

Certainly, the above characterization is not uncontroversial, but it does serve to provide a rough idea of what sort of practices we will be concerned with for the remainder of this thesis. Furthermore, it must be clarified that it is neither the aim of this project to offer an account of what scientific modeling is, nor to endorse any particular proposal. All that we will need to keep in mind for our purposes here is that scientists construct some tools (models), which are to be manipulated and applied to the study of something other than the tools themselves. This set of activities is what will be meant by "modeling practices".

What *is* uncontroversial, is that modeling involves constructing models. For that reason, model-building is an activity to which we will pay a lot of attention in the next chapters. But what can be tentatively said about it? An informative characterization is offered by Boumans (1999), who suggests a useful analogy between building "first-generation mathematical models" and baking a cake without a recipe. The latter involves some trial and error in order to figure out which ingredients to use, their appropriate proportions, and a way of mixing them together so the expected result can be obtained. Yet, there is always some background knowledge which can be used as an starting point. For example, even if we haven't baked any cakes before, it might happen that we have experience preparing other baked goods, or that we have some vague sense of possible ingredients. This hints at the fact that model-building is an activity that aims at integrating disparate elements from different sources or domains. For this reason, any set of ingredients for these models will require a different recipe, that is, different ways of succeeding at making the elements fit together (1999, p. 67).

Boumans' proposal is inspired by business-cycle models, which can involve "theoretical notions, metaphors, analogies, mathematical concepts and techniques, policy views, stylised facts and empirical data" (1999, p. 94). Accordingly, his account highlights the importance of the choice of mathematical formalism for integrating the disparate elements of the model. The possibilities for shaping the ingredients in a mathematical form will be determined by this choice, which is largely independent of theory or theoretical requirements (1999, p. 94).

2.1 Scientific Models

We have been talking about scientific models for a while now, without saying what they are. This has been done deliberately, since giving a precise characterization of models is still an on-going enterprise within the literature on scientific modeling. Some philosophers have attempted to offer a functional characterization by answering the question of *what roles models perform*, whereas others have tried to provide ontological accounts which are meant to specify *what kind(s) of objects they are*.

Traditionally, it was thought that the key to functionally characterize models lies in their representational role. For example, Hughes claimed that "[t]he characteristic—perhaps the only characteristic—that all theoretical models have in common is that they provide representations of parts of the world, or of the world as we describe it" (1997, S324). Likewise, Giere suggested that our focus should be on the activity of representation, which can be conceived as a relationship of the following form: scientists, as intentional agents, use models to represent some aspect of the world for specific purposes (2004, pp. 743, 747). This is possible since "[w]hat is special about models is that they are designed so that elements of the model can be identified with features of the real world" (2004, p. 747). More recently, it has been noted that some models are not representational in nature, or at least do not perform said role primarily. Some of these other functions include (non-exhaustively) (*cf*. Bokulich, 2011), exploring (*cf*. Gelfert, 2016), or intervening on (*cf*. Keller, 2000) some relevant aspect of the world.

On the more metaphysical side, the efforts to give a unificatory account of models have not been successful either. By taking a look at the history of scientific practice, anyone will notice that some models are material objects—for example, model organisms in biology, ball-and-stick models in chemistry, Watson and Crick's cardboard and metal models of the DNA structure, or fluid models. Yet others, such as the Bohr model of the atom, evolutionary models in biology, and economic models, are constituted by equations, descriptions, set-theoretic structures, abstract entities more generally, or a combination of these. Thus, it seems unlikely that there is a distinctive ontological category to which models belong in virtue of being models.

A different approach focuses instead on giving a characterization which is subsidiary to the functional one by answering the question of what kind of entities models have to be in order to perform their roles. The underlying assumption here is that any ontological account of scientific models must start by recognizing that models are not such in virtue of belonging to a particular ontological category. Instead, models are used by scientists in virtue of having some features which make them suitable for performing the specific roles that scientists intend them to serve. And this suitability could be certainly associated to belonging to a particular ontological kind (Frigg, 2022, pp. 397–407).

2.2 Models and Their Relation to Theories

Historically, the interest in scientific models emerged in the philosophy of science literature as a reaction to the overtly axiomatized conception of scientific theories that "the received view" (also known as "the syntactic view of theories") proposed. This alternative approach, now termed "the semantic view", conceives of a scientific theory as the collection of its models. Models, under this view, are interpretations (*i.e.*, mathematical structures or non-linguistic entities) which make the theory true (Morgan & Morrison, 1999, pp. 2–4; da Costa & French, 2000, S119).

Yet, as Cartwright et al. (1995), and Morgan and Morrison (1999) have pointed out, such an understanding of models is still quite limited: the semantic view takes scientific models to be subsidiaries to theories, but this is not reflective of scientific practice. For instance, Cartwright argues against a "theory-dominated" view of science by offering an alternative instrumentalist conception. According to this "tool-box view", scientific practice consists in putting to use adaptable tools (instruments, techniques, methods) in order to pursue scientific endeavors such as intervening and representing. And if we inquire about which tools are most suitable for these purposes, we will realize that it is not theories which represent the world, but models (1995, pp. 138–139). Furthermore, the joint work of Suárez and Shomar shows via case studies that models which represent phenomena are neither deduced from theories, contained by the latter, nor arrived at by de-idealization—instead, they are purposely built by scientists to represent, with theory only playing a small role in these modeling practices (1995, pp. 140–142). Hence, they concluded that the relationship between models and theories is more distant and complex than what the semantic view originally suggested. Similarly, Morgan and Morrison argued that the fact that models are characterized in terms of their relation to theories gives rise to an inaccurate understanding of modeling practices (1999, pp. 10, 36). For that reason, they suggest that any account of how models provide information about the world must acknowledge the autonomy of models from theory (1999, pp. 7–8). In the next section, I will briefly examine how they implemented such an approach to the study of scientific models.

2.3 Instrumentalism and The Autonomy of Models

For Morgan and Morrison, models "have a life of their own and occupy a unique place in the production of scientific knowledge" (1999, p. 18). Under this account, models are characterized as *autonomous agents* which perform the function of being *instruments* of scientific investigation. Moreover, the autonomy of models is explained by their being partially independent in many respects from both theory and the world (1999, p. 10).

Models are autonomous partly in virtue of the way in which they are constructed. This picture of model-building practices aligns with that of Boumans (1999)—and in fact, is inspired by it. Accordingly, Morgan and Morrison hold that the construction of models is largely independent from theory, even in cases in which most of the elements seem to come from a theory. They acknowledge that models will also involve elements which come from data and/or from outside of the domain of investigation (1999, pp. 13–14). As a result, models are not part of a theory-world axis. That is, they are not part of a hierarchical structure between theory and world (1999, pp. 17–18).

Another important tenet of this view is the claim that models function as tools. As any other tool, they are independent or separate of whatever they operate on, but there must be a connection between the two if the role is to be performed satisfactorily. The same needs to be the case with models if they are tools for learning from both theory and world. That is, models must be somewhat independent from both, but appropriately connected to the two domains. Yet models are not just like any other tool. For Morgan and Morrison, what distinguishes them is their representational capacities. Furthermore, models not only can represent aspect of the worlds, but also of theories. This is how we learn from both domains, and how models are able to mediate between them (1999, p. 11).

Recall that Morgan and Morrison's main motivation for giving an account based on scientific practice was to give an answer to the question of how we learn by using models. So far, we have given an overview of certain features of scientific models that might explain how learning is possible. But how does learning actually occur? For them, learning mainly takes place while the models are being constructed and manipulated (1999, p. 11). Building a model is an activity for which there are no rules, and in which modelers are trying to answer the question of what will fit together and how. As a consequence, one learns both about the world and theory since the practice involves interpreting and conceptualizing both evidence and theories into compatible terms, in addition to integrating them. Moreover, even in cases where a "basic recipe" for building models is inherited, it will occur that the models are gradually built up, since some new pieces will need to be added and the model tweaked in an attempt to deal with perceived problems and omissions (1999, p. 31).

The result of these model-building practices are representative instruments which will enable us to learn both about them and about the world. When we manipulate the models, we learn about the situation that is depicted by them—*i.e.*, what is often called "the model world". This is the starting point for learning or understanding reality, and what is physically possible or impossible (1999, p. 33). Thus, that models serve as agents in the learning process explains a) how models can be both a means to and a source of knowledge, and b) why they are so relevant to scientific practice (1999, p. 35). Morgan and Morrison elaborate on these two points as follows:

> A virtue of our account is that it shows how and why models function as a separate tool amongst the arsenal of available scientific methods. The implication of our investigations is that models should no longer be treated as subordinate to theory and data in the production of knowledge. Models join with measuring instruments, experiments, theories and data as one of the essential ingredients in the practice of science. No longer should they be seen just as 'preliminary theories' in physics, nor as a sign of the impossibility of making [a discipline] a 'proper' science. (1999, p. 36)

The present project will be carried out under the assumption that the theorizing done in formal semantics mostly (if not completely) involves modeling practices. Hence, the goal here will be to offer an epistemological account of how semanticists can draw inferences and derive knowledge by applying semantic models to the study of natural language meaning. In order to do this, some case studies which illustrate these modeling practices will be assessed. As the reader can notice, this approach has been inspired by Morgan and Morrison's insight that paying attention to scientific practice is essential for explaining how models help us to learn about theories and about the world. Moreover, the proposal to be offered later will be more or less consistent with the claim that models are autonomous agents which serve the function of being instruments of investigation.

However, it is not the aim of this thesis to defend Morgan and Morrison's account of scientific modeling. In fact, one aspect in which the account to be offered here will differ from theirs is related to the role that representation plays in learning about the models themselves and about the world. Accordingly, it will be argued later that the modeling practices in formal semantics deliver knowledge even if the models built by semanticists do not primarily serve a representational role. Still, the reason why Morgan and Morrison's proposal has been presented here is because it serves as a good starting point for understanding the spirit of the present project.

Chapter 3

Semantic Theorizing: Preliminaries

Before examining our case studies, a few things need to be said or clarified for getting a sense of what the modeling practices involved in formal semantics are about. To start, we can consider this nice overview that Jacobson (2014) provides: "[semanticists ask] what a compositional semantics might look like: how can we model the tools available (...unconsciously) to speakers of a language that allow them to compute meanings of larger expressions from the meanings of the smaller ones that make them up. What are the formal ways in which meanings combine? And what are the types of objects that we need in order to model that?" (2014, p. 8). In the present chapter, we will briefly review some of the theoretical notions involved in said enterprise: truth-conditions, semantic values and compositionality.

3.1 Truth-conditions

Why are linguists interested in semantic theorizing? An important aim in this project is to reveal something about the semantic competence of the users of a language. The underlying assumption is that speakers possess some knowledge which enables them to make use of their linguistic resources in order to convey information. Since meanings have been traditionally conceived of as the medium for encoding information, we can speculate that communication is possible at least in part because speakers have knowledge of the meanings of linguistic expressions. Otherwise, how would they know which linguistic expressions to use in order to express what they intend to express? Or how would they know what someone else wants to convey by uttering certain expressions? The problem is that it is very difficult to say what meanings are or where to locate them. But if meanings were to play the role of evidence for linguistic theorizing, we would need to find a way to manipulate them in order to test our preliminary hypotheses or findings. And given that we don't want this obstacle to stop us from doing semantic theorizing, we need to find something that can actually play the appropriate role.

The interest in meanings is strongly connected to the interest in the contents that are expressed by, for example, acts of assertion. When speakers assert something, they use a suitable linguistic construction to convey the information that the world, or things, broadly, are in a certain way. Hence, in order to interpret what was meant by some utterance, we need to be able to identify what would it take for the uttered sentences to be true, regardless of whether we know if they are in fact true. These conditions under which sentences would be true are known as *truth-conditions*. Then, competent speakers should possess the knowledge which enables them to figure out what the truth-conditions of particular sentences are. And since the knowledge of sentential meanings is associated with the knowledge of truth-conditions, we can indirectly theorize about the former by directly theorizing about the latter. As a result, formal semanticists will aim to give a systematic account (via their models) of how sentences can be paired with truth-conditions[1].

3.2 Semantic Values

As we have seen, semantic theorizing involves pairing sentences with posits such as truth-conditions. But in the more general case, we will find that semantic modeling also involves pairing a wider set of linguistic constructions with other elements in the semantic models. More technically, it can be said that a semantics for a natural language pairs well-formed expressions from the object language with model-theoretic entities such as objects, sets or functions. We will call such pairing a *semantic value assignment*. This assignment is an interpretation function[2] for the object language

[1]Certainly, it is very unlikely that truth-conditions themselves play the role of evidence or of data. When linguists attempt to test their models, or when they find some data point they want to explain, they look around for linguistic intuitions. That is, they look for the judgments of speakers regarding what sort of information can and cannot be conveyed by using a particular linguistic construction. These intuitions are supposed to reveal something about the semantic competence of the speaker. Consequently, linguistic data are usually associated with intuitions, and these intuitions can be elicited and "manipulated" in order to confirm/disconfirm generalizations. In contrast, truth-conditions appear to be on the more theoretical side of the spectrum—something that is processed or extracted from these intuitions. Still, we won't be taking any particular stand here as to what form the actual linguistic data or evidence takes (for example, as to whether the linguistic intuitions themselves are the evidence or some processed version of it). This is still an open matter for discussion which goes beyond the scope of this thesis. For some tentative proposals regarding linguistic intuitions and the role they play in linguistic/semantic theorizing, the reader is referred to Devitt (2010) and Schindler et al. (2020) (especially, Collins, 2020).

[2]The semantic value assignment is assumed to be a partial function, *i.e.*, a function that doesn't pair every grammatical linguistic expression with a semantic value. That the assignment is partial should account for why grammaticality (being well-formed according to syntactic rules) is different than interpretability (being a meaningful expression).

which takes linguistic expressions as inputs and delivers semantic values as outputs. A *semantic value* just is whatever entity the interpretation function assigns to a particular expression.

3.2.1 Extensional and intensional semantic value assignments

Most of the modeling practices we will examine in Chapter 4 are depicted in Heim and Kratzer's *Semantics in Generative Grammar* (1998) textbook. Since the textbook offers a basic recipe for modeling the semantics of a very limited set of English constructions, it will serve as a useful starting point. Hence, some clarifications about Heim and Kratzer's approach are in order.

First, Heim and Kratzer model and define both an extensional and an intensional semantic value assignment. In an *extensional* semantics, the semantic values or denotations of linguistic expressions are all extensions—or in more familiar terms (following Frege), their "referents" (1998, p. 15). For instance, the denotation of proper names are the individuals to which the names apply, the denotation of intransitive verbs are functions that go from individuals to truth-values, and the denotation of sentences are their truth-values (either true or false). Let's illustrate this with some examples.

In formal notation, whenever we have a linguistic expression α from the object language enclosed in double brackets, we are symbolizing its semantic value, as in $[\![\alpha]\!]$[3]. Here, the double brackets ($[\![.]\!]$) stand for the interpretation function in the metalanguage. Then, we form identity statements to indicate what these expressions get paired with, as in:

(1) $[\![$**Betty**$]\!]$ = Betty

Any statement of the form $\ulcorner [\![\alpha]\!] = x \urcorner$ is supposed to be read as "the denotation (or the semantic value) of α is x". (1) is a statement in the metalanguage that says that the semantic value of the proper name *Betty* is the individual Betty. Now, for the extensional semantic values of the verb *jumps* and the sentence *Betty jumps*, we get:

(2) $[\![$**jumps**$]\!]$ = the function that maps each individual from the domain to 1 iff they jump[4]

(3) $[\![$**Betty jumps**$]\!]$ = 1 iff Betty jumps

[3]Strictly speaking, the object language expressions should be enclosed within quotation marks when they are being mentioned. Here, Heim and Kratzer's convention of using bold-face text when the expressions are enclosed in double brackets will be followed (1998, pp. 22–23). Italics will be used for all other mentions.

[4]Alternatively, we could have *the subset of individuals (from the domain) that jump* in the right-hand side of the statement.

A semantic framework can also assign intensions to expressions if there is a need for more fine-grained distinctions—for example, if there are modal operators or propositional attitude verbs involved. Intensions (following Carnap) are what determine how extensions depend on possible circumstances. Then, an *intensional* semantics is one where semantic values are functions that take circumstances of evaluation as inputs and give appropriate extensions as outputs.

Under Heim and Kratzer's intensional framework, the only relevant circumstances of evaluation are *possible worlds*, broadly understood as "ways the world could be" (1998, p. 302). Consequently, the new inventory of denotations will also include (in addition to our familiar extensions) functions from possible worlds to functions from individuals to truth-values for intransitive verbs, and functions from possible worlds to truth-values for sentences. Now, for our toy examples (notice the superscript w which indicates the point of evaluation):

(4) $[\![\textbf{Betty}]\!]^w$ = Betty[5].

(5) $[\![\textbf{jumps}]\!]^w$ = the function that maps each individual from the domain to 1 iff they jump in world w

(6) $[\![\textbf{Betty jumps}]\!]^w$ = 1 iff Betty jumps in world w

Within formal semantics, a *proposition* is model-theoretically defined as a function from possible worlds to truth-values, or alternatively, as the set of worlds which the function maps to the truth-value 1 or *true*. Hence, in the Heim and Kratzer framework, propositions play the role of being the intensional semantic values of sentences. As we have said, possible worlds are circumstances of evaluation. And the circumstances under which a sentence is true are its *truth-conditions*. So, if a proposition is the set of possible worlds in which it is true, it is also the set of circumstances under which it is true. Hence, by pairing sentences with propositions, the intensional framework is pairing the former with their truth-conditions.

As you may recall from earlier, to offer such a pairing was one of the main goals of semantic theorizing. However, we won't say much in this thesis about propositions, or about the modeling of truth-conditions, for that matter. All of our examples will involve extensional semantic value assignments, to keep matters simple. What must be kept in mind, though, is that semanticists care about these extensional semantic values insofar as they can be proxies for the intensional ones.

[5]Within formal semantics, proper names are assumed to be *rigid designators*. That is, their reference is not relative to circumstances of evaluation, or possible worlds, in this case. Thus, even in an intensional framework, proper names are paired with extensions (Heim & Kratzer, 1998, p. 304).

3.3 Compositionality

Within formal semantics, it is customary to work with compositional semantic value assignments. That is, the interpretation function for a natural language fragment must adhere to the *Principle of Compositionality*[6]:

(7) **Compositionality**

The semantic value of an expression is a function of the semantic contribution of its parts and the way in which these parts are combined by the syntax.

The above principle states that complex semantic values depend on both simpler semantic values and on the syntax of the language under consideration. For that reason, when we say that a language is compositional, what we have in mind is a formal property of its semantic value assignment in relation to a syntax (Pagin & Westerståhl, 2010, p. 253). If the assignment is to have this property, then it must be compositional for every syntactic operation in the language. A common way to interpret this claim is to hold that for every syntactic operation that gives rise to a complex expression, there will be a corresponding meaning-composition rule that is applied to the semantic values of immediate constituents in order to derive the semantic value of the more complex expression (Yli-Vakkuri, 2013, p. 526; Pagin & Westerståhl, 2010, p. 254).

The idea that natural languages must be compositional has been partly motivated by two considerations: *productivity* and *systematicity*. On one hand, arguments which focus on productivity point out that language users are able to grasp the meanings of complex expressions that they have never heard before. Certainly, it would be very unlikely for them to have learned all of those meanings from past experience. Thus, there must be some sort of "knowledge" (which doesn't have to correspond to *knowledge-that* nor *knowledge-how*[7]) that allows them to determine the meanings of complex expressions. On the other hand, arguments from systematicity point out that language users who are able to understand a complex expression can also understand other complex expressions which have been constructed by recombining the former's constituents (Szabó, 2012, pp. 75–76). Then, the compositionality of natural languages appears to be the best explanation

[6]It must be noted that every formulation of the Principle of Compositionality talks about the compositionality of "meaning". Since we will be dealing in this paper with models of formal semantics and these work with semantic values, the principle will be stated here in terms of semantic values.

[7]Within generative linguistics, linguistic competence is thought of as the state of cognizing one's own language, yet this state doesn't have to be conscious or transparent to the language user.

for these facts. That is, the knowledge that enables speakers to determine or construct unfamiliar complex meanings must at least involve the knowledge of a limited number of simple meanings (that of words, for example) and of grammatical structures.

Within the formal semantics literature, it is assumed that these "simple meanings" are stored in the lexicon, or in the corresponding lexical entries of each linguistic item. The *lexicon* is, roughly, a list of words (or any meaningful elements, but we will assume that these are words to keep matters simple) and their properties that are part of a speaker's knowledge of their own language. Hence, the lexicon can be thought of as a mental dictionary. Each item from the lexicon is called a *lexical entry* and should at least contain information about the pronunciation of the word, its "meaning" (*i.e.*, its lexical semantic value) and its syntactic category. Then, the computational component of the language faculty, or the modules of the grammar must access this information in order to construct, interpret and pronounce complex expressions (Carnie, 2013, pp. 236–237). We will have a bit more to say about these modules in the next subsection.

3.3.1 The syntax-semantics interface

Recall that, according to the Principle of Compositionality (7), the semantic value of a complex expression also depends on the way in which its constituents are syntactically combined. So, how do semanticists go about implementing the compositionality assumption as part of their models? Given that the principle establishes that semantic composition mirrors syntactic composition, they need to assume some syntactic analysis for the constructions under consideration. Once the latter is in place, semanticists can proceed to combine functions and arguments in a similar order in which syntactic units are combined.

One way to implement this is is to assume that the input to the interpretation function are uninterpreted syntactic structures. In other words, under this approach, the semantic value assignment pairs syntactic structures with semantic values. Hence, we should replace (3) with:

(8)

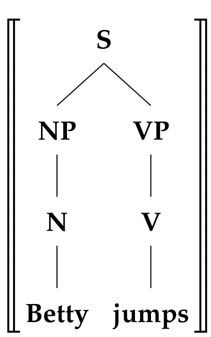

= 1 iff Betty jumps

This presupposition is consistent with how the syntactic and semantic modules of the grammar are thought to interact within generative linguistics, according to the inverted Y-model of grammar (*cf*. Carnie, 2013, ch. 13). The relevant assumption is that the grammar as a cognitive faculty consists of subcomponents which are autonomous from one another, and which only interact at the interfaces. For instance, the syntactic module builds a syntactic structure by applying a set of syntactic rules. This structure is then sent to two interfaces: the Phonological Form (PF) and the Logical Form (LF) levels of representation. In order to receive a semantic interpretation, the syntactic structure is assigned an LF-representation at the LF-level. This representation is the result of applying additional syntactic rules. Then, semantic composition is computed by applying compositional rules to this LF-representation (Heim & Kratzer, 1998, p. 185).

But why should we accept the claim that the syntax is autonomous from the semantics? Although the assumption has not been universally endorsed, there are certainly some data-related and theoretical motivations for it. For instance, acceptability judgments would suggest that grammaticality and interpretability can come apart, *viz.*, a structure can be well-formed but "meaningless", or it can be ungrammatical yet interpretable by speakers. On the more theoretical side, the Logical Form level of representation is necessary to account for movement operations which have an effect on the interpretation a sentence receives (1998, p. 185), such as Quantifier Raising. More about this will be said in Chapter 4.

3.3.2 Frege's Conjecture

For the semantic value assignment to be compositional, sentential semantic values must be derived via the application of compositional semantic rules. For this reason, a key assumption in Heim and Kratzer's system is the principle that they call *Frege's Conjecture*: "semantic composition is functional application" (1998, p. 29). This is rooted in Frege's belief that the meanings of linguistic expressions can be classified either as saturated or unsaturated. Unsaturated entities are missing something—an argument, so it is appropriate to think of them as mathematical functions. Once a function is applied to an element that acts as its argument, saturation occurs (*cf*. Frege, 1997). Hence, very roughly, Frege's Conjecture states that whenever two linguistic expressions combine semantically, the meaning of one is an argument for the meaning of the other. This insight is captured by the following rule:

(9) **Functional Application (in more or less simplified English)**

If α is a branching node with β and γ as its daughters, and β's semantic value is a function that can take the semantic value of γ as its argument, then the semantic value of α is the result of applying the function β denotes to the argument that γ denotes (Heim & Kratzer, 1998, p. 44)[8].

As way of illustration, let's re-consider the representation in (8):

(8) 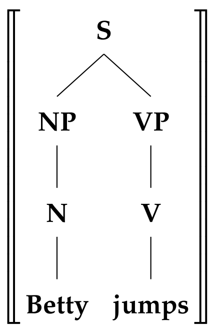

= 1 iff Betty jumps

If we are to apply the Functional Application rule, we must start by noticing that S is a branching node with NP and VP as its daughters. And this means that in order to determine the semantic value of S (*1 iff Betty jumps*), we will need to semantically combine NP and VP by applying the semantic value of one of them to the semantic value of the other. But how do we figure out if that is even possible?

Semantic types

In order to keep track of how the Functional Application rule must be applied, it is customary to include in the formal apparatus semantic types for denotations[9]. e and t are the basic types of denotations, which also correspond to saturated entities: e is the type of individuals, and t is the type of truth-values. These correspond to the extensional semantic values of proper names and sentences, respectively.

Since unsaturated entities are functions from meanings to meanings, their semantic types are built from other semantic types. We write the semantic types for functions as $\langle \sigma, \tau \rangle$ to indicate that they take semantic values of type σ as inputs and give semantic values of type τ as outputs. For example, the semantic values of intransitive verbs are functions of type $\langle e, t \rangle$ since they take

[8]Notice that this isn't the final (pedantic) version that Heim and Kratzer use. I am omitting all references to assignment functions for now.

[9]In the formal semantics literature, this is known as *type-driven interpretation* (*cf*. Heim & Kratzer, 1998, p. 44).

individuals as inputs and give truth-values as outputs. In other cases, we will have that some functions can take other functions as their arguments. Hence, more generally we will say that in order to apply Functional Application, one semantic value must be a function and the other must be something in that function's domain.

As was mentioned before, our examples will only involve extensional semantic value assignments, so all we need to keep in mind for the remainder of this thesis is that (extensional) semantic values can only have types e, t or some other type recursively constructed from those two.

Chapter 4

Modeling Practices in Linguistic Semantics

In this section, we will focus on the modeling practices involved in formal semantics. More specifically, a few case studies which illustrate (some of) the kinds of inferences that formal semanticists draw in order to model derivations of meaning composition will be presented: inferences within models, from data to models, and between models. Semantic models should be understood here as the tools that formal semanticists build and apply to the study of natural language meaning. Even though issues related to the metaphysics of these models will be left aside (for the most part), the examples should elicit the intuition that semantic models involve *at least* a semantic value assignment, a set of semantic rules and a lexicon.

4.1 The Search For Tools

One of the key aims of the practices of formal semanticists is to find suitable tools for modeling. As some of the examples in this section will show, most of these formal tools are extracted or "borrowed" from other disciplines and appropriately adapted so they can be applied to the case at hand. These devices can come directly from mathematical notation or from modeling devices, techniques or methods which have been found to be useful in domains such as formal logic, probability theory, or computer science.

If, as has been vaguely gestured, semantic models contain at least three components, then the model-building practices of semanticists will involve the search and selection of adequate semantic values, semantic rules and lexical entries. Yet these searches are not independent of one another. As we will see, it is the job of the semanticist to make these essential ingredients fit together by finding an appropriate formal notation. For that reason, any decision that semanticists make regarding one component will constrain the possible options for modeling the other two components. That is, the model itself, or what has been built of it so far, dictates to some degree future choices.

4.1.1 Modeling a semantic value assignment

So, how does one model a semantic value assignment? As you may recall from Chapter 3, there are semantic values which appear to be taken for granted given methodological commitments like Frege's Conjecture. For instance, the obvious choice for playing the role of an unsaturated expression which is missing an argument is a function. And functions must be applied to mathematical objects. If this is an appropriate way of modeling, for example, instances where something is predicated of an individual, then a reasonable choice for semantic values are functions in the case of the linguistic expressions which serve the role of predicates, and objects of a domain (*viz.*, the domain of individuals of the interpretation function) in the case of proper names.

But this is not the only possible way to do so. For instance, we can think about predication in a more intuitive way: when we attribute a property to an individual, we convey (amongst other things) that the individual belongs to a specific group—the group of those who share the same relevant feature. Since sets are widely used to represent groups, then it would seem that sets can also be suitable entities for modeling verbs, common nouns, adjectives, and the like[1]. The issue, though, is that working with sets would preclude a literal implementation of Frege's Conjecture. For that reason, Heim and Kratzer opt for using functions as denotations for the aforementioned linguistic expressions. Luckily for them, functions and sets are tightly connected to one another: each set has a *characteristic function*. That is, for each set, there is a function which takes entities from the domain as arguments, and delivers the value 1 if these entities are members of the set, and 0 otherwise. By taking advantage of this fact, they can conveniently (and informally) talk about the denotation of some expressions as sets whenever the function talk becomes too cumbersome (Heim & Kratzer, 1998, pp. 24–26).

[1] The intuition becomes stronger once expressions which involve noun modifiers are considered. For example, if we want to model the meaning of an expression such as *poodles in Lake Park*, we can use a set intersection operation which would be applied to the set of poodles and the set of individuals in Lake Park. But, as we will see shortly, Heim and Kratzer continue to use functions even when dealing with these cases.

At this point, a choice between functions and sets appears to have been made. The next step is to specify what kind of functions can play the roles we need them to perform: the roles of being the semantic values of certain expressions. Now, we will consider a toy example of how the model can provide enough information to make this decision.

Case 1: The denotation of intransitive verbs

When Heim and Kratzer are attempting to model a derivation which involves an intransitive verb, they need to figure out what the extensional semantic value of intransitive verbs should be. In order to do this, they need to consider other elements in addition to the rule of Functional Application. For example, they need to take into account other semantic values (that of proper names and sentences), and other semantic rules (the ones for non-branching nodes).

Let's recreate their reasoning. First, if we want to predict the semantic value of an expression, we must begin by assuming that there's some syntactic representation that we will need to interpret. For instance, (10) corresponds to the sentence *Betty jumps*:

(10)

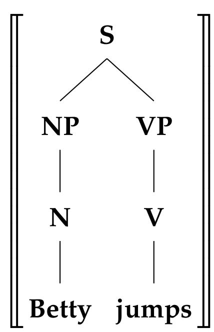

In order to compute the extensional semantic value of *jumps*, Heim and Kratzer used the following semantic value assignment and semantic rules as ingredients (*cf*. 1998, pp. 13–20):

(A) The (extensional) semantic value of the S node must be a truth-value such as 1 iff Betty jumps.

(B) The semantic value of the proper name *Betty* is the individual Betty.

(C) A set of semantic rules:

1. *Terminal Nodes*: The semantic value of a terminal node is encoded in the item's lexical entry.

2. *Non-Branching Nodes*: Nodes with one daughter inherit the semantic value of their daughters.
3. Functional Application (see (9))

Then, the strategy they followed was to answer the question of what semantic value the lexical entry for *jumps* should encode in order to obtain (A), given (B) and (C). In other words, the goal was to accurately predict the datum that *Betty jumps* is true if and only if Betty jumps, and a means to it was to come up with a semantic value (or a semantic type) that could be integrated with the rest of elements that the model already contained.

Since S is a branching node, its semantic value should be computed by applying the Functional Application rule. From this, it follows that its daughters NP and VP should satisfy the requirement that one of them must have a function as its semantic value, and the other must denote something in the domain of that function. Now, (A) tells us that the sentence's semantic value is a truth-value, so it must be of type t. Consequently, one of S's daughters must have as its semantic value a function of type $\langle x, t \rangle$. Our job now is to figure out what x is. Because of (B), we know that *Betty*'s denotation is of type e. If we apply the Non-Branching Nodes rule, we obtain that NP inherits *Betty*'s semantic value. Therefore, the semantic value of NP is of type e. This seems like a suitable argument for a function of type $\langle x, t \rangle$. Hence, it can be reasonable to assign to VP a semantic value of type $\langle e, t \rangle$. Finally, by applying Non-Branching Nodes again, we can obtain the type of *jumps*'s semantic value: $\langle e, t \rangle$. If we have to specify a condition for this function, then we can think of it as mapping entities from the domain of individuals to truth-values depending on whether they jump or not. The denotation of *jumps* is then encoded in its lexical entry.

The assumption that these functions of type $\langle e, t \rangle$ are adequate denotations for intransitive verbs is then extended to the treatment of adjectives, nouns and prepositions. That is, the lexical entries for many of these expressions will encode functions from individuals to truth-values (Heim & Kratzer, 1998, p. 62). In the literature, these functions are known as *1-place predicates*.

4.1.2 Modeling semantic rules

Case 2: Predicate Modification

As Heim and Kratzer examined new syntactic constructions, they noticed that not every instance in which two linguistic expressions semantically combine can be modeled in a straightforward way as the application of a function to an object. For example, sometimes the meaning of complex expressions such as *fluffy cats*, which involve common nouns and their modifiers, must also be computed. As we have just established, within the model, the semantic values of adjectives and common nouns are 1-place predicates, or functions of type $\langle e, t \rangle$. Given that the semantic values of *fluffy* and *cats* are of the same type, we cannot think that either of them is a function to be applied to the semantic value of the other. Hence, we cannot derive the meaning of the expression by applying Functional Application. Instead, the following rule must be applied:

(11) **Predicate Modification (in more or less simplified English)**
If α is a branching node with β and γ as its daughters, and the semantic values of both β and γ are functions from individuals to truth-values, then the semantic value of α is the characteristic function of the set that is the intersection of the sets characterized by the semantic values of β and γ.

Still, Heim and Kratzer acknowledged that these cases could also be modeled by applying Functional Application:

> Predicate modification is a genuinely new semantic composition principle on our list. It is obviously not functional application. If it is really needed, there is more to semantic composition than Frege's Conjecture. Are we forced to this conclusion?... By systematically revising the entries of all adjectives and prepositions, we are able to interpret all phrases containing a noun with one or more modifiers in them by means of Functional Application alone, and so we can eliminate Predicate Modification from the theory. (1998, pp. 66–67)

The problem, though, is that there would be an extra cost for keeping Functional Application as the only compositional rule in the system: they would have to either posit a lexical ambiguity (*i.e.*, two lexical entries) for the copula "be", or keep "be" as before and instead posit a lexical ambiguity for all adjectives and prepositions. In short, there was a trade-off between a literal implementation of Frege's Conjecture and simplicity in the lexicon, and Heim and Kratzer went for the latter.

4.2 Contrasting Models With Independent Data

The examples we just reviewed are certainly not all there is to the Heim and Kratzer model. In fact, the model is eventually expanded to cover cases which involve quantifiers and pronouns (as we will see later), and also propositional attitude verbs. Ideally, models would be flexible enough to be applicable to the study of a wide array of linguistic constructions (it would be great if we could use the same tools for modeling the meanings of pretty much everything!). Yet in reality, the models' applicability will be restricted to specific "domains", and it just won't be possible to extend the models as they are to cover more complex linguistic phenomena. As a consequence, there is an effort in formal semantics to "recycle" components from previous models, especially when it has been proven already that the latter succeed in some respect—perhaps they are simple, elegant or have good predictive power.

But how do we go about finding out whether a model can be more or less extended or whether specific tools can be recycled? We want to construct models which can account for the data, so naturally, a way to figure out if this is the case is to start testing the models' predictive capabilities. In order to examine how models or modeling devices can interact with the data for the purpose of extending the range of applications, we will consider the case of the *predicate vs. formulas debate* in the domain of verb phrase ellipsis.

4.2.1 Case 3: The predicates vs. formulas debate

As you may recall from the time we were examining how Heim and Kratzer modeled the denotation of intransitive verbs, one of the inferences we made led us to think of the semantic values of verb phrases as functions from individuals to truth-values, or 1-place predicates. More generally, we will call this claim *the predicate hypothesis*. *Prima facie*, endorsing this hypothesis seems like an unproblematic move for offering an account of semantic composition in simple cases like those which Heim and Kratzer considered. Yet as Heim (1997) admits, there doesn't seem to be much empirical justification for this choice (1997, p. 197). After all, in our assessment of the Heim and Kratzer model, we noted that most of the decisions in this regard seemed to already be determined by working assumptions as well as other components of the model (the only datum we invoked was that the truth-value of *Betty jumps* should be 1 if and only if Betty jumps). Then, this leaves open the question of whether there is another tool which could do a better job at being the semantic value of verb phrases[2].

[2]A clarification is in order: in the first case study, we were focused on the choice of 1-place predicates as the semantic values of *intransitive verbs*. But, as (10) reveals, a verb phrase isn't the same as an intransitive verb—that is why VP and V are different nodes. Then, it was just a feature of the example under consideration that the verb phrase only contained the intransitive verb. Intransitive verbs don't take any internal arguments, anyway. Yet, in other cases, verb phrases also include the direct object.

Heim (1997) also acknowledged, in light of the data presented in Kennedy (1994), that the predicate hypothesis might be inadequate in the context of a theory of verb phrase (VP)-ellipsis. VP-ellipsis constructions involve "missing" verb phrases, yet they are interpreted as if the verb phrases were present. Moreover, the available interpretations (if any) are constrained by other VPs which are in fact present in the structure—that is, which are pronounced. For example, consider:

(12) John left, and George did too.

We interpret (12) as expressing that George also left. Thus, the elided material in the construction is assumed to be the verb phrase *leave*, as (13) shows:

(13) John left, and George did <~~leave~~> too.

Now, if we maintain that the denotation of these verb phrases are 1-place predicates, then we end up making the wrong predictions about the possibility of eliding VPs in larger constructions. For instance, whereas it is possible to elide *leave* and preserve the intended interpretation in (14a), it isn't possible in (14b):

(14) a. Every man who wants to leave should <~~leave~~>. (Heim, 1997, p. 200)
b. *Every man who wants George to leave should <~~leave~~>[3].

The difficulty here is that the predicate hypothesis incorrectly takes VP ellipsis to be licensed in (14b). Inspired by these findings, Heim explores whether an alternative yet seemingly equivalent notation would be better suited for modeling verb phrases and other expressions such as relative clauses, and those which compose with quantificational determiners. Under this *formulas hypothesis* approach, VPs denote open propositions—*i.e.*, expressions which contain free variables. It must be noted, though, that Heim's proposal for the denotation of VPs is more or less restricted since it can only be integrated with accounts in the domain of ellipsis which presuppose a focus-based semantics such as Rooth (1992)'s (1997, p. 216).

But even if Heim succeeded at accounting for some cases, matters were far from settled yet. For instance, Jacobson (1998) argued from new data regarding quantificational binding that the

[3]The " * " symbol is meant to show that the sentence would be ungrammatical (*i.e.*, that it isn't well-formed) if it had been constructed to express the intended reading. The intuition that the example aims to elicit is that once *leave* is elided, we cannot use the construction to express that what every man who wants George to leave should do is to leave.

formulas hypothesis had some problems of its own. In the case of (15), it would wrongly predict the construction to be unavailable:

(15) Every woman who John loves spoke to every woman who $[\text{BILL}]_F$ does[4]. (1998, p. 89)

As Jacobson points out, we seem to have a dilemma given our two hypotheses and the linguistic data. Proposals which do not make use of variables and hence assume the denotation of VPs to be closed expressions, have an advantage when accounting for cases such as (15). However, these accounts struggle with Kennedy (1994)'s original cases. In contrast, Heim (1997)'s formulas hypothesis appears to make the appropriate predictions for the latter, but not for (15) (1998, p. 90).

In order to address this challenge, Kennedy (2014) offers an account in which we incorporate both predicates *and* formulas to the framework. Before, the choice seemed to be between an analysis of quantification under which quantifiers compose with two predicates or another under which they compose with two formulas. Instead, Kennedy argues that the semantic value of the expression that corresponds to the quantifier's restriction (e.g., *cat* in *every cat purrs*) is a 1-place predicate, whereas the denotation of the expression that corresponds to its scope is a formula (e.g., *purrs* in the previous example). Moreover, the particular implementation that Kennedy adopts is one where we start with formulas and turn them into predicates as needed (2014, p. 265).

To sum up, the debate was motivated by the fact that the original view which took 1-place predicates to be the semantic values of verb phrases could not make adequate predictions for some instances of VP-ellipsis, when coupled with an account of ellipsis phenomena such as Rooth (1992)'s focus semantics. And this finding was revealed when the data provided in Kennedy (1994) were considered. As a result, any comprehensive account of how verb phrases semantically compose with other expressions in these linguistic environments had to involve some changes.

For example, Heim's move was to replace the denotation of verb phrases. Instead of working with 1-place predicates, she opted for using open propositions as semantic values. Still, it is always an available alternative to give up other commitments. One of them is the account of VP-ellipsis to be endorsed, since there could in principle be other models of VP-ellipsis that deliver the right results when integrated with the predicates hypothesis. And in fact, Heim had already examined this possibility by contrasting Rooth's account and a previous one (1997, pp. 200–208). More generally, she was well aware that there were a lot of moving pieces to her proposal:

[4] $[.]_F$ here indicates the constituent which is emphasized by way of a pitch accent in pronunciation, that is, the phrase which is focus-marked.

> And while we search for [independent] evidence [against the predicate hypothesis], we want to theorize about the general principles of syntax and semantics which are ultimately responsible for the fact that we found formulas rather than predicates where we did. What have to be our basic assumptions about the universal inventory of semantic composition operations, about possible meanings of lexical items, about the construction of phrase structure, and about the nature of movement, so that the theory makes available just those LF-structures which we found to be attested? (Heim, 1997, pp. 216–217)

Hence, other possibilities involved the modification of other elements in the system or of the underlying assumptions about the syntax-semantics interface, and this is the approach that Kennedy (2014) followed. As the above quote details, the predicates vs. formulas debate is ultimately about which kind of LF-representations are better suited to account for Kennedy (1994)'s data. That is, whatever approach one endorses, one ends up being committed to different representations at the LF-level for quantificational determiners and relative clauses, in addition to verb phrases. Hence, what we are after is one of many possible implementations of the syntax-semantics interface. This means that another commitment which could have been given up is a movement approach to binding by quantifiers based on Quantifier Raising and/or its semantic implementation which involves the Predicate Abstraction rule. To these issues we will turn our attention in the following sections.

4.3 Integrating Models Across Sub-Fields

4.3.1 Case 4: Variable binding

Even in the case of Predicate Modification, there was a clear alternative to the introduction of the rule—we could have kept Functional Application as the only rule in the system. But could the same be said about more complex cases? For instance, semanticists have noted that sentences like (16) require a special form of analysis that could go beyond Frege's Conjecture. That is, some expressions involving quantifiers pose a challenge to the claim that semantic composition is just Functional Application.

(16) Marcia loves every linguistics class.

As (17) shows, if we attempt to compose the transitive verb *loves* and the quantificational DP *every linguistic class*, the derivation will crash due to a type mismatch. Both constituents denote functions (we haven't reviewed here how Heim and Kratzer reached these conclusions, so you will have to take my word for it), yet neither semantic value can act as an argument for the other (Heim & Kratzer, 1998, p. 179).

(17)

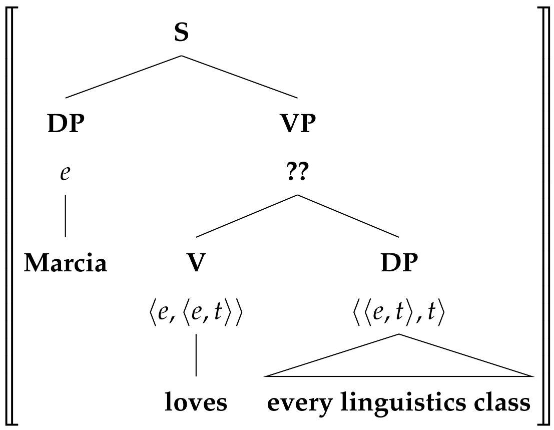

So, what is Heim and Kratzer's proposed solution? They offer a two-step approach to solving the type mismatch: first, they adopt a movement account which was already present in the literature in order to displace the quantificational DP. Second, they introduce a new semantic rule in order to allow for the semantic composition of the DP and the rest of the sentence. The final result is the following LF-representation:

(18)

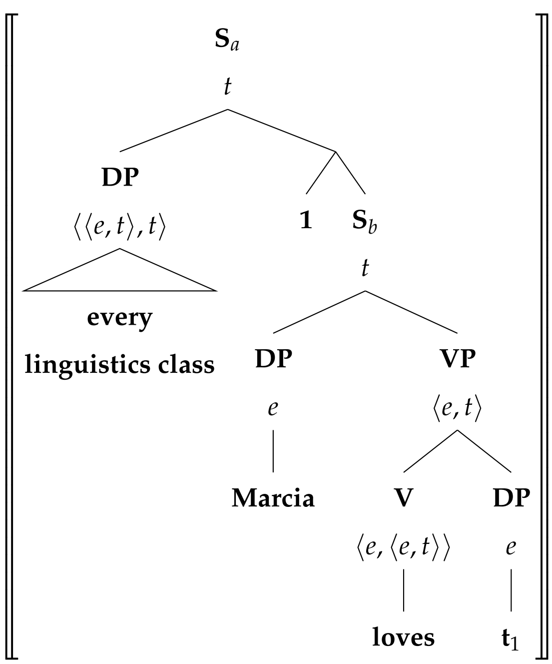

Let's review each step carefully. Notice that in (18), *every linguistics class* has been moved to the "front" of the structure and left a trace (t_1) in its original position. (A *trace* is an unpronounced syntactic element left by displaced elements to mark their original position.) This movement operation is known as *Quantifier Raising* (QR). In more technical terms, quantificational object DPs that undergo QR adjoin to S nodes (S_b in this particular case). But how is it possible that what receives a semantic interpretation is the structure in (18), and not something like (17)?

As you may recall from Chapter 3, the syntactic module of the grammar delivers an uninterpreted syntactic structure to both the phonological and semantic modules. This means that what gets pronounced (what linguists call "surface structure") can differ greatly from the structure which ultimately receives a semantic interpretation. As we saw earlier, once the structure reaches the Logical Form component, it might undergo additional transformations, including "covert" movement which doesn't have a phonetic realization. So, here Heim and Kratzer are taking advantange of the methodological commitments that were already in place within the model: endorsing QR isn't really problematic since it fits the story about the syntax-semantics interface provided by the Y-model of grammar (1998, p. 185). But once they decide to implement a movement approach via QR, they need to make sure that it can be integrated with the rest of the system. This is where the second step comes in.

For instance, notice that the displacement of the DP is not the only consequence of the QR operation. Under the movement approach, traces are treated as variables. For this reason, traces receive a denotation of type *e*, as you can observe in (18). As Heim and Kratzer note, "[v]ariables were invented precisely to be like ordinary referring phrases in the respects we want them to be[5]. . . [a] variable denotes an individual, but only relative to a choice of an assignment of a value" (1998, p. 92). However, we don't want the truth-value of *Marcia loves every linguistics class* to be relative to the semantic value that we happen to assign to the trace. Hence, we also need to introduce a binder for the latter—in the case of (18), the index 1 which is sister to S_b. Otherwise, the variable would remain free[6].

[5]A nice outcome of this decision is that now the semantic values of the trace t_1 and the transitive verb *loves* in (18) can compose via Functional Application.

[6]We are skipping over some details in this explanation. For instance, the reason why the trace must bear an index is not just because we need it to be bound by the displaced constituent. More generally, variables are interpreted relative to variable assignments (more on this later), so they can only be interpreted in the system—*i.e.*, assigned a denotation—if they bear an index (Heim & Kratzer, 1998, p. 186).

The difficulty, though, is that we still can't combine the relevant nodes by using Functional Application, nor implement the binding of the variable with the system in its current state[7]. For example, there would still be a type mismatch between the raised quantificational phrase and S_b in (18) if they were to directly compose with each other. As Functional Application stated, whenever two constituents compose, one must be an argument for the other. In our particular case, *Marcia loves* t_1 (S_b) is of type t. For sure, this isn't a function. The alternative is for it to be the argument of the quantificational DP. But the DP's semantic type is $\langle\langle e, t\rangle, t\rangle$, so S_b isn't a suitable argument either.

To solve these issues, a new rule called *Predicate (or Lambda) Abstraction* (1998, p. 114) is introduced in the semantic framework. The rule's key mechanism is the transformation of some semantic value of type x into a function that takes individuals as arguments and delivers some output of semantic type x. That is, the semantic type of the new function will be $\langle e, x\rangle$. In the particular case that interests us, the rule takes the extensional semantic value of a sentence (a truth-value) and transforms it into a function that goes from individuals to truth-values. Hence, the name of the rule: by applying Predicate Abstraction, we create a "predicate" from a sentence by abstracting over its actual truth-value. Moreover, the binder—the index 1—is expected to trigger the application of the rule. Since the node with the binder has no meaning of its own, we call this a *syncategorematic node*. Then, by applying Predicate Abstraction, we can transform S_b into a 1-place predicate with semantic type $\langle e, t\rangle$ that is able to compose with the semantic value of the quantificational DP.

A few comments are in order. First, Heim and Kratzer endorsed a particular movement approach that was appropriate for the view of the syntax-semantics interface that they were already committed to. That is, a view under which the linguistic construction which gets pronounced doesn't need to match the structure that ultimately receives an interpretation. But they didn't need to resort to movement. Hence, when evaluating how movement approaches fare against *in situ* approaches, they acknowledge that there are no decisive arguments in favor of QR:

> We will present some versions of these standard arguments [in favor of the movement approach]...But we should note at the outset that their ultimate force is very difficult to assess. For one thing, ...the most recent theories that [proponents of *in situ*

[7] This is just a manner of speaking. The Predicate Abstraction rule (see below) was already in place by the time Heim and Kratzer deal with quantifiers, since binding also plays a role in the treatment of relative clauses (e.g., the book that Marcia borrowed) (*cf.* Heim & Kratzer, 1998, chapter 5).

> approaches] have developed differ from the movement-based theories in so many respects at once that comparison becomes a very global and difficult matter. And an even more important difficulty is that the issue is not really an all-or-nothing choice. It is entirely conceivable that there exist both quantifier movement and mechanisms for in situ interpretation (of nonsubject quantifiers), and that the grammars of natural languages employ one option in some constructions and the other in others. *In fact, the majority of scholars who have done influential work in this domain have (explicitly or implicitly) taken such mixed positions of one sort or another. So we could not possibly purport here to give decisive evidence in favor of a pure movement approach.* The best we can hope for is to give you some preliminary idea of what sorts of considerations are pertinent, and what price you have to pay for a particular decision. (Heim and Kratzer, 1998, p. 194, my emphasis)

Second, the underlying theoretical commitments of the model didn't completely determine how the issue should be addressed, even if they made easier the particular implementation of the movement approach that we just examined. Thus, there were different alternatives available to Heim and Kratzer. For example, even after introducing Quantifier Raising, they could have avoided the introduction of the Predicate Abstraction rule. One way to do this is by relativizing every semantic value in the system to variable assignments, so that Functional Application could proceed as usual (*cf.* Glanzberg & King, 2020, §V). Yet there were other ways to deal with quantifiers in object position which avoid movement in general. Some of these approaches introduce type-shifting rules, instead, so specific semantic values can be modified and made suitable for composition. In the next section, we will consider a couple of cases in which modelers made this move instead.

4.4 Foundational Issues

The examples we have examined so far assumed more or less the same set of methodological commitments. However, not all modeling practices in formal semantics involve the same implementation of the syntax-semantics interface. Hence, we will now examine two competing proposals in order to compare how semantic modeling can also involve making more complex choices regarding foundational matters.

4.4.1 Case 5: Pronominal binding and variable-free semantics

A variable-full approach to pronominal binding

Pronouns have referential and non-referential uses. *Referential* uses include cases where pronouns must pick their reference from a linguistic antecedent. For example, in the case of (19), we take the pronoun (*him*) and the antecedent (*Joshua*) to co-refer. In other words, they share a referent, but they can only do so insofar as each of them has a referent on its own.

(19) *Joshua* loves cats, and cats love *him*.

Yet notice that pronouns don't have to inherit the reference from their antecedents, even if they could. For instance, (20) has two possible readings: Willy could have played the piano he owns or someone else's. For the latter interpretation to be possible, the context must provide a suitable salient individual so the pronoun can get a referent.

(20) Willy played *his* piano.

In contrast, *non-referential* uses are those in which the pronoun is bound by a linguistic antecedent such as a quantifier, as in (21). In this case, it can't be said that *no philosopher* and *them* co-refer, since *no philosopher* lacks a referent to begin with.

(21) *No philosopher* likes people who don't agree with *them*[8].

In order to keep a uniform account that applies to all sorts of uses of pronouns, Heim and Kratzer treat pronouns as variables. In this way, referential uses are analyzed as instances of free variables, and non-referential uses as cases of bound variables. Hence, they claim:

> *The simplest assumption* we can make at this point is that all pronouns have the same internal syntax and semantics. They must all bear an index (numerical subscript) at LF to be interpretable, and they are all interpreted by the same rule, Traces and Pronouns. The only thing that distinguishes referring pronouns from bound-variable pronouns is that they happen to be free variables. In other words, *the difference between referential and*

[8]Of course, there is an available reading according to which *them* refers to some contextually salient individual(s). We are not concerned with such use of the pronoun here.

bound-variable pronouns resides in the larger surrounding LF structure, not in the pronouns themselves (1998, p. 242, my emphasis)

Let's unpack their remarks. As we have seen, a treatment for quantifiers in terms of a movement operation approach is already in place, so non-referential uses won't be terribly problematic. For instance, the LF-representation of (21) still requires the movement of the quantificational DP, regardless of the fact that it was already in subject position:

(22) 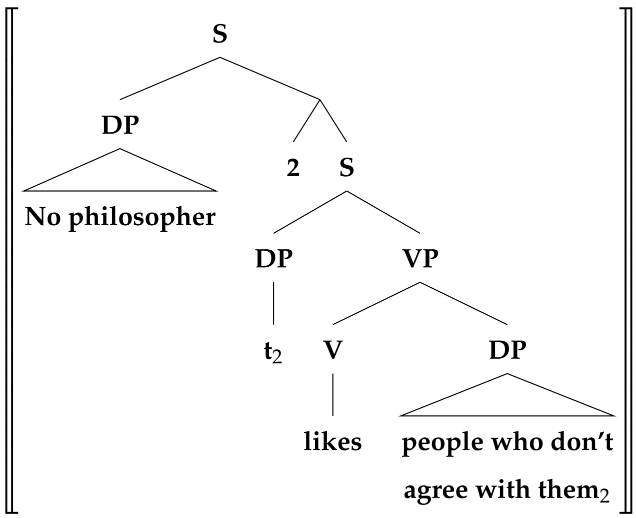

But these are not the only cases that can be analyzed in terms of movement operations. Sentences like (20) which include possessives are amenable to a similar treatment[9,10]:

(23) 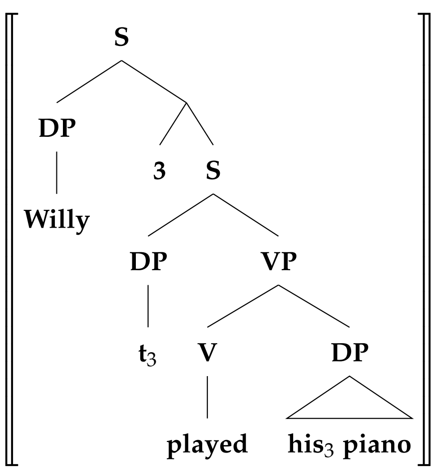

[9] We're assuming here the reading according to which the piano in question is Willy's.

[10] Actually, Heim and Kratzer's treatment of possessives involves definite descriptions (*cf.* 1998, pp. 246–248). For instance, in this particular case, *his*$_3$ *piano* would be semantically interpreted as ⟦the he$_3$ piano⟧. But we are skipping over these details in (23) since they are not necessary for visualizing the movement operation.

The more interesting cases, though, are those involving the free uses of pronouns. If pronouns are variables in the object-language, free pronouns are variables that remain free within the structure. Then, a question arises: how are we to compose them with the rest of the structure in order to account for the fact that their reference shifts with the context of utterance? For Heim and Kratzer, this phenomenon can be captured by appealing to the dependence of their semantic values on variable assignments.

Variable assignments (also called *assignment functions*) are partial[11] functions that take indices as arguments and deliver individuals. That is why a requirement for indices must be enforced at the LF-level. Then, the *Traces and Pronouns* rule that Heim and Kratzer posit lets us assign a semantic value to a free variable under an assignment function. In these cases, sentential semantic values are non-constant functions from variable assignments to truth-values. Then, these semantic values are open propositions (*cf*. Heim & Kratzer, 1998, ch. 9).

A detailed explanation of how semantic composition works for constituents that include pronouns will not be offered here, but some details are worth mentioning. First, free pronouns are able to semantically compose because our Functional Application rule gets amended to be general enough[12]. Once variables get introduced to the system, rules like Functional Application are amended so they can take variable assignments into account. Very roughly, the rule now has to define how two constituents compose semantically, each being interpreted under a particular variable assignment.

Second, by generalizing to the "worst-case scenario" (that is, free pronouns), we can give an account of both referential and non-referential uses. Every pronoun starts as a free variable (that is their lexical semantic contribution), but they can also be bound if needed. In that case, the constituent that contains the pronoun receives an assignment-independent interpretation. This is made possible by our Predicate Abstraction rule, which lets us bind a variable by an index[13]. But as we have also seen, Heim and Kratzer's movement approach entails that variables must be of the same semantic type as ordinary DPs (recall that traces are of type *e*). A natural consequence is that constituents that contain free pronouns are of the same semantic type as constituents that contain

[11]They are *partial* insofar as they don't need to (and probably won't) map every possible index to an individual. All that is required for an utterance with free pronouns to be felicitous is that the context determines a variable assignment that includes in its domain at least those indices that the free variables bear (Heim & Kratzer, 1998, p. 243).

[12]Again, this is a manner of speaking, since Functional Application had already been amended by the time Heim and Kratzer deal with pronouns.

[13]More specifically, the rule can introduce new indices to the domain of the assignment function whenever it is necessary to do so, and also stipulates the value associated with each new index. This is explained in terms of a *modified variable assignment* (*cf*. Heim & Kratzer, 1998, pp. 112–114).

ordinary DPs (Jacobson, 1999, pp. 124–125). All in all, it looks like uniformity and systematicity has been achieved. But at what cost?

Jacobson's variable free semantics

Jacobson (1999) pointed out that even sentences that don't include free pronouns end up also being related to assignment functions, under standard approaches to binding like Heim and Kratzer's. Given that these models attempt to give a uniform treatment of pronouns, the frameworks are adjusted to operate under assignment functions. So, underlyingly, the result is that all semantic values are relativized to variable assignments. Of course, the values that most of these functions will deliver are constant (1999, p. 144). For example, regardless of whether the semantic value of a proper name is a function from an assignment to an individual, said function will always deliver the same individual as output. Similarly, the extensional semantic values of sentences that don't involve free uses of pronouns will be constant functions from variable assignments to truth-values.

What Jacobson argues, then, is that a function's being constant doesn't make it any less of a function. Sentential semantic values *are* functions from variable assignments to truth-values. Some of these will be open propositions, and some will not. Moreover, the fact that these frameworks are not explicit about this dependency doesn't make the issue go away: variable assignments are already part of the object language, and so are variables (1999, p. 127). Hence, a natural question to consider is whether we are indeed forced to embrace such a complex semantic value assignment.

Jacobson replied that we are not, so her strategy was to offer a competing model—*variable-free semantics*—for the semantics of pronominal binding which makes no essential use of variables nor variable assignments as part of the object language. Then, even if variables still appear in semantic derivations, this is done for notational purposes only[14]. That is, variables are just part of the meta-language. Jacobson avoids their introduction as part of the object language by positing a different lexical entry for pronouns. Under the variable-free approach, all pronouns are *identity functions*[15] from individuals to individuals. The upshot is that pronouns now have the same denotation ev-

[14]Hence, Jacobson says: "To verify that these are not part of the actual semantic machinery, one can note that all of the variables in my representations are bound. (Another way to convince oneself that the use of variables is for notational convenience only is to note that all of the formulas with variables could be reformulated without by using combinators. Unfortunately, such reformulations would often be quite reader-unfriendly; hence the use of representations with variables.)" (1999, p. 122).

[15]An identity function is one that takes an argument as input and delivers it as output, too.

erywhere they occur, since their semantic values are not dependent on indices nor on assignment functions (1999, p. 145). Furthermore, unlike the standard approach, pronouns here do not have the same semantic type nor syntactic category as ordinary DPs (1999, p. 147). For Jacobson, that was an unwarranted assumption to begin with (1999, pp. 124–125).

There are some trade-offs associated with embracing variable-free semantics, though[16]. For instance, one of Heim and Kratzer's key assumptions was that the interpretation function mapped syntactic structures (or their representations, to be more exact) to compositional semantic values. In contrast, for Jacobson, this intermediate step is not only unnecessary, but undesirable. If we get rid of variables in the framework, we don't need to explain the interaction between, *e.g.*, pronouns and their quantificational antecedents by appealing to movement (QR), as standard approaches do. By the same token, if QR was one of the motivations for keeping the LF-level of representation, then it seems like we have reason to get rid of this intermediate level of representation, as well.

Here, Jacobson's reasoning is that LF isn't required to give a story about how compositionality works. And this is consistent with her endorsement of the *Direct Compositionality* thesis. According to Direct Compositionality, the syntactic and the semantic components of the grammar are not autonomous. More specifically, semantic composition isn't computed on the basis of an LF-representation (there is no LF-level to begin with), after the syntax has delivered a full syntactic uninterpreted representation. Instead, the syntax and the semantics work *in tandem* (*cf.* Barker & Jacobson, 2007, pp. 1–5). Constituents receive a semantic interpretation at the time that they combine syntactically—this claim is known as *the local interpretation hypothesis* (1999, p. 120).

In order to implement Direct Compositionality, Jacobson assumes a different view of the syntax. Once we get rid of LF, we also dispense with the notion of hierarchical structures—a core feature in traditional transformational approaches to syntax. Thus, Jacobson's actual implementation will be based on what is known as (Combinatorial) *Categorial Grammar* syntax (CG). According to CG, every linguistic expression has both a semantic interpretation and a syntactic category. Moreover, CG allows Jacobson to deal with composition in terms of immediate adjacency between constituents. For that reason, each lexical entry must contain information about the item's combinatorial possibilities. That is, every lexical entry is a triple that includes phonological information, a meaning, and a syntactic category (under this view, semantic types can be deduced from syn-

[16]We will see in the next subsection that this might turn out to be an inaccurate statement in light of Charlow (2019; 2022)'s findings.

tactic types). Then, linguistic rules operate on these triples and map them to other meaningful linguistic expressions (Jacobson, 2000, pp. 98–100).

So far we have only covered the theoretical motivations for replacing the standard variable-full treatment of pronouns. Then, it can be argued that one could switch to a variable-free approach if those considerations seemed compelling enough. Nevertheless, it should be mentioned that Jacobson (1999) also offers empirical motivations for going variable-free, at least with respect to the domain of pronominal binding. In more recent work, Jacobson herself and other semanticists have extended the variable-free approach in order to account for a wide array of phenomena that includes questions and answers, gapped structures, ellipsis and scope. For more on this, the reader is referred to Charlow (2022) and references therein.

Due to space limitations, we cannot go over the full details of how the compositional semantics of pronouns is modeled under this approach. But the following can be said: in the more general case, we find that the rules of the system allow for composition even when one of the semantic values involves a missing argument because the expression includes a pronoun. Additionally, Jacobson's framework includes *unary type-shifting rules*, that is, rules that are applied to a single linguistic expression at a time in order to change both its meaning and its syntactic category (2000, p. 101). Some of these rules will make linguistic expressions more suitable for combination and others will make binding possible. Let's take (20) as an example (*Willy played his piano*): Here, the expression with a missing argument—*his*—can combine with others (for example, *piano*) by carrying the missing argument over until it is ready to be bound. For this to happen, the verb *played* must undergo a type-shift operation such that its subject—*Willy*—can bind the missing argument.

Now that we have presented the model, it is time to consider how it compares to Heim and Kratzer's treatment of pronominal binding. First, it must be noted that once we go variable-free, instead of having open propositions play the role of semantic denotations of expressions involving free uses pronouns, we have functions of type $\langle e, t\rangle$. In other words, these semantic values will look pretty similar to the semantic values of intransitive verbs. Then, expressions with free pronouns don't express truth-conditional contents until the context supplies a denotation for the pronoun. In addition to this, we can note that the introduction of type-shifting rules leaves us with a simpler semantic value assignment. That is, we don't need to think of, for example, the semantic values of proper names or verbs as being (constant) functions from assignment functions to other entities. Type-shifting rules allow for the modification of meanings as needed.

4.5 Unifying Seemingly Incompatible Models

In the last section, we noted that there are two competing more or less successful ways of modeling the binding of pronouns. Moreover, it was pointed out that there were some fundamental disagreements in terms of their methodological assumptions and ontological commitments. Hence, what the case illustrated is how from the construction and application of two models (variable full and variable free accounts of pronominal binding), we could obtain two different stories about how the syntax-semantics interface might work. But where do we go from here? Do we just accept the incompatibility of the models and choose one way or the other (given our purposes)?

As you might recall from the predicates vs. formulas debate that we have already examined, Kennedy (2014) followed a different approach to problem solving. Instead of endorsing either the predicates or the formulas hypotheses, Kennedy offered a proposal in which he synthesized both claims in order to explain how quantificational binding occurred. And this is just an example of how modelers can make progress *in virtue of* (and not in spite of) models or formal tools being inconsistent. Going back to our pronominal binding debate, we will assess how Charlow (2019; 2022) implements a modeling practice which also aims at unifying elements from the two incompatible models at his disposal.

4.5.1 Case 6: A modular approach to pronominal binding

Charlow (2019; 2022) follows an interesting strategy to model-building in the domain of pronominal binding. The essential idea is that some features of either approach (variable-full or variable-free) can be abstracted out and modularized for future integration with the opposite architecture. For instance, the model presented in his (2019) paper adds certain improvements to the variable-full implementation: lexical entries remain simple (as opposed to relativized to assignment functions), and Predicate Abstraction is implemented without syncategorematicity (*i.e.*, nodes with no meaning). Alternatively, his (2022) offers a way to improve the predictive capabilities of Jacobson's variable-free semantics by borrowing some tools from standard approaches.

We will not review the complete details of the models here. But I will highlight a few ways in which the modular approach is implemented. For instance, his (2019) model, despite being a variable-full one, is expected to deal with assignment functions only when needed. Since we want to keep our lexical semantic values as simple as possible, there's a rule in place for relativizing semantic values to variable assignments if necessary (Charlow, 2019, pp. 4–5). And this rule is inspired by the flexible type-shifting rules that were locally applied in the variable-free framework.

Furthermore, recall that Jacobson treated pronouns as identity functions from individuals to individuals. Hence, Charlow points out that this dependency of individuals (that is, of entities of type *e*) on something else mirrors the dependency of variables (also of type *e*) on assignment functions. For that reason, he suggests that his rule for relativizing semantic values as needed can also be applied in a variable-free system. And this fits with Functional Application already being sensitive to such dependencies in Jacobson's system. For these reasons, Charlow argues that "variable-free semantics can be done using the exact same combinatory tools as the modular treatment of assignment-sensitivity in variable-full semantics" (2019, p. 13). What is more, the modular approach's generality allows the user to implement semantic composition with or without variables.

But now that we have seen Charlow's findings, a question remains: is there still a point to variable-free semantics? If we can just extract (or abstract away) the features of interest and apply them to a variable-full framework, one could think that variable-free semantics loses its appeal. Or maybe that the choice becomes superfluous: if both frameworks can now do the same, why would it even matter which one is endorsed? Alternatively, one who is as worried as Jacobson was about simplifying the model's ontology can take this to show that variable free-semantics is still to be preferred. That is, if the unificatory approach is telling us that the same set of rules, when applied to either the variable-full or the variable-free ontologies (that is, either set of model-theoretic entities) can account for the data, then theoretical virtues might compel us to avoid the introduction of variables to the object language.

In practice, however, there can be some benefits to keeping variables and assignment functions as model-theoretical entities within the system. For example, a modeler could choose to incorporate them if they took these elements to be necessary for making the semantic rules or the lexical entries simpler, or the whole formalization less cumbersome. But recall that Jacobson's theoretical argument was meant to show that standard semantics wasn't getting these benefits from the introduction of variables, either. Still, this is not to deny that there might be other reasons to prefer the variable-full alternative. For example, as Charlow also points out, it is clearer how to extend such an approach to cover phenomena other than pronouns, as opposed to variable-free models (2022, p. 184).

So, what can we say about the modular approach? This modeling practice appears to be inspired by the following reasoning: if models have been successful to some degree at predicting the linguistic data, then we can try to extract some elements from each and integrate them in an attempt to construct a better model. Such an attempt will, as a by-product, reveal how different or incompatible the models actually are (recall that for Jacobson, the models were inconsistent in terms of their ontologies). That is, if the elements can in fact be combined, we will discover that there are less obstacles to integration than we thought (for example, one won't need to commit to Direct Compositionality in order to reap the benefits of going variable-free!).

In the particular case under consideration, we started with two inconsistent models from which semanticists could choose based on a mix of theoretical and data-related considerations. Yet now that a modular approach like Charlow's is available, it is unclear to what degree there is an actual inconsistency between the models and what sort of motivations could lead a semanticist to choose one over the other.

Chapter 5

Epistemological Issues

We have already examined the cases. Is there anything left to say about them? A few lessons about the epistemology of these modeling practices are in order, so that's the aim of the present chapter. First, some general features of the modeling practices or of the models themselves that were evidenced by the case studies will be highlighted. Second, an account of how semanticists are able to make inferences from the construction and application of the linguistic models will be offered: it will be argued that the modeling practices of semanticists are *exploratory* in nature (*cf.* Gelfert, 2016), and that semantic models resemble a class of models used in particle physics: *hypothetical perspectival models* (*cf.* Massimi, 2019). Finally, a consequence of the present proposal will be discussed: that the inferences that semanticists draw from their modeling practices are hypothetical, since their modeling practices consist in assessing the plausibility of different accounts of linguistic phenomena at the syntax-semantics interface.

5.1 Features of Modeling Practices in Semantics

5.1.1 They require inter-model interactions

A recurrent finding across our case studies is that semantic models are not built in isolation. As semanticists construct them, they must import some commitments or tools from other models which are already available. This is especially relevant for deciding the intended domain of application of the new model.

As we have seen, even a textbook model like Heim and Kratzer (1998)'s needs to interact with other models: since it is a simplified implementation of a compositional approach (what they call the "Fregean program", *cf.* p. 13), and Compositionality is a property of a semantics in relation to a syntax, the model has to incorporate some elements from syntactic theory. There is no other way to built a semantics which is relative to a syntax. Hence, when Heim and Kratzer construct the model, they import some assumptions from the generative grammar theory such as that phrase structure trees are syntactic representations of linguistic expressions, and the claims from the inverted y-model of grammar (which stipulates the Logical Form level of representation). But they also need to assume a particular syntactic analysis of the linguistic constructions. The same can be said of other models, regardless of whether they do it more or less explicitly. Hence, we can say that modeling practices in formal semantics involve building something on top of syntactic models.

But modeling practices can also involve interactions between semantic models. First, a model can be an "extension" of one which is already at use in some specific domain, especially if the aim is to explore the suitability of a particular tool in said domain. Then, the procedure is to integrate the tool with an already existent model. For example, part of the process that Heim (1997) followed involved assessing whether 1-place predicates or formulas could be integrated to Rooth (1992)'s focus semantics in order to account for Kennedy (1994)'s data. Second, one can use other models as reference in order to construct one which omits as many elements or assumptions as possible, in order to simplify the already existent models or give an alternative account of the phenomena under study. For instance, Jacobson (1999)'s variable free semantics was constructed with the aims of both simplifying the ontology of variable-full approaches, and departing as much as possible from the original assumptions (recall that the model was a way of implementing the Direct Compositionality thesis). Or, one can attempt to integrate elements from two seemingly incompatible models in order to offer a new model or an improved version of the previous ones, like both Kennedy (2014) and Charlow (2019; 2022) did.

5.1.2 Ontological worries are the exception rather than the rule

When we briefly reviewed the theoretical motivations behind Jacobson (1999)'s proposal, it was noted that Jacobson's main complaint was that under variable-full approaches, variables were introduced to the object language in order to account for binding phenomena. In practice, this meant that the semantic value assignment was more complex than it appeared at first sight—*viz.*, every semantic value was relativized to assignment functions. As a consequence, the construction of the variable-free semantics model was an attempt at figuring out whether these extra entities (variables and assignment functions) played any crucial role in the model, or whether they were mere representational artefacts that the models could dispense with. So, it was very clear that the choices that were made as the model was constructed were influenced by matters of ontological commitment.

However, this doesn't seem to be the case for most of the models. After all, it is a fact that most models still rely on variables and assignment functions, even if only for pragmatic reasons. For example, a variable-full approach can be preferred if the goal is to integrate models, given that most modelers keep making use of variables. Moreover, recall that giving up on variables also requires a different implementation of the syntax-semantics interface, so this doesn't just limit integration between semantic models—it also affects the way that semantic models integrate with syntactic models.

Yet, this is not to say that questions about ontological commitment might not arise sporadically. For example, Heim (1997) does entertain the question of whether, assuming the formulas hypothesis is to be endorsed, this would point to the non-existence of 1-place predicates, even beyond the domain of VP-ellipsis: "...the stronger hypothesis is *prima facie* more interesting: the reason why [we can't integrate Rooth's account with] predicates is simply that no such predicates exist. Evidently, we would hope to eventually find confirmation for this claim from other, superficially unrelated, domains of evidence" (1997, p. 216, original emphasis). Still, the default position appears to be that the construction and application of the models are not going to settle the matter.

In general, the modeling practices of semanticists suggest that there is an instrumentalist conception of the components found in the models, such as the semantic value assignment (or the semantic values themselves) and the semantic rules. There seems to be a preference for tools which modelers perceive to be doing a better job at performing their intended roles. But there isn't an explicit consensus on what "doing a better job" might mean. Is it just a matter of getting the right predictions? Could it be their potential for integration/achieving a uniformity across models?

5.1.3 The choice of tools is contextual

A related regularity to be identified across the diversity of modeling practices is that, despite the interest in recycling some elements of previous models, it still acknowledged that the tools (and models) have a limited domain of applicability. For example, from the construction of a model for VP-ellipsis which implements the formulas hypothesis, it doesn't follow that we need to endorse a more comprehensive view of the syntax-semantics interface in which we commit to formulas or open propositions being the semantic values of certain linguistic expressions. Formulas might do a good job, but they only do so in the context of a model which incorporates certain assumptions or elements, such as a particular account of VP-ellipsis. Then, the usefulness of these tools depends on the way in which they can integrated within the model. If it is later found that formulas are not useful for being the semantic values of verb phrases within another model, this won't be conclusive evidence that the former model should be rejected.

Hence, the choice of tools is contextual in two senses: first, in the context of the model itself. As we have seen in our case studies, decisions can be more or less constrained by previous decisions regarding commitments or other elements that are already part of the model. That is why we can't expect a specific device to be suitable for integration with other models, even if they are intended to play the same role. Second, the choice is contextual regarding the domain of applicability of the model. That is, the fact that the tool in question does a good job within a model which accounts for some phenomenon doesn't mean that the tool would also be useful as an element of a model which deals with a different phenomenon.

5.1.4 Most choices are made on pragmatic grounds

Given the previous features, it can be expected that modelers will be faced with many decisions for which linguistic data will not constitute decisive evidence in favor of some option. That is, even if we can test the models by applying them to the analysis of a variety of linguistic constructions, it can still occur that alternative models are as good at predicting the data. So, how do semanticists decide?

Some of these decisions will be motivated by matters of convenience, by appeal to theoretical virtues, or a mix of both. As an illustration, recall that Heim and Kratzer acknowledge that approaches which do not rely on movement can be well suited to solve the problem of quantifiers in object position. However, some of them require the modeler to embrace lexical ambiguities for quantificational determiners, or involve the introduction of additional semantic rules (type-shifting rules). On one hand, the first option is unappealing given that it avoids giving a compositional account, and instead re-assigns the problem to the domain of lexical semantics. And, by doing so, it over-complicates the lexicon in the model. On the other, the second option might require further changes in the system, especially in relation to the implementation of the interaction between the syntax and the semantics modules of grammar. Then, there is a sense in which the choice is up to the preferences of the modeler or of the wider community.

5.2 Scientific Modeling Revisited

Our main objective for the remainder of this chapter is to find (or come up with) a tentative explanation for why these modeling practices exhibit the features we just identified, and ultimately, for how semanticists are able to make progress by building and applying models in the way they do. In order to do so, we will consult the philosophy of science literature for inspiration, and then try to apply some relevant insights to the assessment of the modeling practices under consideration.

5.2.1 Exploratory Modeling

As you may recall from Chapter 2, the philosophical study of scientific practices led to the identification of different functions that models (or modeling practices) could serve, such as representation, explanation, or intervention. More recently, Gelfert (2016) has noted that there is a distinctive kind of modeling practices which is oriented toward exploration. In these cases, the exploratory uses of models might give the impression that modeling is closer to experimentation than to theorizing proper (Gelfert, 2016, p. 5). Then, models can serve an exploratory function when their design, construction and/or application allow the modeler to figure out which tools and techniques are suitable for modeling. In this sense, the modeler appears to be "experimenting" with modeling techniques in order to master and understand them (2016, p. 41). This ultimately contributes to the goal of successful future model-building.

But these practices are not only useful for exploring modeling techniques. Given that exploration doesn't require full clarity about what the object of scientific research actually is, modelers can design models which are "in search of a target system" (*cf*. Gelfert, 2016), or that involve hypothetical target systems or no target system at all (*cf*. Massimi, 2019).

After assessing different case studies, Gelfert identifies four functions which exploratory models can serve:

(i) They can be the starting point for future inquiry that might give rise to more realistic and complex models (2016, pp. 84–85).

(ii) They can show that potential representations of some target phenomenon can be achieved by adopting a particular methodology or that some modeled mechanism resembles the behavior found in an actual system (2016, pp. 85–86).

(iii) They can contribute potential explanations (2016, pp. 87–93).

(iv) They can help in assessing the suitability of a target (2016, pp. 93–94).

Since the contexts in which exploratory models are applicable are those where a fully-articulated theory is missing, oftentimes models perform many of these roles at once. However, this doesn't mean that the models are theory-free or that they are not responsive to observation.

A possible suggestion, then, is to think of the practices of formal semanticists as involving the construction of models which perform the roles identified by Gelfert. If this is correct, it certainly gives us reason to think that the modeling practices of semanticists are primarily exploratory. A further possibility that we will consider in the next section is that semantic models are not constructed so as to represent an actual target system.

5.2.2 Hypothetical perspectival models

After examining some examples of modeling practices in the domain of the natural sciences, Massimi (2019) added two more functions to Gelfert (2016)'s list:

(v) Exploratory models can deliver knowledge of causal possibilities.

(vi) Exploratory models can deliver knowledge of objective possibilities for hypothetical entities.

These functions correspond (respectively) to the two types of models that Massimi identifies: *targetless fictional models* and *hypothetical perspectival models*. By using models like these, it is argued, scientists are able to extract *modal knowledge* about what is possible. Hence, exploratory models can still contribute to scientific progress despite not having a main representational function (2018, p. 349). Here, we will focus on the second kind for the purposes of our discussion.

In virtue of being exploratory models, hypothetical perspectival models do not count as successful in virtue of being accurate representations that map onto a specific target system. Instead, their success depends on how useful they are for exploring and carving out the space of possibilities. Since what is being modeled are objective and epistemic modalities (2018, p. 339), we can still say that the models are *about something*, and this aboutness is enough to ground their representational capacities (2018, p. 349). As a way of illustration, let's consider the modeling practices involved in Beyond Standard Model (BSM) physics.

In contemporary high-energy physics, perspectival models are routinely used to find new particles known as supersymmetric or SUSY particles. If these particles were to exist, this would force particle physics to go beyond what is known as the Standard Model. For this reason, physicists employ models—*BSM models*—that omit as many assumptions from the Standard Model as possible. Their purpose is to enable the finding of entities with features that are incompatible to those that the Standard Model ascribes. In this sense, BSM models are independent of theory. And since SUSY particles are not known to exist by the time the models are built, these models are not descriptions or representations of data either (Massimi, 2018, pp. 349–350).

The key detail about Massimi's example is that BSM models are meant to deliver knowledge about *objective possibilities*—that is, possibilities concerning the very existence of the model's hypothetical target system. In the particular case of particle physics research, the conceivability at play is constrained by laws of nature. These laws fix the "boundaries" within which inferences about the hypothetical target system can take place (2019, p. 875). Hence, these models deliver a type of *how-possible inferences* of the following form: *If some law-bounded physically conceived scenarios were to be possible, how could some entities (SUSY particles) be possible?* As has been mentioned, the possible scenarios depart from what the Standard Model takes to be actual, so evidence must play a role in constraining these live objective possibilities. The how-possible inferences are constantly being refined by introducing new data, since the latter constrains the antecedent of the inference. In practice, this means that some regions in the space of possibilities will get excluded from future searches (2019, pp. 876–878).

Moreover, for Massimi, thinking about these models as *perspectival* is crucial for understanding how in any scientific context we can have many models at play that are performing their exploratory function jointly. This is possible since each model provides only a partial account of the phenomenon at stake (2018, p. 350). Thus, despite the fact that each of them is strictly speaking *false*, their falseness is not as problematic as it would be for models whose representational function is essential for scientific research (Gelfert, 2016, p. 41). In the early stages of inquiry, the question of which model is the one that accurately represents the target system seems premature. After all, scientists might lack the resources for appropriately assessing truthfulness or empirical adequacy (Gelfert, 2016, p. 84). But more importantly, such a concern might not even make sense given that it is still an open question which is the target system and whether it is actual.

5.3 An Epistemological Account

Ultimately, it will be suggested that an appropriate way to understand the modeling practices of formal semanticists is to think of them as exploratory, given that the function of semantic models is to deliver modal knowledge about a hypothetical target system—in this case, the syntax-semantics interface. But first, it will be argued that the modeling practices under consideration are not attempts to model the semantics of linguistic constructions by implementing theories, even if the modelers are endorsing certain theoretical commitments.

5.3.1 Semantic models are not models of theory

Under the *theory-driven view of models* (*cf*. Cartwright et al., 1995), models are supposedly *derived* from a coherent theory. For example, models which apply Newtonian mechanics to the study of some empirical phenomenon are thought to be the perfect candidates for fitting this characterization. The purpose of these models is to find confirmation for the theory by introducing correction terms which de-idealize more abstract models (1995, p. 142). More generally, this view assumes that theories provide modelers with the fundamental building blocks that are needed to construct models. Accordingly, models must import commitments from the underlying theory[1].

Yet, this doesn't seem to be the case in formal semantics. For instance, in the case of Heim and Kratzer's semantic framework, there doesn't seem to be an established theory from which the model is derived. And for sure this isn't a defect of the model—the history of science has already shown that a fully-developed theory doesn't need to be available for scientists to start building models[2]. But even if there were such a theory, thinking of semantic models as theory-driven would still be misguided. Semantic models are tools for figuring out which commitments semanticists want to eventually include as part of their future models and theories. And, as the debate between variable-full vs. variable-free approaches showed, matters related to foundational commitments are far from settled yet.

[1]Morgan and Morrison (1999) take this approach to be exemplified by the semantic view's understanding of theoretical models as "models of theory".

[2]Of course, whether or not there was in fact an available fully-developed theory by the time Heim and Kratzer devised their textbook system can be up to debate. In any case, the point made here isn't totally dependent on this particular claim.

Moreover, as Morgan and Morrison have argued, "theory does not provide us with an algorithm from which the model is constructed and by which all modelling decisions are determined" (1999, p. 16). And this is what we have seen in practice. Frege's Conjecture for sure is an important guiding assumption for constructing semantic models, but as the case studies have evidenced, it would not be productive for modelers to implement it in a literal way. Moreover, it can't be denied that Heim and Kratzer introduce rules other than Functional Application based on pragmatic considerations, and this is not an illegitimate move when building models. Hence, the case of formal semantics is consistent with the following picture of scientific modeling: models are independent from theory in virtue of their construction (Morgan & Morrison, 1999, p. 14), because theoretical assumptions are only tools for scientific practice (Suárez & Cartwright, 2008, p. 43).

5.3.2 Semantic models deliver modal knowledge

As we just have seen, semantic models are not built by applying a theory. Even if semanticists rely on some theoretical assumptions such as the compositionality of meaning, an articulated and specific notion of compositionality is still absent. This is why there are models such as Heim and Kratzer (1998)'s and Jacobson (1998)'s which vastly differ in their implementations of the principle. Since the compositionality of a semantic value assignment is dependent on the interaction of the syntactic and the semantic components of the grammar, their frameworks differ in their choice of syntax, too. While Heim and Kratzer assume a transformational grammar approach, Jacobson adopts a categorial grammar syntax. Hence, even though semanticists are not yet sure about what the syntax actually looks like (*i.e.*, which particular syntactic models are representationally accurate), they can still build models that import some presuppositions from theories of syntax.

In light of this picture, one might wonder how it is that semanticists can still make progress and derive knowledge from their modeling practices. If the practices involved in the construction of the aforementioned models are exploratory in nature, then they serve the purpose of figuring out which modeling techniques to adopt. But more importantly, they serve the purpose of assessing the suitability of the syntax-semantics interface as a target-system for their models. But how does this occur in practice?

Hypothetical inferences appear to serve an important role in the modeling done within formal semantics. For example, recall that Jacobson (1999) had pointed out that it is in virtue of introducing variables to the system that we end up being committed to model-theoretic entities like assignment functions. Hence, what the construction of the Heim and Kratzer model shows is how, if we were to generally treat pronouns as variables, we would be forced to adopt a subsequent commitment. Then, the construction of Jacobson's variable-free semantics model is an attempt to explore whether it would be possible for the semanticist to reject the theoretical commitments of standard variable-full approaches and adopt others instead. In a nutshell, this methodological approach would appear to involve the alteration of the set of conceived scenarios included in the antecedents of relevant how-possible questions. For example, the modeling practice in question can be characterized as attempting to answer the question of how pronominal binding can be possible if Direct Compositionality and local interpretation were possible (*i.e.*, no LF-level, and no autonomy of the syntax), and no variables were introduced in the object language. Then, in order to give a hypothetical account of pronominal binding, the model gives us an alternative story of how the syntax-semantics interface works.

Still, it can always happen that given a set of assumptions about the syntax-semantics interface, the whole space of possibilities gets ruled out. In that case, one would have reason to think that pronominal binding just isn't possible in these conceived scenarios. Furthermore, one can extrapolate from this and infer either that variables and assignment functions are necessary to explain the phenomena under consideration or that they are not. But notice that theoretical assumptions don't do all of the work. Part of this exploratory process involves testing whether the predicted truth-conditions that each model delivers match the truth-conditions that native speakers associate with sentences. Moreover, as new data are taken into consideration, some conceived scenarios can be ruled out. For example, once Heim and Kratzer examine the semantics of expressions involving the modification of nouns by adjectives or prepositional phrases, we rule out scenarios in which the syntax-semantics interface is such that Frege's Conjecture is literally true (1998, p. 65). Something similar occurs once they take into consideration the data from quantificational DPs, since now compositionality must be more flexible than before.

But matters can be even more complex than that. After all, Charlow's modularity approach reveals the following findings:

> . . . both sorts of theories can readily exploit flexible compositional architectures (though variable-full theories typically do not), and both sorts of theories can have a rule for binding oriented around scope (though Jacobson's does not). The status of object-languages variables, by contrast, is definitional and, therefore, non-negotiable. Nevertheless, the two kinds of theories can be stated in ways that make them seem much more alike than we might have initially supposed. This suggests in turn that one may

not need to be a variable-free semanticist in order to reap the empirical insights of variable-free semantics[3]. (2022, p. 18)

To put it in more familiar terms, what Charlow is suggesting is that by introducing some assumptions in the antecedent of their how-inferences, the variable-full and variable-free modelers are not in fact ruling out as many scenarios as they might have initially thought. For example, the assumptions in the variable-full system do not rule out that semantic composition is possible when we have simple lexical entries and type-shifts are introduced only when needed. Likewise, endorsing Direct Compositionality, Categorial Grammar syntax, and the hypothesis of local interpretation doesn't rule out scenarios where binding occurs in virtue of scope relations. But what does get ruled out beforehand is the introduction or omission of variables in the object language, since that was already part of the assumptions.

The lesson to be extracted here is that the usefulness of semantic models lies in that they serve to reduce the space of objective possibilities in a manner that is constrained both by working assumptions and linguistic data. Even though there are some theoretically motivated assumptions involved in the modeling practices of semanticists, semantic models are not fully theoretical in nature. There is no fully-developed theory to be applied yet. Moreover, linguistic data plays a crucial role in that it serves to reduce the space of open possibilities: if the introduction of a compositional rule makes the wrong predictions, then that is reason to doubt that it can capture some mechanism by which the syntax and the semantics interact. So, all in all, semantic models appear to also be useful tools for deciding which theoretical commitments to include in future model-building and theorizing.

[3]"Flexibility" in this quote is meant to stand in opposition to the standard semantics approach's *uniformity*. As we saw before, everything is underlyingly relativized to assignment functions (the semantic values and rules) because the model-building practice generalizes to the worst-case scenario.

Chapter 6

Final Remarks

After assessing some examples of modeling practices involved in formal semantics, a story of how semanticists can derive knowledge from these practices was offered. The modeling practices of semanticists are exploratory in nature, and they can serve many purposes such as figuring out the appropriate tools that they will incorporate in their models, being the starting point for building more complex models, assessing the suitability of a target system such as the syntax-semantics interface, and providing modal knowledge about what this target system might look like.

But there are many questions in the vicinity which unfortunately, we didn't get the chance to address. Since the account offered here is a tentative one, further refinements could be expected once the following issues are also considered:

What is the object of study of formal semantics? Here, my account has provisionally suggested that it is the syntax-semantics interface. After all, we have been talking about compositional semantics, and compositionality (as we reviewed before) is a property of a semantics relative to a syntax. But someone could be dissatisfied with this response, given that the syntax-semantics interface and compositionality appear to be theoretical constructs. Shouldn't formal semantics aim to study something that is "out there in the world"? But what could this be? Semantic competence? Mental/cognitive states? Regardless of whether that is what semantics *should be doing*, the approach followed here has started from the assessment of the modeling practices of linguists. And there isn't much involved in these modeling practices which would suggest a more satisfactory answer. Of course, this isn't to deny that whatever semanticists are learning from building and applying models isn't revealing something about semantic competence, in a broad sense.

In light of what has been said, how should we go on about interpreting the models or products of linguistic theorizing? Here, we have offered some reasons to think that they should not be interpreted as being ontologically committed to the entities that the elements of the models might appear to designate. Yet the account presented here cannot give a more informative answer regarding this issue either.

Are semantic models providing rough explanations of anything? Can explanation rely on hypothetical systems/entities? Can modal accounts be explanatory? There is a sense in which they might be. Here we have extensively reviewed different "accounts" offered by semanticists. So, semanticists must be trying to account for something, and the claims involved in their proposals might be explanations of that something. But it is hard to answer this question without having a stronger intuition of what the object of study of the discipline or of the models actually is. Could they be aiming to explain how semantic composition in fact happens? It is doubtful, since it is usually acknowledged that compositionality is a methodological assumption. Perhaps a better way to frame the original concern is the following: could semanticists be attempting to say something beyond the world of the model, or are they mostly explaining what happens inside the model? Moreover, if these models are not providing explanations of phenomena, then how else could we characterize the findings of semanticists? Is semantics just about finding generalizations?

What are linguistic data? The present account is certainly silent about this matter, but any answer to this question should presumably illuminate the aforementioned issues. If the data are empirical, then the models could be rough attempts at explaining them. But in order to establish this, we would need to have a stronger sense of what counts as linguistic data for the purposes of modeling and confirmation, and whether the data are raw linguistic intuitions, processed linguistic intuitions, or something else.

Ball, D., & Rabern, B. (Eds.). (2018). *The science of meaning: Essays on the metatheory of natural language semantics*. Oxford University Press.

Barker, C., & Jacobson, P. I. (Eds.). (2007). *Direct compositionality*. Oxford University Press.

Bokulich, A. (2011). How scientific models can explain. *Synthese*, *180*(1), 33–45. https://doi.org/10.1007/s11229-009-9565-1

Boumans, M. (1999). Built-in justification. In M. S. Morgan & M. Morrison (Eds.), *Models as mediators: Perspectives on natural and social sciences* (pp. 66–96). Cambridge University Press.

Burgess, A., & Sherman, B. (Eds.). (2014, October). *Metasemantics: New essays on the foundations of meaning*. Oxford University Press.

Carnie, A. (2013). *Syntax: A generative introduction* (3rd ed.). Wiley-Blackwell.

Cartwright, N., Shomar, T., & Suárez, M. (1995). The tool-box of science: Tools for the building of models with a superconductivity example. *Poznan Studies in the Philosophy of the Sciences and the Humanities*, *44*, 137–149.

Charlow, S. (2019). *A modular theory of pronouns and binding*. https://ling.auf.net/lingbuzz/003720

Charlow, S. (2022). On jacobson's "towards a variable-free semantics". In L. McNally & Z. G. Szabó (Eds.), *A reader's guide to classic papers in formal semantics: Volume 100 of studies in linguistics and philosophy* (pp. 171–196). Springer International Publishing. https://doi.org/10.1007/978-3-030-85308-2_10

Collins, J. (2020). Semantic and syntactic intuitions: Two sides of the same coin. In S. Schindler, A. Drożdżowicz, & K. Brøcker (Eds.), *Linguistic Intuitions: Evidence and Method*. Oxford University Press. https://doi.org/10.1093/oso/9780198840558.003.0006

da Costa, N., & French, S. (2000). Models, Theories, and Structures: Thirty Years On. *Philosophy of Science*, *67*(S3), S116–S127. https://doi.org/10.1086/392813

Devitt, M. (2010). Linguistic intuitions revisited. *The British Journal for the Philosophy of Science*, *61*(4), 833–865. https://doi.org/10.1093/bjps/axq018

Frege, G. (1997). On Concept and Object (1892). In M. Beaney (Ed.), *The Frege reader* (pp. 181–193). Blackwell Publishers.

Frigg, R. (2022). *Models and theories: A philosophical inquiry*. Routledge. https://doi.org/10.4324/9781003285106

Gelfert, A. (2016). *How to do science with models*. Springer Berlin Heidelberg.

Giere, R. N. (2004). How models are used to represent reality. *Philosophy of Science*, *71*(5), 742–752. https://doi.org/10.1086/425063

Glanzberg, M., & King, J. C. (2020). Binding, Compositionality, and Semantic Values. *Philosopher's Imprint*, *20*(2). http://hdl.handle.net/2027/spo.3521354.0020.002

Heim, I. (1997). Predicates or Formulas? Evidence from Ellipsis. *Semantics and Linguistic Theory*, *7*(0), 197–221. https://doi.org/10.3765/salt.v7i0.2793

Heim, I., & Kratzer, A. (1998). *Semantics in generative grammar*. Blackwell.

Hughes, R. I. G. (1997). Models and representation. *Philosophy of Science*, *64*(4), S325–S336.

Jacobson, P. (1998). Antecedent Contained Deletion and Pied-Piping: Evidence for a Variable-Free Semantics. *Semantics and Linguistic Theory*, *8*(0), 74–91. https://doi.org/10.3765/salt.v8i0.2800

Jacobson, P. (1999). Towards a Variable-Free Semantics. *Linguistics and Philosophy*, *22*(2), 117–184.

Jacobson, P. (2000). Paycheck Pronouns, Bach-Peters Sentences, and Variable-Free Semantics. *Natural Language Semantics*, *8*(2), 77–155. https://doi.org/10.1023/A:1026517717879

Jacobson, P. (2014). *Compositional semantics: An introduction to the syntax-semantics interface*. Oxford University Press.

Keller, E. F. (2000). Models Of and Models For: Theory and Practice in Contemporary Biology. *Philosophy of Science*, *67*(S3), S72–S86. https://doi.org/10.1086/392810

Kennedy, C. (1994). *Argument Contained Ellipsis* (LRC-94-03). https://semantics.uchicago.edu/kennedy/docs/ace.pdf

Kennedy, C. (2014). Predicates *and* Formulas: Evidence from Ellipsis. In L. Crnič & U. Sauerland (Eds.), *The Art and Craft of Semantics: A Festschrift for Irene Heim* (pp. 253–277). MIT Working Papers in Linguistics.

Massimi, M. (2018). Perspectival Modeling. *Philosophy of Science*, *85*(3), 335–359. https://doi.org/10.1086/697745

Massimi, M. (2019). Two Kinds of Exploratory Models. *Philosophy of Science*, *86*(5), 869–881. https://doi.org/10.1086/705494

Morgan, M. S., & Morrison, M. (Eds.). (1999). *Models as mediators: Perspectives on natural and social sciences*. Cambridge University Press.

Pagin, P., & Westerståhl, D. (2010). Compositionality I: Definitions and Variants. *Philosophy Compass*, *5*(3), 250–264. https://doi.org/10.1111/j.1747-9991.2009.00228.x

Rooth, M. (1992). Ellipsis redundancy and reduction redundancy. In S. Berman & A. Hestvik (Eds.), *Proceedings of the Stuttgart Ellipsis Workshop*. Universitäten Stuttgart und Tübingen in Kooperation mit der IBM Deutschland.

Schindler, S., Drożdżowicz, A., & Brøcker, K. (Eds.). (2020). *Linguistic Intuitions: Evidence and Method*. Oxford University Press. https://doi.org/10.1093/oso/9780198840558.001.0001

Suárez, M., & Cartwright, N. (2008). Theories: Tools versus models. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, *39*(1), 62–81. https://doi.org/10.1016/j.shpsb.2007.05.004

Szabó, Z. G. (2012). The case for compositionality. In M. Werning, W. Hinzen, & E. Machery (Eds.), *The Oxford Handbook of Compositionality* (pp. 64–80). Oxford University Press.

Weisberg, M. (2007). Who is a Modeler? *The British Journal for the Philosophy of Science*, *58*(2), 207–233. https://doi.org/10.1093/bjps/axm011

Yli-Vakkuri, J. (2013). Propositions and Compositionality. *Philosophical Perspectives*, *27*(1), 526–563. https://doi.org/10.1111/phpe.12025