Characterizing Kinds

A Semantics for Generic Sentences

Matt Teichman

October 27, 2015

Acknowledgments

There is no question that Melina Avery deserves my heartfelt thanks for living through this project with me over the past few years. We met as it was in its formative stages, and I could not have asked for a more thoughtful or sensitive partner. It is no trivial undertaking to be in the company of someone who is writing a dissertation, and I must say she pulled it off with impeccable grace. If better people write better dissertations, then you primarily have her to thank for making me into a better person.

I am also especially eager to thank my parents, Harold Teichman and Susan Davis, for far more than can be fit into an acknowledgements section. Throughout my early life, at great personal and financial cost to themselves, they made it their top priority to secure for me the best education they possibly could. And throughout my adult life, they have regularly gone beyond the call of duty when it came to nurturing my varied intellectual interests. When I developed an interest in philosophy, right at the tail end of my undergraduate career, it was my father who told me about Frege's work, bought me copies of Wittgenstein's *Tractatus*, Anscombe's *Intention*, and Kripke's *Naming and Necessity*, and urged me to take my first logic course. When I first developed a serious interest in Wittgenstein, it was he who patiently dissected the rule following argument, the private language argument, and the remarks on pain for me. When I took my first ever graduate seminar in philosophy, then as an outsider to the discipline, I frequently had to phone home for help, but I was able to get myself up to speed remarkably quickly, thanks to his formidable erudition in philosophy, mathematics, the German language, physics, poetry, and countless other subjects. For this unofficial second education, which formed the perfect complement to my formal education, I am profoundly grateful.

During those late undergraduate years, I was a dual major in linguistics and cinema studies, about to pursue graduate work in cinema studies on avant-garde film. My first ever philosophy course, on aesthetics, was with Christopher Williams, who did an excellent job of introducing me to the art of drawing subtle logical distinctions between superficially similar philosophical positions. It was Molly Diesing, Junko Shimoyama, and Dorit Abusch who gave me my first introduction to the philosophy of language as an undergraduate, and I am sure that my take on the field has been strongly influenced by their respective ap-

proaches. And it was Greg Salmieri who introduced to me both to ancient philosophy and to the history of philosophy more broadly. Although it may not be immediately apparent from the direction my research has taken, my overall philosophical world picture derives much of its inspiration from the ancient Greeks. Even today, my views about how philosophical activity integrates into the rest of human life take Greg's important work in this area as their point of departure.

As my graduate work in cinema studies progressed, an interest in systematically characterizing the differences between film interpretation and linguistic interpretation led me to draw more and more on contemporary scholarship in philosophy, and before I knew it, my research had evolved into something that it only really made sense to pursue in a philosophy department. When I finally decided to make the transition from one academic discipline to another, John McDowell, Lucy Fischer, Nicholas Rescher, Greg Salmieri, Cathy Legg, Ronald A. T. Judy, John Whitman, and Abigail Cohn were particularly forthcoming with practical advice on how best to do so. It was on McDowell and Rescher's strong recommendation that I ended up choosing the University of Chicago, and it was thanks to the superlatively warm support of my cinema studies supervisor, Lucy Fischer, that I was able to transition smoothly.

Upon my arrival in Chicago, I was instantly made to feel at home in a vibrant community of first-rate graduate students, each of whom had an impressive combination of diverse philosophical interests and exhibited an desire to share their ideas with others. During my first few years in the program I had a roommate named Mark Hopwood, who it could be said spent much of that time effectively training me to live. Somehow, over scores of extensive powwows that can only be characterized as epic, and despite his downright comical lack of interest in every area of philosophy that I work on, he managed to teach me a nontrivial proportion of what I currently know about the field, along with a substantial proportion of what I currently know about politics, history, and religion. How I managed to talk him into creating a podcast with me is anyone's guess.

During this same period, I was also especially lucky to have access to Aidan Gray and Peter Klecha, under whose intellectual mentorship I was introduced to recent work in the philosophy of language, and with whom I spent numerous hours considering various hypotheses as to the explanatory purport of semantic theories. The deep connections between my previous philosophical interests and this exciting new material were immediately apparent, and before long, I was a regular attendee of the Semantics and Philosophy of Language workshop. Today, Aidan and Peet are still the first people I turn to when I feel my understanding of human language needs to be deepened.

My decision to specialize in the philosophy of language became definitive after I took Chris Kennedy's introductory graduate seminar on semantics. Although I had been exposed

to basic compositional semantics as an undergraduate, most of the content of this course was new to me. Chris has an unbelievable degree of spontaneous energy in the classroom; he thrusts his students right into the thick of the most difficult foundational debates in the field, but in such a way as to make them seem somehow tractable. Each class session of his is impeccably organized, while also retaining a strong improvisational character. I recall one instance in which he asked: 'Which theory of syntactic movement would you rather use? The trace or the copy theory? It's all the same to me.' Someone requested that we adopt the copy theory, and the entire trajectory of the rest of the course was instantly reconfigured, including the content of all subsequent homework exercises. I don't think I've ever encountered that level of extemporaneity in a class ever since.

Shortly thereafter, I sat in on two excellent formal semantics seminars taught by Itamar Francez, wherein we covered such fascinating topics as modality, concealed questions, dynamic semantics, and intensional transitives. One year later, he graciously offered to guide me through an unofficial independent study in the canonical works of that tradition. My entire dissertation topic sprung into existence one afternoon as a result of a conversation I had with him. And a year after that, thanks to the generous letter-writing efforts of Gabriel Lear, Chris Kennedy, and Kevin Davey, I was awarded a Fulbright IIE grant to study logic and semantics in the Netherlands.

It would be difficult to overstate how formative this period was for me. The Institute for Logic, Language, and Computation is incomparable to any other academic environment I've come across. The disciplinary demarcations that dictate the way nearly every intellectual problem is conceived of in the US barely exist there. At the ILLC, researchers in mathematics, philosophy, linguistics, cognitive science, computer science, and engineering think nothing of publishing, teaching graduate seminars, or putting on conferences in each other's fields. In the US, interdisciplinarity typically involves a professor in one department co-authoring a paper with a professor in a different department. But in Amsterdam, it runs much deeper. Most of the work done there is impossible to fit into the familiar academic taxonomy; one has the sense that it could be published in any one of five fields, modulo small variations in terminology. I am reminded of Michel Foucault, who was told by philosophers that he was trying to pass history off as philosophy, and by historians that he was trying to pass philosophy off as history. In my view, to be on the receiving end of such accusations is as sure a sign as any that one must be doing something right.

My entire stay in Amsterdam was a success, but my meetings with Frank Veltman, my faculty sponsor while abroad, are particularly worthy of mention. Over a period of one academic year, he helped me solidify my sense of how the standard rubric of topics in mathematical logic hooks up with the standard rubric of topics in formal semantics. And although my inclinations tend in a metaphysical rather than epistemic direction, he has probably had more of an influence on the way I conceive of normality than anyone else. In

a less official capacity, I have Facundo Carriero, my roommate for the first four months of my visit, to thank for teaching me most of what I currently know about logic, programming, and how to set up an operating system. Though he would surely insist otherwise, I think it's clear that he is a philosopher at heart.

I returned to Chicago in the Spring of 2012, highly energized for dissertation writing. Billy Cobham has said that the *Bitches Brew* sessions with Miles Davis were an exciting new challenge for him, because whereas previously he'd shared a recording booth with no more than one piano genius at a time, now he was expected to keep up with three. That is not far off from how I felt during my dissertation committee meetings, where the amount of sheer mental firepower in the room was readily palpable. Often, I learned the most just by sitting back and listening to everyone converse. I view my dissertation as a collaborative endeavor, and it is certainly the case everything written in the coming chapters bears the indelible imprint of their thoughtful contributions (though the usual paradox-inducing proviso of course applies).

Jason Bridges was an exemplary supervisor in every possible respect. I originally hoped to be able to bring him on board because I had heard great things from other students about his advising, and because I thought he would have a fresh perspective on the philosophical implications of the semantic analysis I wished to develop. I was not wrong. In part due to his breadth of expertise, stretching through the far reaches of M&E, the philosophy of language, and the philosophy of mind, but mostly due to his ability to situate any given philosophical thesis instantly with respect to every other possibility in logical space, every comment I received from him either raised a worry or brought out a connection that I would never have thought of on my own. I quickly learned to expect the unexpected, and whenever we finished one of our meetings, I was immediately propelled to the library to read up on some topic whose relation to my dissertation research I could not have foreseen. One of the finest gifts a supervisor can give their students is assistance in broadening the scope of their work, temporarily lifting them out of their own headspace. And indeed, this work was broadened by his input to a far greater extent than I would have imagined possible going in. It also bears mentioning that Jason's personal demeanor during these one-on-one interactions was delightful; the seriousness of his feedback was matched only by the casual levity of his address. Few students have the privilege of working with a supervisor who puts them instantly at ease.

Chris Kennedy was a natural choice for supervisor of this project not only because of its strong natural language semantics component, but also because his classes are what pushed me in that direction to begin with. Over the past decade, he has almost certainly gotten more philosophy students interested in the philosophy of language than anyone else on campus. He is renowned throughout the Linguistics and Philosophy departments for his ability to soak up the essence of a view outside his area of research, immediately grasp

how it works and why it's interesting, and give an off-the-cuff response that's as deep and polished as a good peer review piece. But one feature of Chris' inimitable style of thinking that receives less attention is its tai chi-like flexibility. He takes positions opposed to his own so seriously that he is happy to turn on a dime and defend them unhesitatingly, the moment the evidence begins to tip in their favor. His enthusiasm for every aspect of the field, from the theories in the latest *Sinn und Bedeutung* or *Amsterdam Colloquium* proceedings, to teaching at the introductory or advanced level, to brainstorming new approaches to difficult problems with his PhD students, is simply contagious. It was largely thanks to our frequent meetings early on that I was able to get up and running with a dissertation proposal in a timely fashion, and his detailed written comments on more mature versions of my dissertation constantly challenged me to rethink what I was doing. These are all virtues that I have striven to emulate, however poorly, in my own teaching and advising.

Malte Willer arrived in Chicago a bit later than everyone else, just as I was getting my dissertation topic together. Given that fact, it is noteworthy how quickly he became a major part of my academic life. Before he even moved here, he was already sending me written comments on my preliminary essay. The moment he moved to town, I was his teaching assistant in an introductory philosophy of language course. And once we started working together in full, I was struck by how carefully calibrated his feedback was. He doesn't feel shy about jumping right into substantial criticisms of work in progress, but perhaps somewhat paradoxically, the effect of these responses is never anything but encouraging. Whenever he formulates a worry, one can rest assured that it will be on target. I think that Malte might be our department's best kept secret when it comes to advising; he gives his students ample attention, and is equally talented at making birds eye view suggestions and working through specific kinks in a piece of argumentation. When it came time to go on the job market, he was unbelievably helpful with practical advice on how to pitch my argument differently for different purposes, what various search committees might or might not be looking for, and many similar situations that called for some subtle reading between the lines. At a somewhat late stage in the project, he made a very interesting suggestion about how to reframe the argument of the entire piece, and I think everyone agreed that the result was a significant improvement. All in all, I feel blessed to have been treated with such seriousness during my time here.

I would be remiss if I did not also acknowledge the feedback of my many incredible colleagues in the field, who generously volunteered their time to workshop drafts of its various chapters with me, including: Maria Aloni, Stina Bäckström, Rebekah Baglini, Johan van Benthem, Rajesh Bhatt, Ryan Bochnak, Emma Borg, Heather Burnett, Fabrizio Cariani, Ashwini Deo, Donka Farkas, Itamar Francez, Michael Franke, Anastasia Giannakidou, Thony Gillies, Aidan Gray, Tim Grinsell, Jeroen Groenendijk, Mark Hopwood, Dhananjay Jagannathan, Hans Kamp, Peter Klecha, Greg Kobele, Nic Koziolek, Michael Kremer,

Ben Laurence, Gabriel Richardson Lear, Sarah-Jane Leslie, Karen Lewis, Robert May, Hannah McKeown, Mathias Madsen, Alda Mari, Santiago Mejia, Joshua Mendelsohn, Patrick Munoz, Bernhard Nickel, Floris Roelofsen, Robert van Rooij, Daniel Rothschild, Jerrold Sadock, Kjell Johan Sæbø, Galit Sassoon, Katrin Schulz, Daniel Smyth, Rachel Sterken, Martin Stokhof, Anubav Vasudevan, and Hedde Zeijlstra. I was originally inspired to write a dissertation in the philosophy of language, in part, by the improbably high standards of argumentation and creativity that prevail in that literature. It is a singular privilege to be in conversation with such a talented group of interlocutors.

For their valuable comments on this project over drinks, dinners, talks, emails, mock interviews, conferences, workshops, and Q&A sessions, I would also like to thank Mahrad Almotahari, Hrafn Asgeirsson, Nicholas Asher, Andrea Beltrama, Nir Ben Moshe, Patricia Blanchette, Silver Bronzo, Octavian Busuioc, Ben Callard, Agnes Callard, Charlie Capps, Ariel Cohen, Jonathan Cohen, James Conant, Ezra Cook, Kevin Davey, Paul Dekker, David Finkelstein, Branden Fitelson, Haim Gaifman, Ephraim Glick, Yael Greenberg, Simon Gurofsky, Nat Hansen, James Hampton, Dave Holiday, Katie Howe, Adam Kimberley, Claire Kirwin, Oskari Kuusela, Brian Leiter, Tim Leffel, Kristina Liefke, Zack Loveless, Ofra Magidor, Marko Malink, Matt Mandelkern, Mary Kate McGowan, Sarah Moss, Pete Murray, Andrei Nasta, Ryan Nefdt, Dasha Polzik, Sandeep Prasada, Kate Ritchie, Ben Rodenhäuser, Mark Schroeder, Greg Salmieri, Maria Serban, Stephen Shortt, Robert Simpson, Sophia Sklaviadis, Eric Swanson, Zoltan Gendler Szabo, Tuomo Tiisala, Andy Werner, Ming Xiang, Noah Zeldin, and Zsofia Zvolensky.

My colleagues in the graduate program who offered me advice and support throughout my time here, including Nethanel Lipshitz, Dhananjay Jagannathan, Jaime Edwards, Katie Howe, Octavian Busuioc, Emily Dupree, Nic Koziolek, Nir Ben Moshe, Francey Russell, Sara Bosworth, Zack Loveless, and Toby Chow, deserve special thanks. Nir gave me tips on how to track down funding opportunities that were absolutely critical during my final year. And I am quite sure that by now, Nic has read and commented on enough pages of my writing to count as an additional committee member. At a number of key moments when I was completely stuck, it was only thanks to his advice that I managed to come unstuck.

As for general philosophical inspiration, I have had no shortage of it whatsoever during my time in graduate school, thanks to the eighty or so distinguished figures in the field who have been kind enough to make an appearance on my *Elucidations* podcast. Though the influence they have had on me collectively is far too complex to recount in detail, I simply can't resist the urge to cite Peter Adamson's 'cost/benefit' model as almost unique, among many previous attempts at characterizing what it is that philosophers do all day, in being on the right track.

Contents

Abstract

In this text, I argue that generic statements—statements of the form *F*s are *G*, such as 'Bears are furry'—are particular statements about kinds, rather than general statements about individual objects. Although statements of this form intuitively seem like generalizations, I claim that in this case, appearances are deceptive. First, I present new linguistic evidence which raises problems for the standard quantificational theory of generic sentences, according to which generic sentences contain a hidden, unpronounced quantifier. Though the simple kind theory has served as a standard alternative to quantificational approaches in the literature on generics since Carlson (1977), it also has a more sophisticated cousin, which has largely been ignored. I develop an extension of the sophisticated kind theory and show how it can neatly account for these phenomena while sidestepping the standard objections to the simple kind theory. At a broader level, I would like to claim that if a kind theory provides the best explanation for the truth conditions of these sentences in English, then it tells us something interesting about English speakers: namely, that in virtue of their speaking English, they implicitly presuppose an ontology with kinds as possible objects. In this way, I suggest, the search for the best semantic theory of generic sentences has the potential to lead us towards a new, philosophically valuable conception of kindhood.

Chapter 1

Introduction

1.1 What Are Generic Sentences?

Most of the general statements that populate everyday speech are what philosophers and linguists call *generic statements*. These are (what appear to be) loose generalizations which take the form *F*s are *G*, such as:

(1) a. Dogs have four legs.
b. Rattlesnakes are poisonous.
c. Policemen wear uniforms.
d. Cars have wheels.
e. Dutch people love bikes.

Generic statements have held the fascination of so many authors because, unlike strict generalizations, they are able to tolerate exceptions. That is to say, it is possible for there to be counterexamples to a generic statement without that statement having to be false. Consider the following cases:

(2) a. Fido lost his leg in a tragic accident.
b. Raquel the rattler has had her venom glands removed.
c. Mandy is a plainclothes detective.
d. My old car is sitting in the junkyard with no wheels.
e. Aafke, to the chagrin of her friends and family, prefers cars.

What makes sentence (1-a)-(1-e) striking is that not only are they all uncontroversially true; sentences (2-a)-(2-e) are just as uncontroversially consistent with them. In general, the second set of sentences would not count as evidence against the first set of sentences. From a theoretical point of view, that could seem odd. Is the purpose of a general statement not to cover all cases? On a fairly intuitive understanding of what generalizing is, a single coun-

terexample ought to be enough to refute a generalization. Contrast sentences with the word *every* in them:

(3) a. Every dog has four legs.
b. Fido has three legs.

(4) a. Every Dutch person loves bikes.
b. Aafke prefers cars.

Unlike counterexamples to generic statements, the counterexamples from (3-b) and (4-b) do suffice to refute the general statements in (3-a) and (4-a). It fairly intuitive why these counterexamples are inconsistent with the general claim. Assume the following standard definition of *every* statements:

(5) Every *F* is *G* just in case the set of all *F*s is a subset of the set of all *G*s.

Clearly, if the set of dogs is a subset of the set of quadrupeds, then it is impossible for any dog to have three legs. So sentences (4-a) and (4-b) are inconsistent (and likewise for (5-a)-(5-b)). But sentence (1-a) is consistent with sentence (2-a)—so at the very least, it must mean something different from sentence (3-a). The difficult question is what.

Over the past 35 years—first in linguistics, then in artificial intelligence, then in psychology, and now in philosophy—there has been a vast amount of research into what these generic statements mean. When a speaker proclaims that cats have fur, what exactly can she be saying, such that her statement can remain true even when certain exceptional cats lack fur? Obviously not that every last cat has fur, without exception. But is there a definition we could give of '*F*s are *G*' that would make it consistent with the possibility of certain *F*s being *G*?

The most common temptation at this point in the exposition is to hypothesize that sentence (1-a) means something like 'Most cats have fur.' It is a natural thought to have, and if it were true, then it would solve the problem just described in a simple way. But unfortunately, matters are not so simple. There are clear differences between 'most' statements (henceforth, *majority statements*) and generic statements. To see why, consider the following inference pattern schema, which is almost certainly deductively valid:

(6) *F*s are *G*
*F*s are *H*
$\therefore$ *F*s are both *G* and *H*

That is, no matter which properties are substituted for *F* and *G*, the result is an argument whose conclusion cannot be false when its premises are true. By contrast, the following inference pattern is deductively invalid:

(7) Most Fs are G
Most Fs are H
$\therefore$ Most Fs are both G and H

The invalidity of (7) can be demonstrated by means of the following counterexample:[1]

(8) $||F|| = \{b, c, d\}$
$||G|| = \{a, b, c\}$
$||H|| = \{c, d, e\}$

In this model, most Fs are G, most Fs are H, but it isn't true that most Fs are both G and H:

(9) $||F|| \cap ||G|| = \{b, c\}$; $||F|| \setminus ||G|| = \{d\}$ *more Fs are G than not*
$||F|| \cap ||H|| = \{c, d\}$; $||F|| \setminus ||H|| = \{b\}$ *more Fs are H than not*
$||F|| \cap (||G|| \cap ||H||) = \{c\}$; $||F|| \setminus (||G|| \cap ||H||) = \{b, d\}$ *fewer Fs are G & H than not*

If 'Dogs have four legs' truly meant the same thing as 'Most dogs have four legs,' then the problem would be resolved, because it is obviously consistent to say that some dogs have only three legs and that most dogs have four legs, as long as the non-quadrupedal dogs do not outnumber the quadrupedal dogs. However, this difference in inference patterns suggests a deep difference in meaning between the two kinds of statement; if 'Dogs have four legs' really were synonymous with 'Most dogs have four legs,' we should expect arguments of the form represented in (7) to be valid. So much for the most obvious answer.

Another initially tempting move is simply to be a skeptic about generic statements. According to the generic skeptic, generic statements are all vacuous. Making a generic statement is a kind of sneaky conversational move made only by those who wants to avoid committing themselves to anything. Another way to think of skepticism about generic statements is that according to such a view, generic statements are nothing but invitations to commit the 'No True Scotsman' fallacy:

(10) **Jean:** Scotsmen are brave.
Joan: What about MacDougal? He's a deserter.
Jean: Well, then he isn't a *true* Scotsman.
Joan: What about Wilkie? He's a fraidy cat.
Jean: Yeah, but that's just because he isn't a *true* Scotsman.
Joan: But what about...
Etc.

In principle, this debate could go on forever. No matter which counterexample Joan tries to cite, Jean insists that it is irrelevant, apparently for no other reason than that it is a coun-

[1] For any predicate F, $||F||$ is its extension. For any two sets A and B, $A \cap B$ is the intersection of A and B, and $A \setminus B$ is the relative complement of A in B.

terexample. The skeptic's thought is that anyone who says that *F*s are *G*, much like Jean in the above dialogue, really means the following:

(11) Every *F*s is *G*, except in cases where it is not.

This, of course, is a tautologous statement—which means that on the skeptic's picture, generic statements are all trivially true. Skepticism is especially tempting when it comes to offensive generic statements expressing prejudicial attitudes about disenfranchised social groups. Part of what makes xenophobic attitudes so vexing is that they seem immune to refutation through evidence, in exactly the way depicted above. So it is understandably tempting to hold that believing any generic statement at all means committing oneself to No True Scotsman reasoning patterns.

The trouble with this solution is that skepticism about generic statements is an extreme position. Even if many of the generic statements made by native speakers are hurtful and offensive, a great many of them clearly are not, and furthermore are not so easily dispensed with. A generic skeptic is committed to thinking that the vast majority of the claims we make are empty and tautologous. But generic statements play a central role in our cognitive transactions; they serve as a vessel for our thinking about many of the topics that hold great importance for us:

(12)	a.	Chicago winters are cold.	*the weather*
	b.	Polar bears are vicious	*the wild*
	c.	Engines have many parts.	*engineering*
	d.	Dutch people love bikes.	*culture*
	e.	Grooms wear tuxedos.	*cultural conventions*
	f.	Logic exercises are difficult.	*mental tasks*
	g.	Abortions are wrong.	*ethics*

Though not a priori impossible, it would nonetheless be quite something if it turned out that most of what was said on these central topics of discussion was just empty and tautologous.

Another problem with skepticism is that there are generic sentences which very clearly seem to be false. If '*F*s are *G*' meant 'Every *F* is *G*, except in cases where it is not,' then there would be no false generic sentences. But what about the following?

(13) a. Birds have five legs.
b. Birds don't fly.
c. Liberals primarily support the free market.
d. Conservatives love public health care.

Surely there is a difference between statements like (13-d) and (1-a). No one would think they were both automatically true. Anyone who went around making any of the claims in (13) would justifiably be faulted for spreading misinformation. For example, imagine a

parent who is trying to teach their two-year-old child what a bird is. Any good explanation of what a bird is should include the information that birds fly. A parent who told their child that birds did not fly would be actively shirking their parental responsibilities. This is almost certainly a minimum prediction that any account of generic sentences should be expected to make. And the biggest problem with skepticism about generic sentences is that it lacks the resources to draw a distinction between generics that are true and generics that are false.

The assumption henceforth will be that generic sentences are used to make substantive claims. The most basic challenge for anyone who wants to explain their truth conditions is to explain how they can tolerate exceptions without drifting into vacuity.

1.2 Definites, Indefinites, D-Generics

So far, the focus has been on bare plural generic sentences, and that is where the focus will continue to be throughout this work. As far as this text is concerned, *generic sentence* will simply mean 'sentence of the surface form '*F*s are *G*,' which has an episodic characterizing meaning. This is a narrower use of the term than is common in the literature, which is typically also concerned with the following kinds of cases:

(14) **Definite Generics**
   a. The grizzly bear has four legs.
   b. The car engine has many parts.

(15) **Indefinite Generics**
   a. A grizzly bear has four legs.
   b. A car engine has many parts.

(16) **D-Generics**
   a. Dodos are extinct.
   b. The dodo is extinct.
   c. Dodos/The dodo are/is widespread.

Definite generics are like bare plural generics, but with a definite noun phrase in subject position; indefinite generics are the same, but with an indefinite noun phrase in subject position. D-generics differ greatly in meaning from the sentences we've been considering, insofar as the properties they ascribe cannot hold of individuals. Sentence (16-a) would never be mistaken for any kind of general claim about dodos; according to the standard view, it ascribes a property directly to the kind *dodo*. It is clearly a very different sort of statement. One might wonder, in that case, why D-generics would be thought of as generics in the first-place. The reason is that they share the surface form '*F*s are *G*.' Therefore, any theoretical approach that takes surface variations seriously will have to explain when

sentences of the form *F*s are *G*s are interpreted in the one way and when they are interpreted in another.

These further variations are certainly of interest to philosophers and linguists. However, the differences in meaning between definite, indefinite, and bare plural generics are extremely subtle, and are as yet only partially understood.[2] It is the author's opinion that complex phenomena are best handled piecemeal, in order to isolate different effects which may interfere with one another. Rather than painting these different flavors of generic statement with the same broad brush, as sometimes happens in the literature, this text advocates getting clear about what is happening in the bare plural case before carefully proceeding to the definite and indefinite cases. The example of bare plural D-generics, on the other hand, is quite simple, and readily accounted for by the analysis set forth in this text.[3]

1.3 Generics vs. Habituals

Astute readers will have noticed that the class of statements marked off as generic in the preceding section were characterized as having an *episodic* meaning. This is in order to contrast them with *habitual* generic sentences. The last section situated the subject of this text in relation to a familiar taxonomy of generic sentences. But this work will insist on a further distinction throughout which, although it not customarily drawn in the literature, is of singular importance.[4] This will be referred to henceforth as the distinction between generic sentences and habitual sentences.[5] The following four-way contrast illustrates the distinction:

(17)	a.	Matt is tall.	*neither*
	b.	Matt smokes.	*habitual*
	c.	Bears have fur.	*generic*
	d.	Bears hibernate in the winter.	*generic + habitual*

Note that this way of using the term *generic sentence* deviates slightly from standard practice, which would classify (17-b)-(17-d) as generic sentences. The reason for collapsing these three cases is that there has always been a strong intuition, following Carlson (1977a), that sentences like (17-b) and (17-c) were engaged in something similar. Sentences like (17-c)

---

[2] The most detailed attempt at bringing indefinite and bare plural generics under a single, unified theory is Greenberg (2003, 2007).

[3] What to say about the definite cases is less obvious, but ought to follow from a good theory of definite generics in the genre of example (14).

[4] Section (38) will discuss one example of how failure to recognize this distinction has led researchers to miss important patterns in the linguistic data.

[5] Note that *habitual* is here intended to mean 'semantically habitual' or 'having a habitual meaning'—not 'receiving habitual morphology.'

loosely generalize over individual objects, and sentences like (17-b) are more or less the same, except that rather than loosely generalizing over individuals, they loosely generalize over events.[6] Just as 'Bears have fur' can be true even if some bears here and there have been shaved, 'Matt smokes' can be true even if there are many occasions on which Matt does not smoke. The one sentence raises interesting philosophical questions about how many bears need to have fur in order for it to be true, and the other raises interesting philosophical questions about how often Matt needs to smoke a cigarette in order for it to be true.

There is definitely something right about this intuition. And indeed, there will turn out for be strong affinities between habituals and my semantics for generics. Nonetheless, it is safest to begin by assuming that they are distinct phenomena. One strong consideration in favor of this starting point is that generic sentences seem more closely tied to the idea of normality. It is not unreasonable to think that sentence (17-c) means something very roughly like 'Every normal bear has fur.'[7] But it is very unlikely that (17-b) means 'Every normal occasion in Matt's life is one during which he is smoking.' Surely, the vast majority of normal occasions in any smoker's life are ones during which she is not smoking. What this shows, at the very least, is that some additional factor not present in generic sentences is present in habituals.

More importantly, after acknowledging the distinction between genericity and habituality, it becomes possible to draw an important distinction between cases (17-c) and (17-d). Because the latter is both generic and habitual, there are two different kinds of exception to it. One way of being an exception to sentence (17-d) is by being a bear that never hibernates—such as a bear in the zoo. Such a bear would constitute an exception to sentence (17-d) *qua* generic. Another way of being an exception to sentence (17-d) is by being an occasion on which no bears in the world hibernated. Such an occasion would be an exception to sentence (17-d) *qua* habitual. A correct semantic analysis of these cases needs to predict that sentences like (17-d) appeal to two different kinds of generality, whereas sentences like (17-c) only appeal to one.[8]

---

[6] Or temporal parts of individual objects, or Kratzerian situations. The are various ways to cash the intuition out formally.

[7] Though as will be observed in section 2.2, there are limits to this synonymy.

[8] Here it is worth remarking that a growing contingent of semanticists have been calling into question the habitual/episodic distinction, arguing essentially that we only make habitual statements. These arguments are based on the fact that in languages such as Hindi, even what we would think of as stative predicates receive habitual morphology in generic sentences (Deo, 2009). This is an interesting and fruitful line of investigation. If it turns out that these arguments are correct, standard assumptions about stative aspect will have to be revised more generally, and one consequence of this revision will probably be that sentences like 'Bears have fur' will receive the same sort of treatment as 'Bears hibernate.' But until then, as long as we are persuaded that some languages exhibit a stative/eventive distinction, it behooves us to recognize the difference between merely generic sentences, and sentences that are both generic and habitual.

Given any particular generic sentence, we can determine whether it is semantically habitual or semantically episodic by checking to see whether the predicate in object position is stative or eventive, using the standard grammatical diagnostics (Vendler, 1957):

(18) a. ??Bears had fur for four hours.
b. Bears hibernated for four hours.

(19) a. ??Bears are having fur.
b. Bears are hibernating.

If a *for* time adverbial can be adjoined to it with no loss of grammaticality, then it is a habitual generic. Otherwise, it is an episodic generic. And if it can be put in the present progressive without turning ungrammatical, then it is a habitual generic. Otherwise, it is episodic.

1.4 Is There a Generic Quantifier?

Generic statements look like generalizations. But are they really? The central semantic question taken up in this text is: are generic statements general statements about individuals, or particular statements about kinds? What is the statement that cats are furry *about*? Is it a statement applied pointwise to a range of individual cats, or is a statement about some other type of object, which might be called the kind *cat*? Broadly speaking, there are two ways to approach this question. The first ascribes to them a tripartite logical form consisting of a quantifier, restriction, and scope. The second construes them as monadic predications:

(20) a. **Quantificational Semantics**
*F*s are *G* iff $Q(F)(G)$,
where Q is a quantifier, and F and G are predicates
b. **Kind Semantics**
*F*s are *G* iff $G(F\text{-}kind)$,
where G is a predicate and *F-kind* is the name of a kind

The following chapters of this text argue for a kind theory of generics. The view holds, very generally, that bare plurals in subject position refer to kinds,[9] that a generic sentence predicates a property of a kind (rather than quantifying over individuals), and that a generic sentence allows us to ascribe properties of ordinary objects to kinds because it contains an operator that shifts object-level predicates to kind-level predicates. Thus far, the theory is identical to that presented in Carlson (1977a), Chap. 5—a fact which sets it apart from the more minimal kind theory of Liebesman (2011). But it builds on Carlson's theory by

[9] Whether they are names of kinds or descriptions that refer to kinds is a question left until Chapter 4.

giving a definition of this predicate modifying operation that is driven by a new philosophical account of kinds as production processes. This interpretation gives it certain affinities with normality theories of generics, which have tended to be quantificational, but it does so while offering a novel approach to what has been the principal challenge for normality theories: gendered generic sentences. The theory put forth here aims to roll the best of quantificational and kind theories up into one package.

The principal empirical motivation for the analysis about to be presented is that generic sentences do not contextually domain restrict. It may perhaps be argued that they can engage in something that superficially resembles contextual domain restriction. But on further inspection, even if we do admit these cases as data, they show clear signs of being something other than contextual domain restriction. A kind theory of the sort developed in this text does the best job of predicting the behavior of generic sentences in these cases, in a way that neatly ties into two differences between them and sentences that contain an overt quantifier.

The argument given also has an important metaphilosophical upshot, which will become relevant at various places in this work. The upshot is that the kinds of questions that come into play in natural language semantics are just as relevant for philosophy as they are for linguistics. Natural language semantics is an interesting beast, historically speaking. It was originally created by philosophers, then largely taken up and developed by linguists. During the 80s and 90s (in the US, in any case) it started to look like philosophers had lost interest in the project. Over the past fifteen years, however, there has been a resurgence of interest among American philosophers of language in both natural language semantics and in the generative paradigm in linguistics as a whole. This is a welcome development, and it is one of the hopes of this text to demonstrate that the discoveries of linguists, though they come about in a different context of inquiry, can be of distinctly philosophical interest. Natural language semantics might then be brought back into contact with its historical roots.

More specifically, this text will argue that in trying to come up with a fully compositional semantics for generics, we uncover logical and linguistic evidence to the effect that speaking English commits one, in a sense to be spelled out in the next section, to the existence of kinds.[10] In addition to being helpful for negotiating the challenges faced by semanticists, this notion of kinds is of independent philosophical interest. The basic metaphilosophical picture here is as follows: we investigate what model-theoretic structures are required to give a semantic theory of generic sentences. We then notice that these

[10] The notion of kind to which they are committed will be spelled out in Chapter 3. The exact nature of the ontological commitment will be taken up in Chapter 5.

model-theoretic structures share some important similarities with what philosophers have called kinds. There are of course also differences. But the similarities are strong enough for us to be able to argue that the English language, at least when it comes to generic sentences, has a new notion of kinds tucked away, as it were, inside its grammar. And this notion of kinds is something that semantic investigation can uncover.

It may sound as though the goal is to revive an old-fashioned method that most philosophers now reject: the idea that the best way to learn about a given philosophical topic is to investigate the grammar of our talk about it. That is not exactly the case. For one thing, it is in the spirit of this work to pursue methodological pluralism,[11] which means that any notion of the philosophy of language as any kind of 'first philosophy' is summarily rejected. The philosophy of language should hold no privileged place over other areas of the discipline. Rather, it is preferable to take the view that investigating language is one valuable way of doing philosophy. Although it is possible to draw philosophical conclusions from a semantic analysis, this project is considerably more vexed than most advocates of the linguistic turn have assumed. If it is indeed possible to draw such conclusions, they are of necessity more tentative than the conclusions that philosophers are used to drawing. But perhaps a little modesty in one's conclusions can be a good thing.

It is the position of this chapter that one can learn something about what kinds are, according to the rules for using generic sentences, by studying the rules that native speakers follow in using generic sentences. In order to make it clear why, a detour into a general discussion of this topic is called for. The general discussion should also clarify why adopting (some of) the methods of natural language semantics can yield great philosophical dividends.

1.5 Natural Language Metaphysics

The idea that studying the workings of human language might be a way of yielding philosophical results has a somewhat vexed history. And yet it is one of the first things that undergraduates hear when they take their first philosophy courses: the 20th century introduced a new method for doing philosophy onto the scene whereby instead of investigating philosophical topics directly, it became possible to investigate the way we talk about those topics. When one thinks of the linguistic turn, one typically thinks of of G.E. Moore's recommendation that ethics be re-envisioned as an investigation into the meaning of the word *good* (Moore, 1903), Gottlob Frege's claim that the most fundamental ontological distinction—

[11] More specifically, it is in the spirit of this work to pursue *conjunctive* pluralism in philosophy, in the sense explicated in Novaes (2014), according to which it is not enough for philosophy simply to pursue a variety of methodologies separately; those methodologies should interact.

that between concept and object—is one that can be extrapolated from the grammar of our mathematical statements (Frege, 1951), or Rudolf Carnap's view that whether a state of affairs is necessary boils down to whether the sentences used to describe it would be true by definition (Carnap, 1947). Today this methodology is making a bit of a return, for example, in recent work arguing that the alternative-sensitivity of deontic modals in natural language has something to tell us about the nature of reasons in metaethics (Snedegar, 2013).

At the same time, it is natural to have reservations about whether arguments of this form can ever work. Why should it be assumed that just because the logic of English commits us to saying that *F* exist, *F*s really do exist, or that just because English makes it impossible to say *P* without also saying *Q*, that *Q* necessarily follows from *P*? If saying something does not suffice to make it so, then presumably neither should presupposing something. That is probably why many contemporary philosophers find the move from 'we cannot help but assume XYZ' to 'XYZ' deeply suspect.[12] Another reason is that Kripke (1980) is typically understood to have brought to light a deep distinction between a priori truths and necessary truths. It can be necessarily the case that Samuel Clemens is Mark Twain even if coreference with the name *Mark Twain* figures nowhere in the definition of the name *Samuel Clemens*, and it can be necessarily the case that tigers are animals even if being an animal figures nowhere in the definition of the word *tiger*. So there is a divide between what can be known automatically, in virtue of how the expressions in a language are defined, and what must necessarily be.

Another approach (Varzi, 2002) is to attack the problem one step earlier, raising problems for the idea that natural languages come with *any* ontological commitments. Perhaps natural languages are fundamentally metaphysically indeterminate—perhaps they are amenable to any number of model-theoretic analyses, some of which suggest certain metaphysical views, and some of which suggest others. If that is the case, then questions regarding whether a speaker's commitments line up with the way the world actually is never arise, because there simply are no such commitments about which to ask them.

The aim of this section is to spell out a specific form of metaphysical argument from natural language, and to show that it holds up against this style of criticism. The idea will be that native speakers of a particular language incur certain metaphysical commitments, merely in virtue of their native competence in that language. In the discussion that follows, 'Sentence ϕ is committed to *XYZ*' should be understood as shorthand for saying that anyone who speaks the language containing sentence ϕ is implicitly committed to the assumption that *XYZ*. It will also be important to distinguish two components or 'steps' in metaphysical arguments from natural language. The first question is: when is the philosopher entitled

[12] See, for example, Gross (2006).

to say that ϕ is committed to the existence of things of category F? The second question is: when is the philosopher entitled to conclude that therefore, Fs do in fact exist? Although the norm in philosophy has been to pose both of these questions in succession, this section will be interested in the first exclusively. The interest will be in how it is possible to determine what speakers of a language are committed to—not in whether what speakers of a language are committed to is actually true.

Of course, even though the second question is not the focus here, there is no denying that it is a compelling question. The position taken in this section is merely that the first question deserves careful attention, regardless of one's views on the second. Some philosophers think it is possible to draw conclusions about what actually is the case from a speaker's commitments, and some disagree. Call the former group *copernicans*, following Kant's famous slogan,[13] and call the latter group *anti-copernicans*. Even if copernicanism is correct, taking the first step is a prerequisite for taking the second step, which in turn means that any problems raised for the possibility of taking the first step are thereby also problems for the possibility of taking the second step. And if anti-copernicanism is correct, then there are still compelling reasons to think of claims about a speaker's commitments as belonging to metaphysics. To see why, consider an analogy to folk metaphysics. Thinking that folk metaphysical intuitions are often misguided is still consistent with thinking that they can serve as a prima facie guide to metaphysical facts, for instance. Even if it turns out that folk metaphysical intuitions are not even a helpful guide to metaphysical facts, they still deserve the attention of metaphysicians, if only because they are compelling illusions for which metaphysicians ought to be on the lookout.

A further reason to remain agnostic about whether the commitments a semantic analysis uncovers are in fact true is that uncovering these commitments takes an extraordinary amount of hard work, performed by a sizable community of inquirers over several generations. The question whether a given commitment is true only genuinely becomes live once it has been determined that native speakers in fact have that commitment. But more importantly, once philosophers turn their attention to the second step, the terms of the discussion undergo a dramatic shift. Specifically, the customary route to the early Wittgensteinian position that linguistic expressibility and metaphysical possibility are coterminous is to mount a transcendental argument.[14] And although such arguments are not without their persuasiveness for those who find their conclusions sympathetic, the fact remains that they are

[13] See *Critique of Pure Reason*, Preface.

[14] I understand a transcendental argument to be one of the form: 'Philosophical fact ϕ is a precondition on our being able to speak or think in the first place. Therefore, ϕ is the case.' The classic example of this strategy in the latter half of the 20^{th} century is the 'descriptive metaphysics' championed by Strawson (1959).

quite difficult to assess, which perhaps explains their limited track record of success in convincing parties who disagree with their conclusions going in.

In short, then, given that there is plenty of work to do already in figuring out whether a given natural language brings with it any ontological commitments—and what some of these commitments might be—it would be worthwhile to stay focused on these questions. The next step is to provide a more detailed specification of what it might mean for a language to come packaged with such philosophical assumptions.

Practical and Theoretical Commitment

This section appropriates the term *ontological commitment* in full knowledge that its usage deviates somewhat from that of Quine (1948). Quine thought that ontological commitments were paradigmatically incurred by scientific theories, which would ideally be codified as sets of sentences in first-order logic. This chapter, by contrast, thinks of ontological commitment as the sort of thing incurred by a native speaker (and derivatively, by a class of sentences in that speaker's language).

The benefit to adopting this somewhat deviant pattern of usage is that the phrase *ontological commitment* is perfect for capturing the characteristic feature of the metaphysical commitments in question: namely, that a speaker can be committed to a claim without occurrently believing it. To take an example that is familiar from the literature on philosophical logic, it is possible for a speaker to be *theoretically committed* to some statement ϕ because it is a logical consequence of her beliefs, without that speaker ever occurrently believing ϕ. But arguably, it is also possible to have what will be called *practical commitments*. Being practically committed to some proposition means behaving *as if* that proposition is true. And as in the case of theoretical commitment, it is entirely possible to be practically committed to ϕ while occurrently believing that $\neg\phi$. Suppose a speaker to be in that situation, and suppose that it is pointed out to her that she has a practical commitment to ϕ. In that case, part of what it means to be practically committed to ϕ is that when such a situation arises, the speaker is now in a tension that can only be resolved either by:

- ceasing to engage in the relevant behavior
- retracting her belief that $\neg\phi$

For an example of how there can be a practical ontological commitment that is compatible with occurrent belief in its negation, consider the following scenario. Imagine a person who occurrently believes that human races do not exist. Perhaps this person has been persuaded by Appiah & Gutmann (1996) that race has no biological basis, and is a disjunctive folk concept that is projected onto groups of people for arbitrary political ends. However, sustained observation of their day-to-day behavior reveals that whenever someone of a given racial minority is within eyesight, they cross the street suspiciously. In such a scenario, the

person in question would be practically committed to the existence of racial groups, in spite of their occurrent beliefs to the contrary. They would be behaving as if races existed. Part of what it means to call this a practical commitment is that in such a situation, there would be a tension between this person's behavior and their beliefs which would not be resolved until they either stopped treating the relevant people with suspicion or retracted their view that races do not exist.

This chapter takes the position that semantic theories of languages like English or French reveal practical commitments of this kind. If the best available semantic theory of English cannot but state the truth conditions of English sentences by making reference to, for example, events, then it follows that speaking English is a way of behaving as though events exist. The principal consideration in favor of this position is that it unclear what it could possibly mean to say that such and such is the best available semantic theory of English, if it did not follow that this semantic theory was the most plausible description of how English speakers behave. In virtue of what would a semantic theory that made indispensable reference to events be the best available description of the linguistic data, were it not thereby the case that English speakers behaved as though there were events?

It makes a difference that practical commitments are incurred by *behavior* specifically. Automatic subpersonal cognitive processes taking place at a low level in the nervous system should not count as incurring commitments, any more than a reflex to jerk one's knee ought to count as a commitment to the claim that something in the vicinity merits kicking. Nonetheless, it is also not the case that only the results of voluntary, deliberative action count as incurring commitments. A clear counterexample is belief—most of a person's beliefs are non-occurrent, and yet they incur theoretical commitments. And, as has been commonplace in the philosophical literature for some time, even occurrent beliefs are not completely freely chosen; they arise as a result of prolonged exposure to compelling evidence. They possess a strong involuntary component, and yet they still manage to incur commitments. Behavior occupies precisely the same status in the practical realm as belief does in the theoretical realm: it is sometimes but not always the result of conscious, rational deliberation. But even when it is not, it incurs commitments.

Is Natural Language Metaphysics Possible?

One of the most penetrating criticisms of the idea that one can draw philosophical conclusions from a semantic analysis was put forth in Varzi (2002). Although Achille Varzi's argument is explicitly targeted against the idea that one can draw properly metaphysical conclusions from a linguistic analysis, it is of particular interest because it raises problems for the idea that a language can even commit its speakers to anything.

Before examining Varzi's criticism, it is worth making precise what a semantic analysis is, and what it might mean to say that a semantic theory makes unavoidable reference to entities of some category F. First, following Russell (1905), assume for the purposes of articulating this criticism that a semantic analysis of some sentence of English is a truth condition-preserving paraphrase of that sentence in first-order logic. Then, following Quine (1948), the following criterion for ontological commitment becomes available:

(21) **Quinean Ontological Commitment**
A sentence of English is ontologically committed to the existence of Fs just in case there is no truth-conditionally equivalent first-order logical paraphrase of that sentence that does not entail $\exists x(Fx)$.

The criticism begins by drawing a distinction between what are called *eliminativst* and *introductionist* arguments. An eliminativst argument aims to show that in spite of appearances, some class of sentences is not committed to the existence of a certain class of entities, by virtue of the criterion in (21). For example, Russell (1905) famously argued against the Meinongian view that negative existential sentences are ontologically committed to the existence of (nonactual) possible objects. Russell's schema for logical paraphrase makes it possible to render negative existential sentences (on their noncontradictory reading) in first-order logic as follows:

(22) a. Santa Claus does not exist.
b. $\neg\exists x(climbsChimneys(x) \wedge ridesSleigh(x) \wedge \exists y(x = y))$

For the purpose of discussing Russell, assume that $\neg$ is defined in the usual truth-functional way, and that $\exists x\phi$ is true just in case there is an object that satisfies the formula ϕ. Paraphrase (22-b) never existentially quantifies over chimney-climbing sleigh riders, then, because there is a negation in front of the main existential quantifier. Thus, it is only really saying that no object in the world satisfies the description 'chimney-climbing sleigh rider.' It would only be ontologically committed to the existence of chimney-climbing sleigh riders if it were true just in case some object *did* satisfy the condition.

Going in the other direction are *introductionist* arguments, which aim to show that even though a sentence makes no surface-level reference to some type of object, any correct paraphrase of that sentence has to entail existential quantification over objects of that type. The classic example of an introductionist strategy is Davidson (1967)'s event semantics for action sentences, which recommends paraphrasing sentences like (23-a) as in (23-b):

(23) a. Bob kissed Biff at 8:00 pm.
b. $\exists e(kissing(e) \wedge agent(e, Bob) \wedge theme(e, Biff) \wedge at(e, 8pm))$

Davidson argued that we should paraphrase descriptions of human actions in the manner above because a) doing so avoids having to make use of multigrade predicates, which tran-

scend the expressive capacity of first-order logic, and b) because rendering sentence (23-a) as a big conjunction correctly predicts it to entail, for example, that Bob kissed Biff. Davidson's argument can be called introductionist in the sense that the paraphrase existentially quantifies over a kind of object that the original English sentence did not.

Varzi suggests that if philosophers only ever made eliminativst arguments, there would be no problem. Eliminativist conclusions are easier to come by, because the only thing required to demonstrate that a sentence can be paraphrased without existentially quantifying over *F*s is a single example of such a paraphrase. They have the form of a constructive existence proof. Introductionist conclusions are general, and therefore harder won. They can generally only be made indirectly, by induction from particular cases, because their conclusion is meant to show that every possible paraphrase of a given sentence existentially quantifies over *F*s.

The worry is that once introductionist and eliminativist arguments are both allowed into the philosopher's toolkit, what counts as the best analysis becomes radically underdetermined. And once there are no clear criteria for determining the best analysis, it becomes doubtful that the process of analyzing language can genuinely lead to discoveries. In other words, the worry is that when the philosopher devises a truth-conditionally equivalent logical paraphrase of some English sentence, rather than investigating the language itself, she is simply *projecting* her antecedent philosophical views into the paraphrase. The philosopher with Russell's temperament will be inclined against nonactual possible objects, and therefore inclined to come up with logical paraphrases that only existentially quantify over actual objects. The philosopher with Davidson's temperament will be positively disposed toward events, and therefore inclined to produce logical paraphrases which existentially quantify over events. But it is never as though the language itself were providing evidence for or against either of these metaphysical positions. At best, the philosopher of language is projecting her previous views into the analysis. At worst, she adding insult to injury by claiming not to have done so.

This is a powerful criticism, which gives many philosophers in the tradition of the linguistic turn something to answer for. It works by raising a worry about whether there really are workable criteria for preferring one semantic analysis over another. And indeed, holding onto the assumption a semantic analysis is a truth condition-preserving paraphrase makes it inevitable that this problem will loom large. But perhaps that is only an indication that a different conception of semantic analysis is called for.

Natural Language Metaphysics

Varzi's worry can be answered by replacing that notion of semantic analysis with the notion operative in compositional natural language semantics. 'Compositional natural language

semantics' should be taken to refer to the theoretical approach to natural language meaning developed by Richard Montague and Barbara Partee, codified in Gamut (1991) and Heim & Kratzer (1998). The goal behind compositional semantics is to do more than accurately logically paraphrase some class of sentences in the language under consideration. Rather, the goal is to outline a fully general method of deriving those logical paraphrases from definitions of the individual words in the sentence, along with the rules for putting those words together. Sentences are represented as syntactic structures—more specifically, binary branching trees which represent how the sentence was constructed. A fully compositional theory takes this syntactic structure as input and shows how to derive the meaning of the rest of the sentence mechanically, using fully general *composition rules*, which show how to compute the denotations of each intermediate node on the tree from the denotations of its immediate subnodes. A semantic analysis, on this conception, is a general method for producing correct logical paraphrases—not any particular set of paraphrases.

Following Bach (1986), we may call the attempt to glean a native speaker's ontological commitments from an analysis of her language *natural language metaphysics*. A fully compositional semantic theory uses a formal metalanguage in which to write word definitions, intermediate phrase meanings, and final sentence paraphrases. The key suggestion from Bach (1986), which drives the entire enterprise of natural language metaphysics, is that the formal metalanguage employed by any given semantic analysis has an entire metaphysical worldview built into it, which can be read off of the models used to interpret it. A compositional semantics for a natural language L systematically translates any sentence S of L into a sentence S' of a logic $\mathcal{L}$, on the basis of lexical entries for every word in L and some composition rules that specify how to derive the meaning of a complex constituent from the meanings of the simpler constituents that comprise it. Truth conditions are encoded model-theoretically. So the truth conditions for S are represented, at the end of the procedure, as the set of models that satisfy S', according to $\mathcal{L}$'s formal satisfaction definition. The models with respect to which a logic is interpreted form a toy microcosm with information about:

- what kinds of things exist
- what relations can obtain between them
- the structural features of those relations

Since Gallin (1975), it has been customary in semantics to use a typed, higher-order lambda calculus called Ty2 to write the definitions of words and logical paraphrases.[15] Whatever is in the atomic domains are the objects the logic assumes to exist. The more kinds of objects the logic has in its atomic domains, the more kinds of things the logic is

[15] Ty2 is a variation on Richard Montague's higher-order intensional logic, whose lack of dedicated modal operators makes it a easier to manage. See Montague (1973); Gallin (1975).

assuming to be possible objects. So a logic that has atomic domains set aside for objects, possible worlds, and times expresses an ontology on which there are three different kinds of most basic thing in the universe. A logic that has an additional atomic domain set aside for degrees or events will express an ontology on which there are four different kinds of most basic thing in the universe. Any addition to the set of atomic domains will constitute an extension in the ontology presupposed by the logic. The relations in the logic will indicate what sorts of relations can obtain between the things it takes to exist. And any structural constraints on relations in that logic will reflect what the logic takes to be necessities. When those constraints are varied, what the logic assumes to be necessarily true varies with them.

The semantics for a logic indicates which distinctions between possible states of affairs a logic can discern. The more expressively powerful the logic is, the more distinctions it can make between possible scenarios. If the domain of a model can contain Fs, then Fs are things the logic can 'see'; as far as the logic is concerned, Fs can exist. If the relation R is in the range of some model's interpretation function, then R is a relation the logic can 'see'; and as far as the logic is concerned, R is a relation in which one thing can stand to one another. If $\sim$ is a special relation defined on the domain of all models which the logic requires (for example) to be an equivalence relation, then the logic cannot make the distinction between situations in which $a \sim b$ and situations in which $b \sim a$, it cannot acknowledge the possibility of situations in which $\neg a \sim a$, and so on. As far as the logic is concerned, it is *impossible* for $a \sim a$ not to be the case, and $b \sim a$ follows *necessarily* from $a \sim b$. The main idea behind natural language metaphysics is that one can investigate what speakers of a language are committed to metaphysically by trying to give a compositional semantics for that language. If accurately predicting the truth conditions and entailment patterns of all the sentences in the object language requires using a logic $\mathcal{L}$, and the models of $\mathcal{L}$ are governed by a structural constraint C, then speakers of the language under consideration are committed to thinking that certain entities in the world are governed by structural constraint C.[16]

To state the view officially:

(24) **Bachian Ontological Commitment**
A sentence of English is ontologically committed to the existence of Fs just in

[16] One important caveat is that the metaphysical commitments of any logic, though revealed by the Tarski-style models with respect to which they are interpreted, are highly abstract and mathematical. Model-theoretic structure only exists up to isomorphism. So a particular logic may contain an atomic type that is informally given the name *events*, but there will not be any more to that logic's conception of what an event is than whatever structural features it ascribes to that type. The notion of 'events' that natural language metaphysics has the potential to reveal commitment to is considerably weaker than what many philosophers have meant by the term. That said, it is substantially more than nothing.

case any fully compositional semantic analysis of that sentence must make use of a logic whose models contain an atomic domain of *F*s.

Natural language semantics requires its practitioners to bring a host of considerations to bear on any given analysis. For example:

(25) a. Are its logical paraphrases entailment preserving?
b. Is the analysis consistent with constraints on natural language syntax?
c. Does it require the use of syncategorematic composition rules?
d. Is the analysis supported by crosslinguistic evidence?
e. Is the logic it uses well-behaved,[17] or does it have bizarre features like non-idempotent conjunction, or quantifiers that prohibit alpha conversion?
f. Is it consistent with what is known about language acquisition?
g. Do each of its lexical entries fit with the distributional facts about other words of the same category?

The first point to make is that this is a formidable range of constraints, which seriously narrows down the range of correct analyses. The second point is that it would be highly implausible to claim that the above desiderata had anything to do with the semanticist's prior metaphysical biases. Even though it is a complicated question which logic (if any) is required for analyzing natural language, the discussion is not trapped in a stalemate. And that is all that is required to dissipate Varzi-style objections.

Interestingly, even though Davidson was writing before compositional natural language semantics had gotten fully up and running, he anticipated this very line of response, insofar as his argument appealed to entailment patterns. The event analysis of action sentences has the virtue of correctly predicting the following entailment:

(26) Bob kissed Biff at 8:00 pm. ⇒ Bob kissed Biff.

The event analysis correctly predicts the entailment in (26) because it analyzes action sentences as long conjunctions under the scope of an existential quantifier, and it is a fact of first-order logic that conjunctions in the scope of an existential quantifier entail their (existentially quantified over) conjuncts. It is simply implausible to argue that our intuitions about entailments like (26) have anything to do with our antecedent feelings about whether events are part of the basic inventory of the world. The position taken by this chapter is that most of the considerations brought to bear on proposals in natural language semantics have exactly this feature. It would be a stretch to try to claim that the desiderata in (25) are just reflections of the theorist's previous views in metaphysics.

[17] These are problems that beset unrestricted inquisitive logic (Roelofsen, 2011) and dynamic predicate logic (Groenendijk & Stokhof, 1991), respectively.

The following example, apparently due to Terence Parsons, nicely illustrates the methodology. Suppose we think that temporal adverbs are some sort of quantifier. What is the word *twice* in the following sentence quantifying over?

(27) Matt was punched twice.

Confining the analysis of that sentence to a logic with only objects, truth values, and times means that there are only two possible sets of truth conditions that can be assigned to it:

(28) a. For at least two times t_1 and t_2 prior to the present (possibly the same two times), Matt is being punched at t_1 and t_2.
b. For at least two times t_1 and t_2 prior to the present such that $t_1 \neq t_2$, Matt is being punched at t_1 and t_2.

But neither of those options accurately renders the meaning of sentence (27). Although both correctly predict sentence (27) to be true in, for example, the situation in which Matt was punched in the left cheek at 7:00 and in the right cheek at 8:00, option (28-a) wrongly predicts sentence (27) to be true in the situation in which Matt was punched once in the right cheek at 8:00, and option (6) wrongly predicts sentence (27) to be false in the situation in which Matt was punched in both cheeks simultaneously at 8:00. An analysis in a logic that could quantify over events, on the other hand, would get all three data points right:

(29) At least two events e_1 and e_2 prior to the current event such that $e_1 \neq e_2$ are punchings of Matt.

Perhaps, if an analysis of sentences like this requires using a logic with an atomic domain for events, sentences featuring temporal adverbs like *twice* could be a basis for arguing that speakers of English are committed to the existence of events.

One might imagine the following counterargument. The truth conditions of (27) may indeed not be capturable using a logic that can only quantify over objects, instants, and truth values. However, an analysis that quantified over location-time pairs might be able to do better:

(30) For two location-time pairs $\langle x_1, t_1 \rangle$, $\langle x_2, t_2 \rangle$
such that t_1, t_2 are prior to the present and $x_1 \neq x_2$:
Matt was punched in the x_1 at t_1 and
Matt was punched in the x_2 at t_2.

The semantics in (30) does seem to capture the three data points that the original two candidate analyses could not.

The conclusion to draw from this example is that the moment one is engaged in the project of fully compositional natural language semantics, an entire route of discussion that would not have been open to the earlier paradigm opens up. Far from leading to a stalemate,

this imagined exchange is the beginning of a productive discussion. Positing the analysis in (30) for sentence (27) requires the semanticist to look for independent evidence that supports it elsewhere. She can ask: is there any evidence that sentence radicals are interpreted not only with respect to a time parameter, but also with respect to a verb-location parameter? Is there any evidence that temporal adverbs like *twice* can bind not only time variable, but time/verb-location pair variables? If the answer to those questions is no, then perhaps English speakers are indeed committed to the existence of events. If the answer is yes, then maybe not. But either way, the analysis is not beholden to the analyst's prior metaphysical intuitions.

Varzi's worry posed a problem for the idea of a 'best' possible semantic analysis, because it threatened the idea that there can be any independent criterion for preferring one semantic analysis over another, other than the semanticist's previous metaphysical inclinations. That is indeed a serious worry. But shifting to a model of natural language semantics whereby an analysis is thought of as a more general method for producing paraphrases instantly avails philosophers of many, many criteria for preferring one theory to another that have nothing to do with anyone's previous metaphysical beliefs. The only reasonable way to look at those criteria is as constraints imposed on the theory by facts about the object language itself. Thus, by complicating the operative notion of what semantic analysis is, it is possible to recover the idea that a speaker's metaphysical commitments can be discovered by investigating a language.

Of course, semantics is such a vast enterprise that any claim about 'what logic is required' for it is usually up for debate. So the claims of natural language metaphysics, on this new conception, will have to be more modest and provisional than claims made based on the old conception. At any moment, it is possible to discover a new language, or a hitherto unobserved feature of a familiar language, or a new dialect that didn't exist when the investigation began, that throws previous observations about natural language metaphysics into question, just as new discoveries can always throw previous observations about linguistic universals into question. This may make the enterprise of natural language metaphysics less spectacular and miraculous than certain philosophical traditions might have led us to expect, but it certainly makes it no less fruitful or interesting to pursue. On the contrary, it provides philosophers with substantial, rewarding work to do.

1.6 Looking Ahead

The rest of this text breaks down as follows. With an argument for the legitimacy of investigating speakers' metaphysical assumptions by means of a semantic analysis in place, Chapter 2 then launches the investigation into whether the linguistic behavior generic statements in English bespeaks any metaphysical commitments on the part of English speakers.

It raises problems for an analysis of generic sentences on which they contain a hidden quantifier, while showing the analysis on which they are statements about kinds to be immune to these problems. The main observation is that an essential feature of quantified sentences is lacking in generic sentences: namely, that natural language quantifiers contextually domain restrict. We find contextual domain restriction not just in determiner quantifiers, but in adverbial quantifiers like *always* and *never*, and also modal auxiliary verbs like *must* and *should*. Since generic sentences do not contextually domain restrict, that provides strong reason to doubt that they contain a hidden generic quantifier. Furthermore, positing that generic sentences are particular statements about kinds offers a nice explanation for that fact: a sentence with a referring expression in subject position has no 'domain,' so it cannot contextually domain restrict.

The next step is to rehabilitate the kind theory so that it can do the work that was set out for it in Chapter 2. Although the kind theory seems like a promising alternative, it faces a number of well-known objections. However, these objections only hold against the more famous simple version of the kind theory—not against its more sophisticated cousin. Nonetheless, given the amount of attention devoted to the simple kind theory in the literature, there is a certain amount of basic foundational legwork that remains to be done before the sophisticated kind theory can be put to use. Chapter 3 develops an extension of the sophisticated theory which fills the gaps in the theory and readies it for action. According to the sophisticated kind theory, generic sentences contain a monadic operator that maps properties of objects to properties of kinds. The extension of this more sophisticated kind theory offers an independently motivated philosophical interpretation of what these properties of kinds might be, along with the principle for deriving them from properties of objects. The core intuitive metaphysical idea behind the semantics for kinds provided in this text is that a kind is a production process. A property characterizes a production process just in case all the ideal outcomes for that process are ones in which it creates some things with the relevant property.

Chapter 4 takes the semantic framework laid down in Chapter 3 and uses it to provide an analysis of the phenomena discussed in Chapter 2. Generic sentences differ from quantified sentences in the relevant respects because they are particular statements about kinds. In effect, generic sentences contain a presupposition that everything in the extension of the common noun in subject position is a member of the same kind—or, in production process talk, that there exists a single process responsible for producing everything in that extension. If there exists no such kind, the sentence is infelicitous. And generic sentences do not contextually domain restrict for the simple reason that they have no quantifier and therefore have no domain.

Chapter 5 returns to some of the methodological issues taken up in Chapter 1. There, it is argued that the considerations adduced in previous chapters provide at least some

evidence that as speakers of English, we are committed to the existence of kinds. But there is a caveat: we are only committed to the notion of kinds that is required for the logic that the analysis uses, which is highly austere and abstract. As discussed in Chapter 1, what it would mean to be a kind is exhaustively specified by the expressive capacities of the logic; and that notion of kindhood is more minimal than what philosophers have usually taken kinds to be. But though minimal, that notion of kinds still has something going for it, and extracting it from our best semantic analysis of generic sentences serves as an interesting case study in how language can point the way towards new philosophical ideas.

Chapter 2

Is There a Generic Quantifier?

This chapter presents three linguistic phenomena which speak against a quantificational semantics for generic sentences, and sketches out how a kind-theoretic approach not only correctly predicts those phenomena, but offers the most satisfying explanation for why and when they arise. More specifically, the case will be made that what will be called the sophisticated kind theory—rather than the simple kind theory advocated in e.g. Liebesman (2011)—is what can do this work. The full analysis of the three phenomena will come in Chapter 4 using the foundations laid down in Chapter 3, which presents an extension of the analysis from Carlson (1977a), chap. 5 developed for this purpose.

However, before presenting the relevant data, there is still more to do. Something needs to be said about the kind of semantic theory this work aims to provide, and how the proposal made here relates to existing theories of generic statements. In the interest of not allowing the main argument of the chapter to be swallowed up in these prefatory sections, here it an outline of what is to come:

- Summary of the methodological assumptions behind natural language semantics.
- Survey of previous semantic theories of generic statements.
- Overview of the choice between a quantificational analysis and a kind-theoretic analysis.
- Three contrasts between generic sentences and sentences with overt quantifiers.
- Why a kind-theoretic analysis is more suited to accounting for these contrasts.

2.1 Methodological Assumptions

This text will continue to work under the same general rubric of theoretical assumptions, adopted more or less wholesale from the generative paradigm in linguistics. That is, it

will assume that the human language faculty is broken up into components: a syntactic component, a semantic component, a phonetic component, a phonological component, and possibly others as well, such as a pragmatic component. The syntactic component is the aspect of a person's linguistic competence that determines which sentences are structurally well-formed. In accordance with minimalist syntax,[1] syntactic structures are understood as partially-ordered, binary-branching trees that sort parts of sentences into constituent structures which preserve grammaticality when substituted for one another. Lexical items, the atomic units out of which these trees are constructed, are rendered as sets of privative features, including (at the very least) syntactic features that determine which lexical items can combine to form syntactic constituents with one another, semantic features that determine the contribution a given lexical item makes to the meaning of the whole sentence, and phonological features that determine the contribution a given lexical item makes to the pronunciation of said sentence. Somewhat less pedantically, a lexical item can be thought of as a word, part of which is a set of constraints on how it can be pronounced, part of which is a set of constraints on how it can pair up grammatically with other words, and part of which is a definition constraining what any sentence containing it can mean.

As in section (23), the going assumption will be that the meaning of a sentence is a set of truth conditions. Intuitively, the truth conditions of a sentence are the set of worldly circumstances that would have to obtain in order for it to be true. Formally, the truth conditions of a sentence will be rendered as a set of many-sorted higher-order models—i.e. the higher-order models that satisfy the logical translation of that sentence. A semantic theory of a given natural language will provide definitions for each lexical item in a typed, higher-order lambda calculus. These definitions map each syntactic tree onto a unique formula in the lambda calculus which, once fully reduced, has a well-defined set of models that satisfy it.

The resulting framework is in the spirit of Montague (1973), which takes place in two stages: first, syntactic structures in the object language are mapped to formulas in a formal language, then that formal language is interpreted into English. There are several advantages to proceeding this way. One such advantage is that using a well-studied formal metalanguage allows us to prove things about our semantic theory, such as its expressive power or computational complexity. It also allows us to directly prove which entailments it licenses—something that is impossible using English as a metalanguage. This could in principle become relevant when trying to sort through judgments about patterns of inference involving generic sentences. But more to the point, using a formal metalanguage allows

[1] See Chomsky (1995). Minimalist syntax is not to be confused with *semantic* minimalism, a view in the philosophy of language that tries to establish limits on how much context sensitivity can be part of the literal content of an utterance. See Borg (2004) and Cappelen & Lepore (2005).

one to specify rather precisely what the model theory is taking to be unexplained primitive notions and what it is using those notions to explain. To the extent that a semantic analysis carries metaphysical commitments, it is in the model-theoretic component of the analysis that those commitments are expressed.[2]

This work adopts this general set of assumptions without any further justification, mostly because they belong to the standard idiom of natural language semantics. Contemporary linguistic theory is a complicated beast, and there are an indefinite number of places in the above narrative where one might either opt for different approaches, or choose to get off the train altogether. The hope is that those engaged in the same overall project but who prefer different syntactic or semantic frameworks (categorial grammar, relational grammar, lexical functional grammar, glue semantics, head-driven phrase structure grammar, distributional semantics and so on[3]) can adapt the argument given here to the theory of their choice without excessive difficulty. Even those who have philosophical reservations about the possibility of doing formal semantics in the first place will still get something of value out of the argument put forth in this text, as long as they think that sentences of the form '*F*s are *G*' can sometimes have this philosophically charged 'characterizing' use. Although critical philosophical inquiry into the foundations of linguistic theory is healthy for both philosophy and linguistics,[4] it is also the case that such inquiry should not hold constructively-inclined semanticists from pursuing whatever lines of investigation they wish to pursue, as long as they prove fruitful. And whatever worries one may have about whether the generative paradigm in linguistics has succeeded in living up to all of its initial ambitions, it would be absurd to deny that it has proven fruitful in all sorts of ways, some less expected than others.

How many of the above assumptions does one need to adopt in order to find the idea that language has a set of assumptions built into it about what kinds are? Essentially, one has to believe that one can study the rules that any speaker of a given language obeys in speaking it. There is some set of data about the grammar of English to which a philosophical analysis of generic statements is accountable, and with a view to following the argument through in as much detail as possible, those data will be couched in terms of the standard fully compositional theory. But that there be *some* such set of data is what is most important to the view developed in this text.

The main argument of this chapter is put forth on what will be called *compositional* grounds—compositional, because a fully compositional semantic theory must concern itself

[2] This is essentially just a restatement of the main argument in Chapter 1.

[3] See, respectively, Lambek (1958), Rosen (1984), Dalrymple (2001), Asudeh (2004), Pollard & Sag (1994), and Harris (1954).

[4] Stokhof & Van Lambalgen (2011) and Stokhof (2013) are exemplary in this regard.

with all the details about how the meaning of a sentence is built up out of its parts, and is thus beholden to a massive range of data. The driving idea behind linguistic arguments made on compositional grounds is that if we put forth a hypothesis, that hypothesis is bound to have repercussions elsewhere in the language. If someone wants to argue that sentences of class *X* all contain a hidden pronoun, one way to test my hypothesis is to see whether there are features that pronouns exhibit throughout the language. For example, maybe pronouns give rise to a certain pattern of agreement in the language under consideration. If sentences of class *X* regularly fail to exhibit that pattern of agreement, then the hypothesis that sentences of class *X* have hidden pronouns in them at least has some further explaining to do.

Such is the general form of the argument this chapter will pursue. If generic sentences were implicit general statements—if they contained a hidden quantifier, one would expect them to exhibit other behavior that we associate with quantifiers. Since they fail to exhibit this behavior, the quantificational theory of generic sentences has some further explaining to do. Furthermore, a kind theory can give a satisfying explanation not just for why generic sentences fail to behave like sentences with quantifiers, but why they behave as they do.

2.2 The Landscape of Previous Approaches

Recall from section 1.1 that the task of specifying the truth conditions of statements like *F*s are *G* is making brings with it two principal difficulties: a) explaining why generic sentences can tolerate exceptions, and b) doing so in a way that allows them to make substantive claims, rather than condemning them to triviality. The semanticists who have weighed in on this topic are interested not just in precisely stating the truth conditions of generic statements, but in explaining how those truth conditions are derived systematically from the definitions of the individual words that comprise them and their mode of combination. So in other words, the project of understanding the truth conditions and logical form of generic sentences is usually carried out in the context of compositional natural language semantics. This means that the literature on generic statements has been far more concerned with the surface grammar of the sentences it is analyzing than is customary in certain traditions in philosophical logic.

To take one familiar example, Kripke (1980) is not concerned to give a systematic account of what the meaning of everything that takes the surface form of of a proper name might be. Rather, in those lectures, Kripke assumes that his audience has an intuitive feel for which class of cases he is talking about, and then goes on to inquire into the modal and referential features of proper names *in those cases*. There is no denying that what look like proper names in natural language can have non-rigid uses. But Kripke's point is that they also have rigid uses, perhaps even paradigmatically, and that this rigid behavior in

intensional contexts is a phenomenon that philosophers need to study. Someone writing from such a point of view, when confronted with a counterexample to her semantics, can always exclude those counterexamples from the purview of her consideration.

The authors who have previous written about generic sentences, by contrast, are working under the general theoretical assumption that one begins with a surface form, and then goes on to tell a general story about when sentences with that surface form end up having which truth conditions. On this picture, a theory of what '*F*s are *G*' means when it has a generic meaning should explain when sentences with that surface form have a generic meaning and when they have some other kind of meaning—if it turns out that they can). This fact is important to keep in mind when thinking about the motivations behind some of the approaches to be discussed.[5]

Quantificational Theories

Thus far, it has been argued that that '*F*s are *G*' can mean neither 'Every *F* is *G*' nor 'Most *F*s are *G*.' The next possibility to consider is that there is some other quantifier—call it *Q*—for which '*F*s are *G*' means '*Q F*s are *G*.' Broadly speaking, any semantics for generics which opts for this route will be called a *quantificational theory*:

(1) **Quantificational Theory**

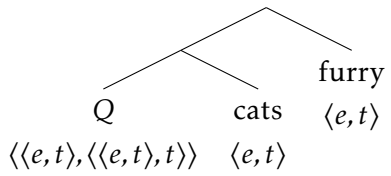

The intuition behind quantificational theories is that generic statements generalize somehow over individuals. But in what way do they generalize? Put in terms of the above analysis, this question becomes: what might *Q* be? In view of the considerations adduced in Chapter 1, it likely cannot be a simple extensional quantifier. That is, it will probably have to be able to take into account nonactual states of affairs.

An intuitively compelling option is to invoke the idea of normality. After all, statements of the form '*F*s are *G*' seem very close in meaning to statements of the form '*F*s are normally *G*.'[6] So perhaps the *Q* the analysis needs is just the quantifier 'all normal.' Although the word *normal* isn't easy to define, it is at least easy to see how 'All normal cats have fur' could be made consistent with 'This shaved cat has no fur.' Intuitively, a normal

[5] In particular, it should be borne in mind when considering the examples involving sexual dimorphism in section 2.2, which would pose no problem for a philosopher with Kripke's orientation.

[6] Though see Veltman (1996), pp. 258-259 for some important differences between them. As a quick example, consider the statements 'French people eat horsemeat' and 'French people normally eat horsemeat.' Surely the former is true but the latter false. We'll return to the horsemeat example in a moment.

cat is, among other things, one that no one has shaved. Normal cats are the ones that no one has interfered with. This approach, then, has it that generic sentences tolerate exceptions in precisely those cases where the proposed exception is abnormal. Apart from that, according to the quantificational theory, there isn't anything unusual about them. Generics behave just like strict universal quantifications, only with a restriction to normal cases.

The deepest worry about normality theories comes from Carlson (1977a), pg. 38, and has to do with gendered generic sentences—sentences which mention an entire biological kind, but ascribe to it a property that only characterizes either males or females of that kind. Consider the following pair:

(2) a. Ducks have brightly-colored feathers.
b. Ducks are oviparious.

Both of these sentences are true, even though (pre-theoretically) the first only concerns male ducks, and the second only concerns female ducks. But if '*F*s are *G*' literally means 'All normal *F*s are *G*,' then (2-a) would seem to entail that female ducks are abnormal. Already it is counterintuitive to think that male cattle are somehow the normal ones. But worse still is the fact that (2-a) and (2-b) are both true. If so, and if our simple normality semantics is correct, then the conjunction of these two sentences entails a contradiction: that female ducks are both normal and abnormal.

At first glance, these examples may seem somewhat arbitrary. Why worry about cases like (2-a) and (2-b) specifically, when generic sentences give us such a wide range of cases to be worried about? Here is why the challenge raised by these examples continues to be of such significance in the debate. Once it has been established that generic sentences cannot be majority statements, there is a strong pull in favor of thinking of them as characterizing what is normal for members of a kind. 'Bears are furry' sounds as though it is saying something about what it is to be a normal bear. Isn't that what generic sentences are all about? But there is a problem: sentences like (2-a) and (2-b) sound completely true, and yet they each ascribe properties to ducks that are distinctly *not* normal for ducks as such to have. At best, for each of those two properties, it is only only normal for about half of the duck population to have it.

The following equally vexing example brings out the general shape of the problem raised by the gendered cases:

(3) French people eat horsemeat.

Once again, one can arguably take this sentence to be true without thinking that it is at all normal for French people to eat horsemeat—if anything, a French person who eats horse-

meat is downright *abnormal*.[7] Therein lies the difficulty: on the one hand, generic statements typically seem to be saying something about what the normal members of a category are like. On the other hand, in a substantial minority of these other cases, they seem to be saying something about what the *aberrant* members of a category are like. And in the examples involving sexual dimorphism, they seem to be saying something about what members of the category that are neither normal nor abnormal are like. These considerations push in one of two directions, depending on what one decides to see as being held fixed and what one decides to see as varying. Examples (2-a) and (2-b) can be understood as both being about what is normal for a group, while varying as to whether they describe the normal features of the male group or the female group. Or they can be understood as both being about ducks as such, where one case is talking about male normality and one the other is talking about female normality. Put a bit more abstractly, one can picture the attempt to get at the normal features of some group of individuals as remaining constant across the two cases, and then think that what varies is which more specific subgroup of individuals each sentence is talking about. Or one can picture the explicitly mentioned group of individuals as constant across the two cases, and then think that what varies is the specific flavor of normality being ascribed to that group.

We now turn to a number of more sophisticated normality theories, which have tried to respond to the challenge posed by gendered generic sentences in different ways. In general, they opt for one of the above two strategies.

Conditional Normality Theories

The approach taken in Asher & Morreau (1995) and Asher & Pelletier (1997) is to construe normality as a notion that applies not to individuals, but to possible worlds. Here, normality is not a simple property that an individual animal might lack or possess, but rather a feature of a circumstance that might obtain. Technically, Asher, Morreau, and Pelletier express the notion of normality via a conditional operator. The conditional operator is defined in terms of a modest selection function semantics in the style of Stalnaker (1968), which is characterized by the following constraints:

(4) **Generic Conditional Selection Function** $f : W \times \mathcal{P}(W) \to \mathcal{P}(W)$

a. $f(w,p) \subseteq p$

b. $f(w,p \cup q) \subseteq f(w,p) \cup f(w,q)$

(5) **Generic Conditional Operator**

$M,w,g \vDash A \rhd B$ iff $f(w,\|A\|_{M,g}) \subseteq \|B\|_{M,g}$

[7] Anecdotally, the author has almost never encountered a French person who has even seen horse on the menu at any restaurant, to say nothing of having actually tried it.

The truth conditions for *F*s are *G* are then given as follows:

(6) **AMP[8] Semantics for Generics**
'*F*s are *G*' is true under any M, w, g that satisfy $\forall x(F(x) \rhd G(x))$

An example should spell out how this is different from the simple normality analysis. The selection function maps any world-proposition pair to the set of worlds that are normal with respect to that proposition. Under the simple normality analysis, 'cats have fur' is true just in case every normal cat is furry. Under the AMP analysis, 'cats have fur' is true just in case for every actual cat, all worlds which are normal with respect to that cat (according to what counts as normal in the actual world) are worlds in which it is furry. Assigning a different normality metric to different propositions is an ingenious subtlety, and allows this theory to avoid postulating a single ranking of all possible worlds from abnormal to normal.

Asher & Pelletier (1997) propose to understand what is happening in (2) as a kind of accommodation. When a speaker hears two generic statements that conflict, she accommodates each of them so that they're talking about a restricted subcategory. So, the examples in (2) would be accommodated into the following:

(7) a. $\forall x((duck(x) \wedge male(x)) \rhd coloredFeathers(x))$
b. $\forall x((duck(x) \wedge female(x)) \rhd oviparous(x))$

This is an interesting idea, though Asher and Pelletier leave their account of gendered generic sentences in a fairly sketchy state, providing no way to integrate it into their formal semantics. There are also further problems with this workaround, as observed in Leslie (2008). For instance, if the truth conditions of (2) are as in (7), then why are these sentences infelicitous?

(8) a. #Ducks have brightly-colored feathers because they are male.
b. #Ducks are oviparous because they are female.

Surely it is true to say male ducks are male, and that female ducks are female. So it seems that Asher and Pelletier's solution needs more formal detail before its ability to handle these counterexamples can be properly assessed.

Multidimensional Normality Theories

Nickel (2010b, forthcoming) seeks to address this challenge by understanding generic sentences not as universal quantifications over normal objects, or objects in normal worlds, but as second-order existential quantifications over *ways of being normal*:

[8] Short for Asher, Morreau, and Pelletier.

(9) **Multidimensional Normality Theory**
'Fs are G' is true under any M, g that satisfy $\exists N(\forall x((F(x) \wedge N(x)) \rightarrow G(x)))$

In other words, cats are furry just in case for some way of being normal, every cat which is normal in that way is furry. Nickel's theory offers an elegant solution to the problem of gendered generic sentences. Ducks do have brightly-colored feathers, because for some way of being normal (namely, the male way), all ducks that are normal in that way have brightly-colored feathers. And they are also oviparous, because for another way of being normal (namely, the female way), all ducks that are normal in that way have egg-laying capacities.

The multidimensional normality semantics has one potentially controversial consequence, which is that generic sentences cannot be argued to support default inductive reasoning patterns.[9] This seems quite counterintuitive; one of the most common uses for generic statements is to support default inductive arguments. Consider the following statement:

(10) Don't go near that rattlesnake! Rattlesnakes are poisonous!

This is surely a paradigmatic use of the generic construction, and it seems to be underpinned by the following line of inductive reasoning:

(11) Rattlesnakes are poisonous.
This animal is a rattlesnake.
$\therefore$ It is reasonable to presume that this animal is poisonous.

According to the multidimensional theory of normality, this inference is not even inductively valid. Why? Because saying that rattlesnakes are poisonous only entitles one to think that for some way of being normal, every rattlesnake that is normal in that way is poisonous. But it provides no reason to think that this animal is a rattlesnake that's normal in the relevant way. For all the semantics in (9) indicates, the premises in (11) might very well be true in a situation where only female rattlesnakes were poisonous, and in which the animal in front of the speaker was a male rattlesnake. If the premises indicate only that Ralph is a rattlesnake and that for some way of being a normal rattlesnake, every rattlesnake that is normal in *that* way is poisonous, it is simply unreasonable to conclude that Ralph is poisonous. What is to prevent Ralph from being normal in some entirely different way that would not lead to his being poisonous?

[9] It should be noted that Nickel sees this controversial consequence as a plus, because he sees it as a mistake to think that generic sentences support default inductive reasoning patterns.

The multidimensional theory of normality thus does an excellent job of handling the gendered examples, but at the cost of making it seem strange that they could play the role they very much seem to play in our commonsense reasoning.

Probabilistic Theories

Cohen (1999b) proposes a semantics for generics with two main components: probability and alternative-sensitivity. We begin by presenting the probabilistic component. Part of what a speaker is saying when she says that cats have fur might very roughly be glossed as: given any particular cat, the probability of its having fur is greater than the probability of its not having fur.

Of course, the main lesson from the gendered examples is that that cannot quite be what the speaker is saying. To deal with the gendered cases, Cohen also hypothesizes that generic statements are interpreted not in isolation, but relative to a contextually-given set of alternatives.[10] Conversational context associates with each predicate a set of other predicates, which are marked as alternative to it. For example, the set of alternatives to having fur might be having some other kind of skin. Suppose there is a function *Alt* which maps any predicate in its context to a set of alternatives, so that *Alt*(having fur) = {having feathers, having scales, having glandular skin, ...}. Then another thing a speaker is saying when she says that cats are furry is that for every alternative A to being furry, more cats are furry than A.

Combining those two intuitive ideas yields the following semantics:

(12) 'Fs are G' is true under every M, g that satisfy:
$\forall x, \forall G' \in Alt(G)$: $Pr(G(x)|F(x)) > Pr(G'(x)|F(x))$

In other words, cats are furry just in case for every cat, and every other skin-type, the probability of its being furry exceeds the probability of its having that skin type.

This semantics offers a highly inventive route to the right result in the gendered cases. The explanation it offers is that ducks have brightly-colored feathers because for any duck, it has a higher probability of having brightly-colored feathers than having drably-colored feathers, or brightly-colored fur, or drably-colored fur, or brightly-colored scales. Why are ducks oviparous? Because for any duck, it has a higher probability of being oviparous than giving birth to live young or reproducing by mitosis. And so on.

Cohen's probabilistic semantics for generic statements is probably the most flexible of the quantificational theories considered so far; it makes the correct predictions about a

[10] The idea of alternative-sensitivity is drawn originally from Rooth (1992), and is now being applied to other areas of formal semantics, such as deontic modals (Cariani, 2013).

wide variety of cases that have confounded earlier theories. But this approach still leaves some crucial philosophical questions unanswered. What mechanisms are responsible for mapping a predicate to its set of alternatives? What constraints do they obey? Are generic sentences alternative-sensitive in the same way as other well-known cases, such as utterances with topic/focus intonation? The probabilistic explanation of the gendered examples is intuitively appealing, but it would be helpful to have a more detailed theory of how predicates are mapped to sets of alternatives in context. That way, the explanation for cases like (2) would look more principled.

Kind Theories

The main competitor to the quantificational approach is the classical kind theory, first put forth in Carlson (1977a). Recall how the theories in the previous section came up for discussion. We began by considering why generic sentences tolerate exceptions. The majority analysis and skepticism about generic sentences both had insuperable problems. The best way to preserve the intuition that generic sentences generalize over individuals was to view them as quantifying over normal individuals, or over individuals in normal circumstances. However, pushing further on that idea led to trouble reconciling the intuition that generic sentences quantify over normal individuals with the fact that certain generics seem true even though it is not normal, *per se*, for the objects in question to have the relevant property.

The approach taken in Carlson (1977a) is to leave it to one side. Some generic sentences require 100% of the objects in question to have the relevant property, and some far fewer—even fewer than half. Some generic sentences can even be true when *none* of the objects in question have the relevant property: consider the sense in which 'cats have fur' would be true even in a situation where a maniacal dictator had captured all the cats in the world and ordered them to be shaved. Rather than trying to figure out what quantifier could fluctuate the proportions it checks for so dramatically from context to context, the classical kind theory draws the opposite moral: generic sentences involve no quantification. It has little to say about why generic sentences tolerate exceptions or what defeasible inference patterns they sanction. Instead, it draws support from a variety of semantic and syntactic distributional evidence.

The driving idea behind the classical kind theory of generic sentences is that bare plural noun phrases, in general, are names of kinds. So the word *cats* in 'cats have fur' denotes cat-kind. What is cat-kind? That is a substantial philosophical question, but in the

the context of Carlson's original theory, cat-kind is understood to be the mereological sum of all cats, and each cat is the mereological sum of its temporal parts.[11]

From here, the kind theory can go in one of at least two directions, both of which were proposed in different chapters of Carlson (1977a). The first, which will be referred to as the *simple* kind theory, is the version from Chapter 4 of that text. The second, which will be referred to as the *sophisticated* kind theory, is the version from Chapter 5. Here is the distinction:

(13) **Simple Kind Theory**

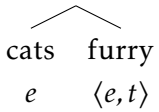

(14) **Sophisticated Kind Theory**

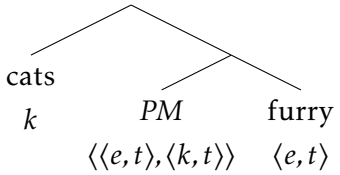

The simple theory in (13) works well enough for Carlson's initial purposes, but he eventually gives it up in favor of the sophisticated analysis in (14), which introduces a new atomic type for kinds and posits a predicate modifier *PM* which shifts properties that hold of ordinary objects to properties that hold of kinds.[12] According to the simple kind theory, the truth conditions of 'cats are furry' are, roughly, that cat-kind be furry. According to the sophisticated kind theory, the truth conditions of 'cats are furry' are that cat-kind be furry', where *furry'* is the kind-level predicate to which *PM* maps the predicate *furry*, whatever that may be.[13] The simple kind theory will be examined in a moment in section (28). For simplicity of exposition, Carlson's own reasons for adopting (14) over (13) will be discussed only after a positive proposal is put forth.

One initial motivation for the classical kind theory is that bare plurals are often interchangeable with 'kind of' phrases:

(15) a. Dolphins are my favorite kind of cetacean.

[11] Carlson's motivation for wanting to construe kinds in this way has to do with the fact that they are organized taxonomically into subkinds and superkinds, and that it is just as easy to make generic statements about these intermediate kinds as it is to make generic statements about individuals. This topic will be reprised in section 5.2.

[12] Not to be confused with the composition rule from Heim & Kratzer (1998) called 'predicate modification.'

[13] Readers who are puzzled as to what a predicate like *furry'* might mean will have all their questions answered in the following chapter. Although Carlson is deliberately silent about what such a predicate would intuitively mean, Chapter 3 of this text is devoted to providing a philosophical interpretation of such kind-level predicates.

b. What kind of animal did you encounter on your hike last week?
—Snakes!
c. Cockroaches/that kind of pest is/are quite common in the city.

Further support comes from the fact that bare plurals pattern grammatically with names rather than quantified noun phrases. For instance, when a bare plural takes a noun as its antecedent, you can substitute the bare plural for the noun without changing the meaning of the sentence:

(16) a. Birds are loud. They squawk a lot.
⇔ Birds are loud. Birds squawk a lot.
b. Bert is a quiet guy. He never talks.
⇔ Bert is a quiet guy. Bert never talks.

Not so for quantifiers:

(17) a. A bird is squawking outside. Let's go feed it.
⇎ A bird is squawking outside. Let's go feed a bird.
b. There are two people at the door. Could you say hello to them for me?
⇎ There are two people at the door. Could you say hello to two people for me?

Carlson also observes a second rather striking analogy between bare plurals and names, which is that bare plurals do not give rise to scope ambiguity when they are in the scope a quantifier. For example:

(18) Somebody loves ponies.
$\exists > Gen$
$\# \; Gen > \exists$

This sentence only has a reading on which the existential quantifier takes wide scope. That is, it has a reading on which for one single person, there are different ponies that she loves; but no reading on which for different ponies, there is a different person who loves each of them. Likewise for names in the scope of a quantifier, which only have one scope interpretation.

(19) Somebody loves Fatima.

However, one disanalogy which Carlson fails to note is that bare plurals do lead to scope ambiguity when the bare plural is in subject position and the quantifier is in object position:

(20) Swans have a favorite resting spot.

This sentence has one interpretation on which different swans have the same resting spot, and another interpretation on which different swans have different resting spots. Sentence

(19), on the other hand, gives rise to no such ambiguity. This example is standardly considered to represent a major challenge to Carlson's theory, whose predictions regarding it are somewhat hazy.[14] But like most objections to the classical kind theory, this one fails to distinguish between simple kind theories, for which cases like this do indeed pose a problem, and sophisticated kind theories, for which examples like this do not. It is also worth bearing in mind that even though this example shows there to be a disanalogy between bare plurals and proper names, Carlson's main point—that there is a disanalogy between bare plurals and proper names with respect to scope behavior—stands. The real lesson from this example is that some explanation for why bare plurals *only* give rise to scope ambiguity when the quantifier is in object position is called for. Section 3.7 will show that, even though it is unclear how a simple kind theory could accommodate either of these data points, a modest extension of the sophisticated kind theory can do so straightforwardly using standard theories of scope ambiguity.

Finally, Carlson observes that bare plurals pattern with proper names in their ability to appear in 'so-called' constructions:

(21) a. Alexander the Great was so-called because he was a powerful conqueror.
b. Scarface was so-called because of a knife wound.

(22) a. Rice Crispies are so-called because they crackle in your mouth.
b. Ladyfingers are so-called because of their shape.

These constructions sound odd with quantified noun phrases:

(23) a. #All rice crispies are so-called because they crackle in your mouth.
b. #Most ladyfingers are so-called because of their shape.

So Carlson's ultimate view, then, is that since bare plurals pattern distributionally with names, why not just say that they are names of kinds? If grammar seems to treat these things as names, then maybe that's what they are. And if they really do function as names, then there must be something in the world to which they refer. Call those things—whatever they may be—kinds. If that is correct, the argument goes, then there is good reason to think that speakers of English are committed to the existence of kinds.[15]

The great advantage to this theory is that it offers a unified account of two sentence classes that are identical in surface form but quite different in truth conditions: sentences of the form '*F*s are *G*' that have a *generic* meaning and sentences of the form '*F*s are *G*' that have an existential meaning:

(24) a. Dogs are great for hunting. *generic*

[14] See, for instance, Cohen (2001), pg. 193 and Cohen (2013), pg. 3.
[15] Understood in Carlson's work as as mereological sums of individuals.

b. Dogs are yapping outside. *existential*

The former is a generic statement, making some sort of claim about dogs as such. The latter, on the other hand, only asserts that some dogs are yapping outside. The most obvious hypothesis about this phenomenon would be to stipulate that statements of the form '*F*s are *G*' are systematically ambiguous, giving rise to both generic and existential interpretations. However, as Carlson notes, sentences of this form typically *do not* exhibit such ambiguity: each particular case usually allows either one reading or the other. Also, bare plural noun phrases pass the above three grammatical tests even in existential sentences of the form '*F*s are *G*.'

The conclusion Carlson draws from this is that there is no generic/existential ambiguity. Bare plural noun phrases denote kinds in *any* statement of the form '*F*s are *G*,' regardless of whether it has a generic or existential meaning. Instead, it is the *predicate* that determines whether the sentence receives a generic or existential interpretation. If the predicate is individual-level, the resulting sentence is a generic; if the predicate is stage-level, the resulting sentence has an existential meaning.

What makes this explanation so compelling is that the stage-level/individual level contrast is independently motivated. These two kinds of predicate can be distinguished by way of the following tests:

(25) a. There were several pencils available.
b. #There were several pencils broken.

(26) a. Yesterday, I saw Winston drunk.
b. #Yesterday, I saw Winston talented at chess.

(27) a. Matt was American. (Kratzer, 1995)
$\rightsquigarrow$ Matt is no longer alive.
b. Matt had the flu.
$\not\rightsquigarrow$ Matt is no longer alive.

(28) a. Being in Boston, Cindy can stop by MIT whenever she wants. (Stump, 1985)
b. Being unusually gifted, Cindy can solve a Rubix cube in 30 seconds.

Intuitively, stage-level predicates like *available* or *in Boston* apply to something only provisionally, and individual-level predicates like *broken* or *unusually gifted* describe properties with a certain level of permanence. However, that is only a rule of thumb; the real distinction is in terms of the above two grammatical tests. Unlike individual-level predicates, stage-level predicates can appear in there-insertion constructions. Unlike individual-level predicates, stage-level predicates sound natural as arguments of perception verbs. Unlike individual predicates, stage-level predicates implicate that a person is dead when ascribed

to her out of the blue. And unlike individual-level predicates, stage-level predicates allow for an additional interpretation in constructions like that in (28). Imagine a context in which Cindy is trying to decide whether to move to Boston. In that case, sentence (28-a) has an interpretation on which it means 'If Cindy is in Boston, then she can stop by MIT whenever she wants.' (28-b), on the other hand, yields no such conditional interpretation; it cannot mean 'If Cindy is unusually gifted, she can solve a Rubix cube 30 seconds.'

Thus, the classical kind theory offers compelling evidence that the bare plurals in generic sentences, much like other bare plurals, behave like names. Rather than trying to explain our judgments about gendered generic sentences, show what commonsense reasoning patterns they make possible, or illuminate why they tolerate exceptions, the classical kind theory relegates these problems to another part of the theory.[16] But in actuality it is mostly just happenstance that has led kind theorists to reject these problems. In principle, there is nothing preventing a kind semantics for generics from addressing the very same questions that quantificational theories have taken it upon themselves to address.[17] And indeed, that is precisely what this text will propose: extend the classical kind theory in such a way that it can do more (or perhaps all) of the things that quantificational theories can do, along with certain new things that neither can do.

None of that is to deny that kind theorists have put forth interesting arguments for the view that a semantic theory of generic sentence should avoid getting into the territory that interests quantificational theorists. Before turning to a positive proposal, it is worth considering some of these arguments. First, we consider a kind theory that seeks to address the topics of interest to quantificational theorists non-semantically; next, we consider a quantificational theory that also seeks to deal with them non-semantically.

Deflationary Theories

One very general approach that has caught on in recent years is to argue that generic sentences have proven so challenging for semanticists because they have been trying to use semantic theory to explain a set of data whose messiness just isn't semantic in nature. Accordingly, one school of thought on generics recommends assigning them a simple semantics to cover most standard cases, then handling the cases where it seems to make incorrect predictions in some other part of the theory. I'll call any theory of generics with these features *deflationary*.

[16] In Carlson's case, that other part of the theory is epistemology.
[17] **?** advances a similar challenge to the status quo.

The Deflationary Kind Theory

David Liebesman kind semantics for generic sentences has been steadily gaining attention among philosophers. It is very close to the semantics proposed in Carlson (1977a), chap. 4,[18] and, like that approach, also shares certain affinities with the view of definite generics expressed in Thompson (2009). Liebesman argues that generic statements have the following truth conditions:

(29) *F*s are *G* just in case *G*(*F-kind*)

Like the theory to be developed later in this chapter, the deflationary kind theory holds that generic statements are particular statements about kinds, rather than general statements about individuals. Unlike the theory to be developed later in this chapter, it holds that (for example) cat-kind can be furry in exactly the same sense as Felix the cat. What does it mean for cat-kind to be furry? Cat-kind is furry just in case cats are furry—end of story. Further explanation is neither required nor possible.

Liebesman's argument that no further explanation is required is quite interesting. First, he assumes that members stand to their respective kinds as parts to a whole. If that is the case, then the difficulty in saying how many furry cats it takes for cats to be furry is just a special case of a more general philosophical problem: how do wholes inherit properties from their parts? Consider the following three things that might be said about a table:

(30) a. The table is wooden. *most*
b. The table is blue. *the visually salient parts*
c. The table is touching the wall. *some*

In (30-a), nearly the entire table (except, perhaps, for some screws, washers, dowels, and adhesive material) needs to be wooden in order for the entire table to be wooden. In (30-b), only the top of the table's surface needs to be blue in order for the table to count as blue. (30-c) can be true even if only a tiny corner of the table is in contact with the wall. Coming up with a theory of how complex wholes inherit properties from their parts is challenging because it seems that in different cases, different proportions of parts are required to have the relevant property in order for the whole to inherit it. Here, Liebesman observes that the same thing is true of generic sentences:

(31) a. Turtles are reptiles. *100%*
b. Turtles lay eggs. *50%*
c. Turtles grow old. *very few*

[18] Interestingly, in the following chapter, Carlson revises his earlier proposal, shown in (13), and decides it needs to be replaced with the more complicated analysis represented in (14). Carlson's revised semantics, presented above, is the starting point for the semantics defended later on in this chapter.

Just as in the case of the table, different proportions of turtles—thought of as parts which comprise the whole that is turtle-kind—need to have the relevant property in order for turtle-kind to inherit it. And just as the difficulty of specifying the relevant inheritance principle should not shake our confidence that sentences (30-a)-(30-c) are monadic predications, we should not let the fact that it is unclear how many *F*-ish turtles it takes for turtle-kind to be *F*-ish shake our confidence that sentences (31-a)-(31-c) are monadic predications.

The deflationary kind theory provides a straightforward explanation for why conjunctions of D-generics and the generics under consideration in this text are grammatical, which is a challenging data point for quantificational theories:

(32) Diamonds are rare and shiny.

Assuming this is a case of noun phrase ellipsis, then sentence (32) in fact looks more like the following, with strikethrough indicating that the second noun phrase is 'there' in the syntax, but not pronounced:

(33) Diamonds are rare and ~~diamonds are~~ shiny.

It is generally agreed that in order for ellipsis to take place, the unpronounced and pronounced syntactic constituents must be identical, apart from their phonological features. So in order for the second *diamonds* in (33) to become silent, it must be identical to the earlier occurrence of the word. This is exactly what the deflationary kind theory would predict. But that cannot be the case according to the quantificational analysis, because the predicate **rare** is inapplicable to individual diamonds—so 'Diamonds are rare' does not contain a quantifier. And if it contains no quantifier, then the logical form of (33) must be something like:

(34) Diamonds are rare and ~~every normal diamond is~~ shiny.

Based on what we know about ellipsis (Sag, 1976), this possibility should be blocked, because 'Diamonds' and 'every normal diamond' differ in far more than their pronunciation (among many other differences, the latter is a phrase composed of three words).

However, as was noted above, this analysis has a difficult time explaining the scope ambiguity that arises in generic sentences with quantifiers in object position. Consider these cases:

(35) a. Swans have a favorite resting spot.
b. The table has a buyer.

Generic sentences like (35-a) systematically give rise to the scope ambiguity we discussed above. But sentence (35-b), like other sentences describing complex wholes, only has an

interpretation in which the indefinite takes wide scope. So the deflationary kind theory still owes us an explanation for this difference.

This objection only has traction with a simple kind theory—the further logical structure provided by the sophisticated kind theory allows for the possibility of scope ambiguity. Section 3.7 will explain how the sophisticated kind theory can neatly accommodating both of these data points.

Deflationary Psychological Theories

Leslie (2007, 2008) advocates an adverbial quantificational theory of generic sentences, mostly on the grounds that that is the consensus view in the literature.[19] According to this view, generic sentences contain an unpronounced adverbial quantifier. That much is often taken for granted.[20] But then comes an interesting and highly original twist: rather than giving this adverbial quantifier a formal definition in the semantics, Leslie takes it essentially as a pointer to a more basic, pre-linguistic mode of categorization. In other words, her semantics does look like this:

(36) *F*s are *G* just in case *Gen*(*F*)(*G*)

But whereas in a typical semantic theory, *Gen* would be a logical operator that came packaged with some further model-theoretic definition, Leslie gives it the following simple definition, which effectively outsources the task of defining it to another component of the theory:

(37) A person judges *Gen*(*F*)(*G*) to be true just in case given any *F*, her default mode of generalization would ascribe the property *G* to it.

Do not be misled by the term *default*: no allusion to default logic is intended here. Rather, Leslie's core proposal is that we all have a fundamental capacity for making inductive generalizations based on our experiences of individual things. This basic capacity is something we develop before we acquire language, around the same time we develop our prelinguistic psychological capacity for categorization. It is our 'default' method of generalizing in the sense that we engage in other varieties of generalizing—for example, the kind of generalizing that is triggered by quantificational determiners like *every* or *some*—only when our language faculty gives us explicit instruction to do so. We become competent in other modes of generalizing—which lend themselves to mathematical definition but are more cognitively demanding—only after mastering a nontrivial amount of our native language.

[19] That claim is potentially debatable. If it is true, then Chierchia (1998); Krifka (2004); Cohen (2004) are certainly significant voices of opposition.

[20] The suggestion was first made in Heim (1982), §4.3 for indefinite generics, and in Farkas & Sugioka (1983) for bare plural (as well as indefinite) generics.

It may be somewhat surprising to see the theory defended in Leslie (2007, 2008) discussed under the same heading as the theory in Liebesman (2011), given that the two theories differ in nearly all of their details, and some of the most serious criticisms of Liebesman's view have been advanced by Leslie herself in Leslie (2013). Nonetheless, Leslie's theory and Liebesman's theory have in common the fact that they favor a minimal semantics for generic sentences, and the fact that they are both motivated by the need to explain why almost no natural language has an overt generic operator.[21]

Here, however, the similarities end. Leibesman favors an austere semantics for generic sentences and argues that the further complexities in the data have to do with the metaphysical question of how wholes inherit features of their parts. Thus, the challenge of dealing with the canonical set of examples in the literature is outsourced to metaphysics. However, it is also part of Liebesman's view that there can be no systematic metaphysical theory of how wholes inherit features from their parts. This is where he parts company most starkly with Leslie. Liebesman's theory outsources the standard difficulties to the metaphysics, in order to mark them off as theoretically intractable. Leslie's theory outsources the standard difficulties to a cognitive science, because she thinks they are better suited to be dealt with in that part of the theory. Thus, it would be fair to say that Leslie's theory is only *semantically* deflationary; not *psychologically* deflationary. In fact, she has a substantial amount to say about how this cognitively default mode of generalization works, and even tells a fairly precise story about when speakers will and will not assent to generic statements—a story that looks a lot like truth conditions. However, she insists that telling such a story means describing a set of truth *specifications* for generic sentences, as opposed to truth *conditions*. Leslie ascribes the following truth specifications to generic sentences:

(38) '*F*s are *G*' is true just in case:
  a. some *F*s are *G*, if:
    (i) *G* is the kind of property in virtue of which we identify *F*
    (ii) *F*s are not an artifactual kind
  b. to be *G* is the function of *F*, if *F*s are an artifactual kind
  c. some *F*s are *G* and others are disposed to be *G*, if *G* is striking
  d. most *F*s are *G*, otherwise

This is a suggestive way to phrase things, but without any further explanation as to what exactly the truth condition/truth specification distinction consists in, it is difficult to know what to make of it. It is not entirely clear whether Leslie thinks these truth specifications

[21] Though see Quer (2011), where it is argued that American Sign Language features an overt generic operator. There is also some evidence for a dedicated generic operator in Estonian and Turkish, where it takes the form of an intersubjectivity modal. It would seem that the crosslinguistic evidence on this is still coming in.

are *impossible* to express as truth conditions, whether they could, but doing so would be incorrect, or whether she simply means to leave it open whether they can. I think it would be difficult to make the first case; there is nothing about the expressive idiom of model-theoretic definitions that prevents her truth specifications from being formulated in it. In the second case, it would be helpful to hear more about what is at stake theoretically in explaining some of a speaker's truth judgments by appeal to her extralinguistic faculties, and why the messiness inherent in the data on generic is more suited to being modelled using that kind of theory. In the third case, semanticists have their work cut out for them, the doors are open for anyone who would like to incorporate Leslie's considerable insights into their latest analysis.

More importantly, even though Leslie advocates a quantificational semantics for generic sentences, the psychological experiments she has been conducting in collaboration with Sandeep Prasada, Susan Gelman, Sam Glucksberg, and others could be argued to offer strong support for the analysis on offer in this text. Leslie and her colleagues have uncovered evidence to the effect that children become competent at using generic sentences before they become competent at using first-order quantifiers. And they become competent in both *well* before they become competent at using proportional quantifiers. Leslie draws the following moral from these experiments: human beings have a fundamental mode of generalizing, which comes more naturally to us than the kind of generalizing involved in using quantifiers. This moral is generally quite plausible, but a kind semantics for generics can also just be viewed as a way of representing that fundamental mode of generalizing. We may be able to do it while prelinguistic, but even if so, why not think that this kind of reasoning is encoded linguistically once we start engaging it in by means of language? Again, whether Leslie would disagree with that suggestion is not entirely clear. As was mentioned above, her decision to use the quantificational analysis is not an enthusiastic endorsement. Rather, it is chosen merely because it has the virtue of being the least controversial analysis. But for Leslie's purposes, it seems that she could just as well have gone with a Carlsonian analysis.

A Positive Proposal

The existing accounts of generic sentences seem to assume a trade-off between giving a fully compositional semantics for them on the one hand, and capturing our judgments as native speakers about particular cases and inference patterns, on the other. The former is what the kind theories of Carlson and Chierchia do well, and the latter is what they have essentially nothing to say about. The latter is what quantificational theories, such as those of Asher & Morreau (1995) and Asher & Pelletier (1997) do really well, and the former is what *they* have little to say about.

One response to this implicit trade-off is to bite the bullet, as Sarah-Jane Leslie does, and hold that when it comes to generic sentences, properly semantic considerations can only shed a limited amount of light on our judgments about particular cases. Essentially, this is to say that only the former goal should play a role in determining the best semantic analysis. However, biting the bullet also comes at a certain cost: namely, that it makes natural language semantics a bit odder and more difficult to motivate than it would otherwise be. In particular, taking that approach makes it less clear in what sense the semanticist is providing truth conditions.

Suppose we represent the truth conditions of an English sentence as the set of models which satisfy the logical translation of that sentence. A model, in logic, can be thought of as an abstract mathematical encoding of a worldly circumstance. But if semantics can proceed fully indifferently to whether any philosophical interpretation of the models that comprise some set of truth conditions is ready to hand, then it is time to rethink what natural language semantics is even supposed to be. If there is no straightforward way to understand those models as worldly circumstances, then we have reached a point where semantics is now doing something quite different from mapping syntactic structures to truth conditions. And if so, then we are faced with the rather difficult task of figuring out what that is.

The coming chapters opt for a simpler alternative. Why not just adopt *both* desiderata? A semantic theory should offer plausible truth conditions for some fragment of a language, make accurate predictions about inference patterns involving sentences in that fragment, *and* respect facts about the syntactic distribution of the relevant words. This set of assumptions is much more in line with how philosophers and linguists explain to their undergraduates what semantics is. Furthermore, it is more in line with standard practice in the field. For example, imagine that a philosopher would like to give a semantic analysis of an English word, *X*, favors the hypothesis that it is synonymous with the word *some*. Then showing that it patterns distributionally with weak quantifiers—that 'There are *X* cars in the parking lot' is well formed—and showing that *X* is upwards monotonic in its restrictor and scope would both count as evidence in favor of their analysis. The first is an argument from distributional patterning, and the second is an argument from inference patterns. That both kinds of argument are standard in natural language semantics is a fairly strong indication that most semanticists strive to make accurate predictions about both kinds of phenomena. Given that semanticists already have these ambitions, the default instinct should be not to scale them back unless there is some compelling reason to do so.

From a perspective which seeks to explain the distributional data, Carlson's original theory has a lot going for it. Carlson is still one of the only authors to have gotten his hands dirty with the details of a fully compositional theory. Where other theories provide logical paraphrases of generic sentences in English, Carlson's theory shows how to derive his logical paraphrases from three ingredients: some uncontroversial syntactic structures, lexi-

cal entries, and a small number of composition rules. It offers a unified semantics for bare plurals that predicts when they give rise to the existential interpretation and when they give rise to the generic interpretation. And Chierchia (1998) goes a step further, showing how a Carlsonian kind semantics can be used to account for crosslinguistic constraints on generic morphology. To the author's knowledge, no quantificational semantics for generic sentences, whether it employs determiner or adverbial quantifiers, has been able to provide any alternative account of these phenomena. In addition, further data which will be advanced in section 2.4 pose significant further problems for the quantificational theory. And so, the approach will be to develop an extension of Carlson's (sophisticated) kind semantics for generics that a) can capture all the data that Carlson originally intended to capture, b) explains why and when generic sentences tolerate exceptions, c) answers the principal semantic objections against Carlson's proposal, d) accounts for the new data to be presented, and e) lends itself to a natural philosophical interpretation. Methodologically, this approach has the advantage of showing that there is no trade-off between getting distributional facts right and explaining a statement's truth conditions.

2.3 Simple vs. Sophisticated

One last question to be addressed is: why prefer the sophisticated kind theory over the simple kind theory? What is the justification for positing a predicate modifying operation? Why not think that the noun phrase in subject position of a generic sentence is just the name of a regular object, rather than the name of some new variety of object we're going to call a kind? That is precisely what the simple kind theory, repeated below, holds. Generic sentences are nothing more than garden-variety monadic predications.

(39) **Simple Kind Theory**[22]

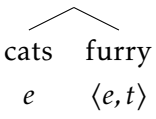

For the simple kind theorist, cat-kind is an object just like any other, and when a speaker says that cats are furry, she is saying that cat-kind has the property of being furry. Individual felines can be furry, and so can cat-kind itself. (Though of course, *what it takes* for cat-kind to be furry is different from what it takes for an individual feline to be furry.) Shouldn't we not just say this, and bypass the additional complications attendant upon giving our logic a new atomic type and positing this new predicate modifier?

[22] See Carlson (1977a), chap. 4 and Liebesman (2011).

First, it is worth nothing that there are compelling philosophical reasons for favoring the former kind theory with a predicate modifier over this simpler kind theory. These will be more fully taken up in section 3.3, but briefly, there is something intuitively rather strange about thinking that a kind can be furry in the same sense as a cat. Even the idea of a kind being furry at all is at least a little bit strange. Suppose that our theory of kinds dictates that cat-kind is the set of all cats. Clearly, a set is not the sort of thing that gets to be furry—a set is a mathematical object that gets to be the subset or member of another set, have a certain cardinality, and so on. What if cat-kind is a Platonic universal embodying the ideal of cathood? The problem is that Platonic universals, on most accounts, are located outside space and time, and being located in space and time is surely a precondition on having fur. None of the main philosophical contenders for a theory of kinds makes them out to be the sort of thing that can bear all the ordinary properties one might apply to individual objects.

Thus, the predicate modifying operation already arguably has a strong philosophical motivation. On the one hand, any property meant to apply to an individual object can be used in a generic sentence.[23] On the other, it is highly unclear what it might mean for a *kind* to have such a property. Of course, no definition of *PM* has yet been proposed.[24] But including it as part of a semantic analysis is at least a way of acknowledging that if kinds can be furry, it is only in a very different sense of the term. Doing so leads to the correct prediction that any object-level property can be mapped to a kind-level property.

The second reason this chapter will reject the simple kind theory is Carlson's original reason for rejecting it, which is that a semantic analysis of generics needs more than the simple kind theory can provide in order to account for bound variable interpretations in generic sentences. Consider this example:

(40) Dogs are loyal to their masters.

Sentences such as (40) reveal something important about generic sentences: any property that applies to individual objects can characterize a kind. But without any additional logical structure in the form of an operation which that maps properties of individual objects to properties of kinds, there is no way to ensure that sentence (40) receives the interpretation on which each dog is paired with *its* master, rather than some other dog's master. There must be a way to lambda abstract over the sentence 'Fido knows is loyal to his master,' form the property 'x is loyal to x's master,' and derive truth conditions in which *that* property characterizes dog-kind.

[23] That was one of the reasons motivating Carlson to put forth a sophisticated kind theory in the first place.
[24] That will come in section (7).

How exactly a semantic theory should account for the interpretation of pronouns is a difficult and contentious topic,[25] but for the purposes of this example, we may treat *loyal to its master* as a single predicate with the following denotation:

(41) $[\![\textbf{loyal to its master}]\!]^{M,g} =$
λx . x is loyal to x's master =
λx . $loyalTo(x, \iota y(masterOf(y, x))$

In other words, it is a predicate denoting the property of being loyal to one's master. Next, suppose a simple kind theory, according to which this property is predicated of dog-kind, which is an object like any other:

(42) $[\![\textbf{dogs}]\!]^{M,g}$ = dog-kind $\in D_e$

If we try to create the lambda-abstracted property of being loyal to one's master, then apply dog-kind directly to that, the truth conditions of sentence (40) come out to the following:

(43) λx . $loyalTo(x, \iota y(masterOf(y, x))$(dog-kind) =
1 iff $loyalTo$(dog-kind, $\iota y(masterOf(y,$dog-kind))) =
1 iff dog-kind is loyal to its master

And those are clearly the wrong truth conditions. One reason is that it was already perplexing what it might mean for dog-kind to be furry; it is even more perplexing what it might mean for dog-kind to have a master. Nonetheless, suppose we grant that we can make sense of that idea, if only for the sake of argument. Even then, the truth conditions in (43) entail that whatever it means for dog-kind to have a master, it has only one. But of course, the truth conditions for sentence (40) are not only such that different dogs need to be able to exhibit loyalty to different masters, but such that each dog must exhibit loyalty to its own master. Our semantics needs to be able to distinguish between these three situations:

(44) a. If d is a dog, then d's being loyal to d's master is a matter of course.
b. If d is a dog, then d's being loyal to some dog's master is a matter of course.
c. If d is a dog, then d's being loyal to dog-kind's one and only master is a matter of course.

More specifically, it needs to be true only in the first situation—not in either of the second two.

Positing a predicate modifying operation resolves this problem. For the time being, assume that that the predicate modifier means something like 'is characterized by the fol-

[25] For an interesting exploration of some of these difficulties, see Dekker (2012).

lowing property.'[26] Think of it as a kind of dummy definition, to be replaced with something more substantive later on, but which has the general shape of the function that is called for:

(45) $[\![\textbf{PM}]\!]^{M,g} = \lambda f_{\langle e,t\rangle} \,.\, \lambda k \,.\, characterizedBy(k, f)$

Now suppose we assign *dogs* a special atomic type dedicated to kinds:

(46) $[\![\textbf{dogs}]\!]^{M,g} = \text{dog-kind} \in D_k$

If the syntactic structure of (40) is as in (45), then its truth conditions come out as follows:

(47) $(\lambda x \,.\, loyalTo(x, \iota y(masterOf(y, x)))(\lambda f_{et} \,.\, \lambda k \,.\, characterizedBy(k, f))(\text{dog-kind}) =$
$(\lambda x \,.\, loyalTo(x, \iota y(masterOf(y, x)))(\lambda f_{et} \,.\, characterizedBy(\text{dog-kind}, f)) =$
1 iff $characterizedBy(\text{dog-kind}, (\lambda x \,.\, loyalTo(x, \iota y(masterOf(y, x)))) =$
1 iff the property of being loyal to one's master is characteristic of dog-kind

Of course, the viability of this as a solution to the problem turns on how we choose to understand the notion of a property characterizing a kind. So far, nothing has been anything about what it might mean for a property to characterize a kind; and indeed, Carlson argues that semanticists *should not* try to say anything about what it might mean. The next chapter parts company with Carlson on this point and provides a philosophical account of what it means for a property to characterize a kind. Without going into that account just yet, it may help to regard this characterizing relation in the terms suggested by Thompson (2009): loyalty to one's master characterizes dog-kind just in case the life form of a dog—sort of a script for how a dog's life is supposed to go—involves being loyal to one's master. Granted, going all in on this definition would mean specifying many details about what life forms are, what it means for different properties to be involved in them, and so forth. It would also involve addressing the fact that the above derivation concludes with an unreduced lambda expression, which is undesirable. Nonetheless, even this quick formulation is enough to show that the truth conditions given in (47) satisfy those situations in which different dogs are incorrectly paired with different masters. Why? Because the property of being loyal to one's master is in the scope of the characterizing relation.

In short, there are compelling reasons, both philosophical and semantic, for favoring a kind theory that uses a predicate modifier over its simpler cousin. Of course, it is reasonable to have reservations at this juncture about what the predicate modifying opera ton is actually supposed to do. The predicate modifying operation still has the feel of a formal cheat; a black box which, by stipulation, magically outputs the precise result we want without our having to explain how. But section 3.3 will argue that the predicate modifier has an independent philosophical motivation; that rather than being an ineffable skyhook

[26] This assumption will be substantially revised in Chapter 3.

brought in to save a theory that seems to require it, there is an important fact about the truth conditions of generic statements that it is uniquely suited to model.

2.4 Quantificational vs. Kind Theories

Next, we turn to some new linguistic data. This section proposes the following strategy for determining whether generic statements are general statements about individuals or particular statements about kinds. Quantification is the paradigmatic tool for making generalizations in natural language. So if generics are true generalizations, we should expect them to exhibit the behavior of statements with quantifiers in them (hereafter, *quantified sentences*). If they do not behave in this way, that already is a clear indication that they do not implicitly contain a quantifier. Furthermore, the distinctive feature of quantification is that it is indifferent to the set of things being generalized over; as long as the restrictor predicate has a well-defined extension, that is enough. It will be argued that whatever it is that generic sentences are about cannot be any old collection of things; there are constraints on which predicates can appear in generic sentences. This is an additional further indication that generic statements are not generalizing over individual objects; if they were, then they could generalize over any set of objects. The fact that they cannot generalize over any old set of individual objects is a sign that they are not really generalizing in the first place. Rather, they are referring to a different sort of object.

The remainder of this chapter will present three contrasts between generic sentences and quantified sentences that have received no discussion in the literature, as of yet. First, generic sentences do not contextually domain restrict. That is the most important contrast: contextual domain restriction is the hallmark of natural language quantification, from determiner quantifiers to adverbial quantifiers and even modal auxiliaries. If generic sentences do not domain restrict, that should give the quantificational analysis serious pause. But there are two other contrasts as well, which are arguably related. One is that generic sentences are more selective than quantified sentences about what kind of predicate can occupy the subject NP. Quantifiers will accept more or less any predicate with a denotation, but generic sentences seem to require something more, which will hereafter be labelled *cohesion*. The other contrast is that generic sentences exhibit a variety of context sensitivity that quantified sentences do not—they can vary as to whether they are interpreted artifactually or non-artifactually.

Contextual Domain Restriction

One of the characteristic features of natural language quantifiers is a certain kind of context dependence.[27] Suppose a teacher walks into class and says:

(48) Is every student here?

In that context, the teacher is obviously not asking whether everyone in the entire universe is there; she is asking whether everyone who is enrolled in the course is present. Of course, she never actually said 'who is enrolled in the course.' That part was left implicit. This phenomenon is referred under the heading of *quantifier domain restriction*, because of the intuition that in sentences like (48), rather than generalizing over the extension of the explicit restrictor predicate—*student*, in this case—the speaker generalizes over a subset of that extension. Thus, the domain of quantification—the set of all students—is *restricted* to one of its subsets—the set of all students in this class. Quantifiers in natural language are rarely, if ever, used unrestrictedly; they nearly always exhibit this type of context dependence.[28]

A natural question to ask when considering whether generic sentences contain a quantifier, then, is whether they also contextually domain restrict. Tellingly, they do not. To see the contrast, imagine a reporter for an animal rights magazine. This reporter hears that Wayne Newton has opened up a new annex in his ranch for exotic animals that is expressly dedicated to jaguars. The magazine flies her over to the ranch to do some investigation; after all, it is important to her that these animals live a happy life. The reporter then spends a few days doing a thorough tour of the ranch, eventually discovering, to her great chagrin, that Mr. Newton's jaguars have been given identificatory tattoos on the inside of their ears, to aid the facility's managers in keeping track of them. From the reporter's point of view, of course, this is needlessly painful and thus morally abhorrent. So her editor flies over to the ranch in order to be able to observe what is happening first hand. The reporter picks her editor up at the airport, then drives straight to the ranch. While opening the door, she proclaims:

(49) (*opening the door*)
It's really something. Every jaguar has a tattoo.

This discourse is perfectly felicitous (and true). Why? Because although it is false that every jaguar in the world has a tattoo, it is true that every contextually salient jaguar—every jaguar

[27] For some classic discussion of contextual domain restriction, see Stalnaker (1970), pg. 276 and Lewis (1979), Example 3. More modern treatments can be found in von Fintel (1994) and Stanley & Szabó (2000).

[28] An exception may be mathematical statements, which the quantifiers of first-order logic were originally developed in order to model.

on the ranch—has a tattoo. But now compare the following alternative discourse, with the corresponding generic in place of the quantified sentence:

(50) (*opening the door*)
It's really something. #Jaguars have tattoos.

In this context, the corresponding generic sentence is at least false, and probably also infelicitous. Why? No interpretation on which it concerns only jaguars on the ranch is available. Saying that jaguars have tattoos in this context sounds like a non-sequitur, because its only possible interpretation is one on which it concerns jaguars in general.

Here is a second example. Imagine Willem de Vlamingh's first mate, at the moment they first stumbled across the Swan River in Australia in 1697. In that context, de Vlamingh's first mate would not have been able to use a generic statement about 'swans' to express his astonishment as they both stared dumbfounded at a flock of black swans. But a quantified statement about swans would be perfectly appropriate:

(51) a. I can't believe my eyes! Every swan is black.
b. I can't believe my eyes! #Swans are black.

A crucial distinction between generic sentences and quantified sentences is encapsulated in these examples.[29] Still, one might wonder there is anything special about them. Perhaps there are other special circumstances in which generic sentences contextually domain-restrict. It is not unreasonable to think, on first glance, that examples in the vein of Condoravdi (1992, 1997) are indeed cases in which generic sentences contextually domain-restrict.[30]

We now turn to one of these potential counterexamples. In most contexts, the following sentence is false:

(52) Squirrels are friendly to people.

But mentioning a location at which squirrels behave unusually may suffice to give the above sentence a restricted interpretation:

(53) Washington Square Park is quite a place. Squirrels are friendly to people.

Why might it be true to say that squirrels are friendly to people in this new context? Presumably it would have something to do with the fact that squirrels in this location are subject

[29] For a related example, see Asher & Pelletier (1997), pp. 1165-1166.
[30] Condoravdi's original example was the following:

(i) A ghost has been haunting campus. Students are afraid.

This particular example is less than ideal as a counterexample, because *afraid* is a stage-level predicate, and so it is difficult to hear the second sentence of (i) as a generic. But closely-related examples such as the one presented below suggest themselves.

to an unusual amount of tourist traffic, which has led them to evolve a distinctive set of behavioral habits over the past few generations. Somehow, mentioning the park earlier in the discourse has made available a reading on which the speaker is talking about squirrels in the park, rather than squirrels in general.

The first thing to say about these Condoravdi-style examples is that judgments about them are shaky. A weak majority of native speakers the author consulted found them well-formed, but a significant minority find them ill-formed, strongly preferring variations that begin with 'squirrels there' or 'in Washington Square Park, squirrels...' So their ultimate status as counterexamples to the claim that generic sentences don't domain restrict is at least somewhat in question.

But of course, to leave the argument there would be to proceed in bad faith. Suppose for the sake of argument that sentences like (53) are admissible as data. Even then, it seems unlikely that they involve anything like *domain restriction*. One reason is that is was observed above, a key feature of quantifier domain restriction is that it happens by default. The above phenomenon, whatever it may be, only happens in particular circumstances. Essentially, it only happens when the speaker marks a location off as remarkable earlier on in the discourse.

A natural thought for a quantificational theorist to have at this point is that there is a good exmplation for all this: namely, that the generic quantifier cannot domain restrict *deictically*; it can only domain restrict *anaphorically*. The examples involving jaguars and swans were all deictic, in the sense that the information about how the predicate explicitly mentioned is being restricted comes from information perceptually available to the speaker and listener. Example (53) is an example of anaphoric domain restriction, in the sense that information about what to restrict the predicate *squirrels* to—Washington Square Park squirrels—comes from a location mentioned earlier in the discourse—Washington Square Park. So perhaps it is possible to say that there is a generic quantifier, but the reason it doesn't contextually domain restrict in the examples we originally looked at is that there is nothing earlier in the discourse on which the contextual restriction can base itself. Though not incredibly common, this distinction between anaphoric-deictic and anaphoric-only phenomena has been known to arise from time to time.[31]

The problem with taking that route is that if there were a generic quantifier that was anaphoric-only (as it were), then mere mention of the location in advance would suffice to license the Condoravdi phenomenon. But mere mention of the location in advance does not suffice to license the Condoravdi phenomenon:

[31] Klecha (2011) argues that *gonna* can domain restrict either deictically or anaphorically, whereas *will* can only domain restrict anaphorically.

(54) a. Yesterday I bought peanuts in Washington Square Park. ??Squirrels are friendly to people.
b. Yesterday I bought peanuts in Washington Square Park. Every squirrel is friendly to people.

Further variations prove comparably awkward, and the quantifier is always a more natural fit:

(55) a. Have you been to Never Never Land? ??People/Everyone can fly.
b. In Washington Square Park, squirrels are given euphoria-inducing drugs. ??Squirrels are/Every squirrel is friendly to people.
c. In Washington Square Park, tourists have been feeding animals for years. ?Squirrels are/Every squirrel is friendly to people.

Thus, there is independent reason to think that something other than domain restriction is happening in the Condoravdi cases. If this truly were an example of contextual domain restriction, it would have to occur in a much wider range of environments.

The Cohesiveness Presupposition

The second distinguishing feature of generic sentences involves additional constraints on what kind of predicate they will accept in subject position—constraints which are absent from quantified sentences. In particular, the NP in subject position of a generic is resistant to certain kinds of heavily modified predicates. That is not to suggest that it is resistant to *all* heavily modified predicates. The following, for instance, all sound just fine:

(56) a. Rabbits are skittish.
b. Rabbits from Mexico are skittish.
c. Rabbits with fluffy fur are skittish.

These sentences may or may not be true, but they are certainly equally felicitous. However, Carlson (1982) observed that the moment one modifies the NP in subject position of a generic with an indexical expression, the generic in question begins to sound awkward:[32]

(57) a. #Toppings on this pizza are vegetarian.
b. #Chairs in that house are made of oak.
c. #Desks that I am looking at right now have metal tops.

The first remark to make in this connection is that quantified sentences exhibit no such selectivity:

(58) a. Every/some/most topping(s) on this pizza is/are vegetarian.

[32] See Carlson (1982), pg. 153.

b. Every/some/most chair(s) in that house is/are made of oak.
c. Every/some/most desk(s) that I am looking at right now has/have a metal top.

The sentences in (57) are at least false and probably also infelicitous, even in the circumstances in which the sentences in (58) are true.

Why might that be? The position of this chapter is that these are signs that generic sentences are not engaged in true generalizing, where generalizing is the kind of purport we associate with a quantifier. Quantificational generalizing is maximally general, in the sense that any set of objects is fair game for a quantificational generalization. Any group consisting of more than one object can be generalized over. Generic sentences, by contrast, cannot be made about any old set of objects. The set of objects under discussion in a generic are required to have something further in common. What is this additional something extra? Intuitively, this additional something extra has something to do with the distinction between predicates with rhyme and reason to them, and predicates that seem more concocted or arbitrary.

The contrast is most stark with demonstrative subject NPs, but there is arguably at least a little bit of a contrast with other kinds of predicate as well. Take, for instance, predicates the members of whose extensions have nothing very interesting in common:

(59) a. ?Rabbits who lived three houses away from Matt are skittish.
b. Every rabbit who lives three houses away from Matt is skittish.

(60) a. ?Rabbits located at prime-numbered longitudes are skittish.
b. Every rabbit located at a prime-numbered longitude is skittish.

These generics sound considerably better than their demonstrative counterparts, but still sound at least somewhat awkward in most contexts. The corresponding quantified sentences sound a good deal more natural. There is a lot more to be said about this, but for a first pass at generalizing from these observations: predicates like *rabbit from Mexico* just seem more kindlike than predicates like *rabbit located at a prime-numbered longitude*. Rabbits from Mexico, there is a temptation to say, form a kind, whereas rabbits located at prime-numbered longitudes just happen to be located at prime-numbered longitudes.

To mark this (still rather murky) distinction, let us introduce *cohesive* as a term for predicates with the feature under discussion, whatever it is, and *haphazard* as a term for any predicate that is not cohesive. We will say that generic sentences come with a presupposition that the predicate in subject position is cohesive. Take the above three indexically-tinged predicates as paradigm cases of haphazard predicates. There is a great deal to be said about what exactly the objects falling under the extension of the predicate in subject position of a generic sentence need to have in common, in order for the generic sentence to

be felicitous, and going into it in detail here would take us rather far afield. However, in this context of this discussion, two observations will suffice. First, the class of predicates that quantified sentences will accept as restrictions is quite inclusive in comparison with the class of predicates that a generic will admit in subject position. Second, the class of predicates that a generic sentence will admit in subject position lines up rather closely with what philosophers have called *sortal predicates*.

The sortal/non-sortal distinction comes out of Strawson (1959), Geach (1980), and Wiggins (2002), who trace it back to Gottlob Frege, Thomas Aquinas, and Aristotle.[33] It very roughly lines up with the distinction between nominal predicates, on the one hand, and verbal and adjectival predicates, on the other. Here is what Peter Geach has to say about sortal predicates (which he calls *substantival*):

> This brings us, not yet to the notion of a substantial term, but at least to that of a *substantival* term. Aquinas calls out attention to a feature of Latin grammar—that substantives are singular or plural on their own account, whereas adjectives 'agree in number' with substantives. This suggests to him a logical distinction between two sorts of terms: substantival terms, to which the question 'how many?' applies directly, and adjectival terms, to which this question applies only insofar as they are used to add a qualification to substantival terms. One may ask how many cats there are in a room; but not, how many black *things* there are in a room; only how many black *cats* (say) there are in the room. The basis of this distinction is that the sense of 'cat' determines a sense for 'one and the same cat,' whereas the sense of 'black thing' does not in the least determine what shall count as one and the same black thing.[34]

It is helpful to break the thoughts expressed in this passage down into two components. First, the philosophical intuitions behind sortal predicates have to do with what is required to be competent at deploying them. In order to be competent in the use of a non-sortal predicate F, one needs to know how to determine what falls under F's extension. But in order to be competent in the use of a sortal predicate K, one (at least) also needs to be able to:

(61) a. Determine, for any two objects, whether they are *the same K*.
b. Count Ks.[35]

[33] For a good contemporary overview, see Lowe (2009).
[34] Geach (1961), pg. 86.
[35] Certain accounts of sortal predicates collapse these two conditions, but I will follow Geach (1980) in assuming that competence in (61-a) is a necessary but not sufficient condition for competence in (61-b).

The second component of Geach's ideas lies in their linguistic ramifications. The aforementioned authors very interestingly noted that this set of philosophical intuitions also manifests itself as a set of grammatical tests:

(62) Jean: What is that on the horizon? *Metaphysics* 7.1, Wiggins (2002)
Joan: It's a whale.

(63) Jean: What is that on the horizon?
Joan: ??It's a green thing.

(64) a. This and the animal you saw last night are in fact the same whale.
b. ??This and the tennis ball you saw yesterday are in fact the same green thing.
Geach (1980), Gupta (1980)

Sortal predicates can be used to answer *What is it?* questions, and they can appear in *the same K as* constructions.

What does it take for a predicate to be cohesive—to qualify as the subject of a generic sentence? As mentioned above, this is a substantial question, which will be taken up in detail in section 4.1. However, at this point it can at least be said that the class of cohesive predicates, the category selected for by the indexical-generic diagnostic given earlier, lines up rather closely with the class of sortal predicates, the category selected for by the above three diagnostics. Cohesive predicates must at least be sortals.

To see this, we may run several predicates we already know to be haphazard through the tests for sortal predicates:

(65) a. Bob: It's dark in here. Can you see what that is in front of us?
Biff: ??It's a chair in this house.
b. Bob: It's dark in here. Can you see what that is in front of us?
Biff: ??It's a desk that I'm looking at right now.

(66) a. Can you believe that this and the heavy thing in the box you helped me move yesterday are actually the same chair?
b. #Can you believe that this and the heavy thing in the box you helped me move yesterday are actually the same chair in this house?

Given the way those examples go, we might expect non-sortal predicates to be more awkward in subject position of a generic sentence than they are in restrictor position of a quantified sentence. And that is just what we find:

(67) a. Plants are inanimate.
b. ??Green things are inanimate.
c. Most green things are inanimate.

In ordinary contexts, adjectival predicates make for odd generic sentences. One reason is that to even form a generic sentence with an adjectival predicate in subject position, it needs to be fit into the mold of a noun phrase through combination with a vacuous nominal predicate like *thing*, *object*, or *stuff*, which makes such generic sentences sound stilted. But more to the point, in most contexts, there is something woefully underspecified about the predicate *green thing*. Green *what*, one wants to ask.[36] Philosophers typically try to capture this underspecification by saying that green things have no *identity criteria*: no principle in virtue of which any particular green thing is the same as or different from any other.

It is important to recognize that this not the case in all contexts. Given the appropriate situation, a predicate like *green thing* can take on the kind of explanatory significance necessary for it to perform in a generic sentence. Suppose, for example, that a pair of friends is in the path of a charging bull. In that situation, one friend might turn to the other and say:

(68) Quick! Hand me something red! Anything red! Red things are useful for fending off charging bulls.

The same goes for the earlier cases. In a situation where the pizza before the speaker is of special significance, a generic sentence about toppings on it sounds remarkably improved. Suppose the world's greatest pizza chef has prepared some dough with a little oil, cheese, oregano, and tomato sauce, and has left it up to the speaker to decide how to top it. In that situation, it would not be unreasonable for her to say something like:

(69) Toppings on this pizza are strictly optional. It'll be great no matter what.

Whether a predicate is cohesive, then, is highly context-dependent; it varies with the explanatory purposes of the conversational participants. This is a nontrivial wrinkle in the data regarding cohesion.

If the cohesion of a predicate is context-dependent in this way, and the diagnostics for cohesive predicates agree in output with the diagnostics for sortal predicates, then this observation brings with it a fairly significant consequence: namely, that the sortal/non-sortal distinction is context-dependent in just the same way. And indeed, it would seem that this prediction is borne out. Imagine that Jean and Joan are observing someone at whom a bull is charging, and Jean is color-blind. One might then imagine either of the following two conversations taking place:

[36] Surely that must be what Aquinas had in mind when he wrote that nominal (or substantival) predicates 'carry their subject with them,' whereas adjectival predicates 'add the thing signified to the substantive.' See Aquinas I: Q. 39, Art. 5.

(70) Jean: What is he pulling out of his pocket?
Joan: It's a red thing. He's going to distract the bull and made a break for it!

(71) Joan: He's using the same red thing he used last week to distract another bull.

Assuming that is correct, the distinction between sortal and non-sortal predicates is just as context-relative as the distinction between cohesive and haphazard predicates. It is not quite as clear-cut as saying that such-and-such are the sortals, and so-and-so are the non-sortals, *point finale*. Rather, convincingly demonstrating a given predicate to be sortal or non-sortal in a given hypothetical context will require us to make sure we don't accidentally underspecify that context.

But even though the phenomenon of predicate cohesion is somewhat subtle, the contrast between generic sentences and quantified sentences is perfectly stark. It may indeed be that for any sentence featuring a haphazard predicate, one can, given sufficient time and imagination, dream up in a context in which it would be cohesive. Nonetheless, given any predicate that is clearly difficult to use in a generic sentence, in a particular context, one can always use that same predicate felicitously in a quantified sentence. And so, this is another respect in which generic sentences are choosier than quantified sentences about what kind of predicates they are willing to accept.

Artifactual Interpretations

The third contrast this section will discuss involves a kind of context-sensitivity that is present only in generic sentences—not in quantified sentences. Generic sentences are often susceptible to both artifactual and non-artifactual interpretations. But quantified sentences tend to allow for only one of the two. The relevant examples are cases where the same kind seems to be ascribed different properties depending on whether it is being regarded *qua* artifactual kind or *qua* non-artifactual kind.

One of the most interesting examples of artifactual interpretations in generic sentences is due to Nickel (2008):

(72) a. Dobermans have floppy ears. *true when uttered by biologists*
b. Dobermans have pointy ears. *true when uttered by dog breeders*

Presumably, the first sentence would be true in a context where evolutionary biologists were comparing Doberman Pinschers with, for instance, German Shepherds, whose ears naturally grow to be pointy. Dobermans are born with floppy ears that are traditionally then cropped at a young age to come out pointy. So the first sentence, when true, would capture that fact about their phenotypic characteristics. It would be true on the non-artifactual interpretation. The second sentence would be true in a context where the conversational participants were comparing the features of different breeds in view of their cultural role.

This is the more familiar context, given that it is, in fact, relatively uncommon to see a Doberman with natural ears. The second sentence, when true, would be used to capture the fact that in order to fulfill the cultural roles we have prescribed for them (being recognizable as dobermans, being threatening, having an enhanced ability to hear intruders), Doberman Pinschers are typically given pointy ears. So here we would have the artifactual interpretation, in the sense that the property is being applied to dobermans *qua* cultural artifact.

These readings are not so easy to achieve with quantifiers. Quantified sentences are resistant to artifactual interpretations:

(73) a. No Dobermans have pointy ears. *false*
b. No Dobermans have floppy ears. *true*
c. Every/most Dobermans has/have pointy ears. *true*
d. Every/most Dobermans has/have floppy ears. *false*

The effect is especially dramatic with *no*. In most contexts, it is quite difficult to hear sentence (73) as stating that no doberman is biologically disposed to grow pointy ears, at least not without considerable coercion.

With some effort, it is possible to finesse artifactual interpretations out of quantified sentences, but usually only with the help of illocutionary operators like 'really,' which effectively signal to the hearer that the utterance she is about to hear should be interpreted in some nonstandard way:

(74) a. Really, every doberman has floppy ears. *true*

But the takeaway point should be that generic sentences require considerably less contextual nudging in order for the artifactual interpretation to become available. The only thing required is a certain shared context of inquiry between speaker and hearer:

(75) a. Dobermans aren't like German Shepherds. German Shepherds are the ones with pointy ears. Dobermans have floppy ears. *true*
b. Dobermans aren't like German Shepherds. German Shepherds are the ones with pointy ears. ??Every Doberman has floppy ears. *awkward/confusing*
c. Every German Shepherd is different from every Doberman. Every German Shepherd has pointy ears, but every Doberman has floppy ears. *false*

These examples bring attention to a third difference in how generic sentences pick up information from the conversational context and the way quantified sentences do so.

2.5 A Kind-Theoretic Analysis

In view of the above considerations, then, the prospects for a quantificational analysis of generic statements look somewhat dim. Generic statements are susceptible to artifactual/non-artifactual interpretations, they require cohesive predicates in subject position, and they do not contextually domain restrict. A quantificational analysis would wrongly predict generic statements to pattern with quantified statements on these three fronts. A kind-theoretic analysis, on the other hand, not only makes the correct predictions about these three data points, but offers a unified explanation for them. Here is a broad outline of how such an analysis would look in its most minimal form.

First, all generic sentences would come with a presupposition that the noun phrase in subject position refers to a kind. At this stage of the argument, we need not commit ourselves to any particular metaphysical theory of what a kind is. For the purposes of this semantics, we need only assume that for every cohesive predicate S, there exists a principle that determines whether anything in its extension is the same S as anything else. Call this principle *S-kind*, and call the function that maps predicates to their kinds f. For other predicates—the haphazard ones—there exists no such principle, and therefore f will not be defined on them.

So testing whether a predicate is cohesive will amount to testing whether any kind is associated with it. Giving a predicate the plural suffix but no article when it is in subject position will map that predicate to the kind associated with it; and where there is no such kind, the semantic derivation will crash. In other words, the generic construction comes with a cohesiveness presupposition. For example, sentence (56-b), repeated below, presupposes that the predicate *rabbit from Mexico* has a kind associated with it; that the noun phrase *rabbits from Mexico* refers to a kind. And indeed there is such a kind, so it is felicitous:

(76) Rabbits from Mexico are skittish.
Truth Conditions:
true if $f(\llbracket$**rabbit from Mexico**$\rrbracket)$ satisfies the predicate *skittish'*
false if $f(\llbracket$**rabbit from Mexico**$\rrbracket)$ doesn't satisfy the predicate *skittish'*
anomalous if $f(\llbracket$**rabbit from Mexico**$\rrbracket)$ is undefined[37]

Sentence (57-a), repeated here, presupposes that the predicate *topping on this pizza* refers to a kind. Given that there is no such kind, in ordinary contexts, it is infelicitous:

(77) #Toppings on this pizza are vegetarian.
Truth Conditions:

[37] Read $\llbracket\cdot\rrbracket$ as the denotation function. So for any English phrase 'A B C,' $\llbracket$**A B C**$\rrbracket$ refers to the denotation of the phrase 'A B C.'

true if $f(\llbracket$**topping on this pizza**$\rrbracket)$ satisfies the predicate *vegetarian'*
false if $f(\llbracket$**topping on this pizza**$\rrbracket)$ doesn't satisfy the predicate *vegetarian'*
anomalous if $f(\llbracket$**topping on this pizza**$\rrbracket)$ is undefined

As for the relation between the cohesiveness presupposition and the availability of artifactual interpretations, there is a natural connection to be drawn. To see it, consider examples (72-a) and (72-b), repeated here:

(78) a. Dobermans have floppy ears. *true when uttered by biologists*
b. Dobermans have pointy ears. *true when uttered by dog breeders*

Earlier, it was hypothesized that sentence (78-a) is true because the property of having floppy ears characterizes biological dobermans, and (78-b) is true because the property of having pointy ears characterizes artifactual dobermans. A natural way to account for this phenomenon is to suppose that artifactual dobermans form a subkind of biological dobermans. The subkind relation can then be defined as follows: K_s is a subkind of K just in case:

(79) a. For some predicate S_s, K_s is a principle determining whether anything in the extension of S_s is the same S_s as anything else.
b. Every member of K_s is also a member of K.[38]

Assume that in sentences (78-a) and (78-b) everyone is aware that Dobermans do not naturally have floppy ears. Then it becomes possible to say that sentence (78-a), is false under its most obvious interpretation, one on which it's talking about artifactual dobermans. Artifactual dobermans are not characterized by having floppy ears. But biological dobermans, a closely-related superkind of artifactual dobermans, is characterized by having floppy ears. So sentence (78-a) is reinterpreted so that it is about biological doberman-kind, rather than artifactual doberman-kind.

The full story about how this works will come in Chapter 4, where we replace this sortal-inflected sub/superkind relation with a new notion that is easier to manage, which will be called *comprisal*. But for the purposes of this discussion, the following should at least be a suggestive sketch of how a kind theory will account for these three phenomena. 'Bears are furry' will be true just in case the (contextually salient) sub or superkind of the (contextually salient) sortal principle associated with the predicate *bear* is characterized by furriness. Leave what being characterized by furriness is to one side for now—that will be the concern of the following chapter—and focus instead on the 'principle' clause. On this

[38] Note that this definition presupposes a Geachian conception of sortals, which allows for relative identity. This definition isn't possible under Wiggins' conception, because Wiggins is committed to saying that the principle which makes e.g. this giraffe the same giraffe as the other giraffe has to be the very same principle that makes it the same mammal as the other giraffe.

account, 'Toppings on this pizza' is infelicitous, because *topping on this pizza* is not a sortal predicate; there is no sortal principle underlying it. 'Dobermans have floppy ears' comes out true because a superkind of the sortal principle underlying the predicate *doberman* is characterized by floppy eared-ness. And 'Squirrels are friendly to people' comes out true when preceded by 'Washington Square Park is quite a place' (assuming we accept that sentence) because although squirrel-kind isn't characterized by friendliness, a new subkind of squirrel-kind that was introduced into discourse by the previous sentence is. That subkind is Washington Square Park squirrels.

The purpose of the next two chapters is to unpack the account just sketched into a semantic analysis.

Chapter 3

The Predicate Modifier

The previous chapter examined three key differences in the behavior of generic sentences and quantified sentences. If anything is the hallmark of determiner quantifiers, it is the ability to undergo deictic, anaphoric, and bound variable contextual domain restriction.[1] But generic sentences do not contextually domain restrict. And so, any significant evidence that generic sentences lack this ability is evidence that they do not contain an implicit determiner quantifier.[2] Furthermore, unlike quantified sentences, generic sentences also give rise to artifactual interpretations and come with a cohesiveness presupposition.

These observations necessitate a re-examination of the traditional idea that such sentences are statements about kinds, rather than generalizations over individuals.[3] Of course, most of the considerations adduced thus far speak more clearly against a quantificational theory than in favor of a kind theory. The task of Chapter 4 is to provide the second half of the story, showing not only that the kind theory avoids the incorrect predictions of the quantificational theory, but that it makes its own correct positive predictions in an explanatorily satisfying way.

But before getting there, there is more work to be done. Although the kind theory seems like a promising alternative to the quantificational theory, it faces a number of well-known objections. If it is to be of use in accounting for the data presented in Chapter 2, it

[1] This typology is drawn from Partee (1989). Note that it is unclear how to check for bound variable domain restriction in generic sentences, since the usual constructions involved in such examples involve contextually restricting the domain of a quantifier in object position, and we have restricted our attention to generic statements with bare plurals in subject position.

[2] Advocates of a quantificational theory are sometimes inclined to argue that this is because the generic quantifier is intensional. However, it is unclear what the basis for so arguing would be. The best known example of an intensional determiner quantifier, *many*, patterns with the extensional quantifiers when it comes to the data presented in Chapter 2. For more on the intensionality of *many*, see Bastiaanse (2014).

[3] As originally proposed in Carlson (1977a; 1977b; 1982), and later defended in Chierchia (1998), Thompson (2009), and Liebesman (2011).

must first be refined so as to answer the standard objections that authors have raised against kind theories in general. The task of this chapter is to prepare the sophisticated kind theory to address these worries, before then putting it to work. It will demonstrate that the standard semantic objections do not hold against the more sophisticated (and largely ignored) version of the kind theory from Carlson (1977a), chap. 5. It will then develop an extension of that theory which has a natural philosophical interpretation, leaving semanticists free to enjoy the advantages of a kind theory in accounting for the data from Chapter 2 without incurring any further costs.

The trajectory of the argument is as follows. Chapter 2 claimed that the sophisticated kind theory is not vulnerable to any of the serious semantic objections that one might make against the simple kind theory, and should therefore be preferred. At the current stage of the argument, however, the status of the sophisticated kind theory still looks purely hypothetical: if there were such a theory, it would be a convenient response to the standard objections. However, there are serious questions about whether the sophisticated kind theory is a 'black box' solution, which magically outputs the correct answer by stipulation, without any satisfying explanation as to why. This chapter thus begins by fleshing the 'black box' worry out. It then uses this worry to introduce a philosophical picture which, when used to interpret the sophisticated kind theory, will demonstrate it to have independent philosophical motivation, in addition to its independent semantic motivation. By following this method, the semantics utilized by the sophisticated kind theory will not merely make the correct predictions regarding the truth conditions of generic sentences—which is the principal goal—but as a sort of bonus, there will also be a clear way to make intuitive sense of what it understands the truth conditions of a generic to be.

It would be a mistake to confuse the semantic theory with its philosophical interpretation. It is not part of the view defended here that the behavior of generic sentences in English provides evidence in favor of the metaphysical picture about to be sketched, in all of its detail. Rather, the principal claim of this text is that the semantic behavior of generic sentences in English offers reason to adopt the semantic analysis put forth in this and the subsequent chapter. That semantic analysis, if correct, indicates that the linguistic behavior of generic sentences in English bespeaks an implicit commitment on the part of native speakers to a set of metaphysical commitments—again, not the intuitive metaphysical picture about to be sketched, in all of its detail, but to the austere, abstract set of metaphysical commitments that the model-theoretic analysis brings along with it. Finally, it is entirely possible to disagree, either with the intuitive metaphysical picture about to be described or with the set of abstract metaphysical commitments the model-theoretic semantics brings along with it, and yet be in complete agreement with the main argument of this text. The philosophical interest of uncovering these abstract metaphysical commitments lies in their ability to stimulate our intuitive thinking about the metaphysical issues, leading us to con-

sider viewpoints we would previously have overlooked. There is a fruitful dialogue to be had by going back and forth between our pre-theoretical metaphysical intuitions and the logical features of natural language, each informing the other in successive cycles.

Having laid the philosophical foundations for the sophisticated kind theory, this chapter will then provide a lexical entry for the predicate modifying operation posited in Chapter 2. In order to issue in the standardly accepted result that generic statements are intensional in purport, *PM* will be given an intensional semantics.[4] With this semantics for *PM* in place, it will be possible to show that the sophisticated kind theory has all of the logical power and flexibility of the quantificational theory without any of the quantificational theory's disadvantages, thus leaving the way open for Chapter 4 to put it to work as an analysis of the data presented in Chapter 2.

The final section of this chapter will discuss the ultimate philosophical difference between quantificational theories and kind theories. The sophisticated kind theory defended here, which uses the intensional semantics for *PM*, is, in a certain sense, a version of the popular normality theory of generics (Nickel, 2008, 2010b). But it is the normality theory with a new twist. In particular, the notion of normality that this semantics throws into relief enjoys certain advantages over previous notions. As with most philosophical decisions, the choice between different normality theories involves a cost/benefit trade-off. Adopting the new notion of normality suggested by the semantics for *PM* makes it easier to explain certain philosophical puzzles, and harder to explain others. Indeed, part of the interest in doing natural language metaphysics in this way is that it prompts us to explore new approaches to these philosophical problems which would never otherwise have occurred to us.

3.1 The Black Box Worry

A kind theory of generic statements holds that they are monadic predications, rather than generalizations over individuals. According to the simple kind theory, 'cats are furry' means 'cat-kind is furry.' According to the sophisticated kind theory, 'cats are furry' means something like 'cat-kind is characterized by furriness,' where the expression 'being characterized by furriness' is prosaic shorthand for 'having the property to which *PM* maps the object-level property of being furry.' Both versions of the kind theory immediately raise all sorts of philosophical questions. What exactly is cat-kind supposed to be, and in what sense can it be furry? For whatever one's preferred philosophical theory of kinds—whether they be Platonic forms, Aristotelian universals, cognitive prototypes, cultural stereotypes, linguis-

[4] Once again noting Nickel (forthcoming) as an exception to the consensus view.

tic constructs, or mereological sums of individuals—it is not at all clear that a kind is the sort of thing to which the concept of furriness applies. Perhaps a kind can be furry, in some sense of the term; but certainly not in the same sense as an individual cat. What is cat-kind, and what does it mean for the property **furry** to hold of it?

The kind theory given in Carlson (1977a) was deliberately and explicitly silent on these questions, arguing that it is not the responsibility of the semanticist to answer them. Cats are furry just in case cat-kind is furry in some sense or other; and saying what it takes for cat-kind to be furry is the kind of question that falls under the purview of epistemology or metaphysics, rather than natural language semantics. Carlson argues that these questions are of no concern to semanticists essentially because, on his view, they are questions about how to evaluate the truth of generic sentences—not about their truth conditions. To illustrate the distinction, he suggests an analogy with red apples. All the truth-conditional semanticist needs to say about the meaning of 'That apple is red' is that it is true just in case the apple denoted by the NP in subject position is a member of the set of red things. There are any number of further questions one might raise about what it takes for an object in three-dimensional space to count as red, of course. In the case of an apple, it would seem that only its skin has to be red. In the case of a car, only part of the surface (excluding windows, headlights, license plate, etc.) needs to be painted red. In the case of plasticine, the substance needs to be red all the way through. Etc. With a good supply of time and imagination, it is possible to devise variations on these examples indefinitely.[5] Given that they fall outside of the subject matter of semantics, they are nothing for the semanticist to worry about.

A more recent defense of the simple kind theory (Liebesman, 2011) takes this position to a further extreme, arguing that explaining what it takes for cat-kind to be furry not only falls outside the semanticist's job description—it falls outside of everyone's job description, because it cannot be done in the first place. Rather than drawing an analogy to color adjectives, this view proceeds by raising a more general problem. Liebesman thinks that the problem of specifying how kinds inherit properties from their members and the problem of specifying how complex wholes inherit properties from their parts are difficult for the same reason. For a table to be wooden, most of its parts need to be made of wood. But for the table to count as touching the wall, only a small portion of its surface needs, in fact, to be touching the wall. Imagine that for any property that can be named, there is a checklist indicating, for each of the kitchen table's parts, whether it has that property. And imagine that for any imaginable property, there also exists a second checklist indicating, for every

[5] This presentation of the argument is slightly different from Carlson's, which frames the discussion in terms of habituals. But the point is essentially the same. See Carlson (1982), pp. 159-161 and Carlson (1977a), pp. 64-65.

cat, whether it has that property. Liebesman's view is that just as there is no general rule by which to proceed from the first checklist to a conclusion about whether the table possesses or lacks a given property, there is no general rule by which to proceed from the second checklist to a conclusion about whether cat-kind possesses or lacks a given property.

These two arguments deserve separate attention. First, consider Carlson's position. For a long time, semanticists who were confronted with problematic linguistic phenomena would try to deal with them by farming them off to another branch of linguistics or philosophy. Pragmatics was traditionally the unhappy subfield to which all the skeletons in the semanticist's closet were remanded, which is how it came to be nicknamed the wastebasket. However, it is a mark of the great progress that has taken place over the past thirty years or so that pragmatics is no longer looked upon as sort of gulag for intractable phenomena. It has evolved into a field in its own right, with all the sophistication of syntax, semantics, philosophical logic, and the rest. Nonetheless, whenever it is claimed that it isn't the job of semantics to offer some component of an explanation, it is important for the claimant to make a gesture of good faith. There needs to be a principled explanation for why certain sorts of data are best accounted for within another component of a theory; otherwise, it is difficult to avoid giving the impression of shying away from a good challenge. It is in recognition of this danger that Carlson he deploys the clever argument just described to demonstrate that he is not merely throwing the problem of inheritance into the wastebasket.

However, while Carlson's gesture of good faith is to be commended, there are limits to how satisfying it can be. The problem is that it leans heavily on implicit intuitions about where semantics ends and something else—he calls it epistemology, but it could just as well be metaphysics, pragmatics, or the theory of general intelligence—begins. But there simply is no consensus right now on where to draw the border between semantics and pragmatics. Though there are a number of competing positions on that topic, philosophers and linguists are still in the process of determining exactly how various components of a theory of the human linguistic faculty ought to demarcate themselves from one another.

This text will take the view that anyone who wants to draw on the semantics/pragmatics distinction in an argument about how to deal with a particular phenomenon owes her audience an account of the distinction that is both reasonably worked out and independent of the phenomenon under consideration. More specifically, any claim of the form 'phenomenon X should not be treated via a semantic theory, but rather via a pragmatic/epistemological/metaphysical/general intelligence theory' is of optimal use for natural language semantics only when it is supplemented with additional metaphilosophical backstory about:

- the exact data that need to be explained

- given any general informal explanation X–assuming X captures the empirical data perfectly–what criteria determine which component of a theory X belongs in

- whether the worry about trying to capture a set of data in the wrong component of a theory is that we might fail to completely capture the data in the first place, or whether it is that we might capture the data correctly but imperspicuously

Carlson's work gives no indication as to which exact version of the semantics/pragmatics distinction it has in view. Instead, it provides an example of a question—how speakers of a language assign extensions to color predicates—with the expectation that everyone will share the same intuitions about whether it lies within the demesne of semantics. But if anything is clear from the recent literature on the semantics/pragmatics distinction, it is that few authors share any of the same intuitions about where or how to draw the line. Some think the empirical evidence speaks in favor of a separate cognitive module for pragmatics (Borg, 2012), and some think the empirical evidence speaks against it (Prinz, 2006). It could very well turn out that information about what it takes for something to count as red across different contexts *is* part of what should be included in the lexical entry for *red*.[6] If, at a certain point, an elegant semantic theory that were to make exactly the right predictions about what is red when, should it be dispensed with merely because, for whatever reason, it is the wrong kind of theory? Such a move would surely be perverse.

This is not to say that there are no reason for thinking that semantics should where Carlson says it does—that for any given predicate, the only fact a truth-conditional semantic theory ought to keep track of is whether the object under discussion is a member of the set denoted by the predicate. Borg (2004), for example, cites some empirical findings which suggest that the component of our cognitive faculty that proceeds compositionally from syntactic structures to truth conditions is distinct from the component of our cognitive faculty that decides which particular objects fall into which categories. Most suggestively, it seems to be possible to impair one faculty without the other.[7] But it is worth bearing in mind that research on this topic is still in its formative phase; and questions about how, if at all, the human cognitive faculty is partitioned into modules remain open.

Most generally, the further toward the 'lexical' end of the methodological continuum semanticists find themselves, the more categorization-related machinery they are inclined

[6] Already in Kennedy & McNally (2010) we can see a move in this direction: in their analysis, the semantics recognizes a distinction between gradable color adjectives that track how closely an object's color resembles a prototype and gradable color adjectives that track what proportion of an object's visible regions have the relevant color. Hansen (2011) goes even further, adding additional parameters for frame of reference, observation conditions, and observer. These are just two examples of work in compositional, truth-conditional semantics that don't seem to share Carlson's assumptions about how semantics is intrinsically demarcated.

[7] See Gopnik & Crago (1991); Langdon *et al.* (2002); Caplan (1996); Rossen *et al.* (1996).

to build into their semantic frameworks.[8] And the more moved they are by experiments of the sort just discussed, the less categorization machinery they are typically inclined to build in. This debate is not without its interest. Nonetheless, in the midst of interesting foundational discussions it is easy to lose sight of the semanticist's primary task, which is to correctly model the data. All other considerations, as fascinating as they are, need ultimately to take a back seat. Otherwise, these discussions run the risk of being overly speculative.

This chapter takes the position that it is most interesting to decide whether a given phenomenon calls for a semantic or pragmatic analysis only *after* an analysis that accounts for the data has been put forth, and not before. The collective talents as semanticists and philosophers would be better served by adopting a policy of either reflecting on the status of already existing proposals or advancing new proposals, rather than trying to weigh in definitively on complicated foundational issues in advance of getting their hands dirty with empirical facts. In that way, the optimal criterion for demarcating semantics from everything else can emerge out of the study of natural language in all its detail.

A reasonable response to the claim that how kinds inherit properties from their members is not a question for semantics is thus: the exact form a detailed account of how kinds inherit properties from their parts can be settled only once there are some options to choose between. One could begin by accounting for the data via a semantic analysis, then eventually decide that even though that approach gets the data correct, other considerations demand that it be ported into the pragmatic component of our theory. Alternatively, one could begin with a pragmatic analysis, and eventually decide that the same insight would be best expressed through the lens of semantics. Either way, the primary task should always be to account for the phenomenon in some way or other.

Regarding Liebesman's argument, there is a natural response along similar lines. Perhaps it should be the business of semantics to explain how complex wholes inherit properties from their parts. Or perhaps it should not. This is a challenging, multifaceted question. The best way to answer it is to look at an actually existing account of how complex wholes inherit properties from their parts, then deciding whether *that* should be located in the semantics or elsewhere. In the absence of such an account, the question spins idly. Is such an account impossible? Impossibility claims of this sort are not at all straightforward to assess. Most serious challenges seem impossible before anyone has attempted to meet them, and many continue seem impossible after many attempts. And yet, every once in a while, there are breakthroughs, and the absurdity of trying to predict future breakthroughs is self-evident. There may be various instinctive feelings about which approaches seem promising

[8] An interesting example of such an approach is Sassoon (2009).

and which lack promise, but it is in the nature of such feelings not to lend themselves to discussion.[9]

The worry that Carlson's original version of *PM* (variously called *Gn*, *G*, and *G'*) might be a black box solution is that it is silent on all the most interesting questions about what generic statements mean. It says nothing about what the difference between a permissible and an impermissible exception to a generic statement might be. It makes no predictions about which inference patterns a generic statement licenses. It offers no explanation as to why differing proportions of *F*s that are *G* serve to confirm the statement that *F*s are *G*. Whether it is an intensional or extensional operator is left open. It is just a function that maps object-level predicates to kind-level predicates. Which kind-level predicates? Not any kind-level predicates that are attested in any natural language, nor any kind-level predicates that have an intuitive conceptual foundation. Rather, by fiat, it maps each object-level predicate to whatever kind-level predicate it would have to in order to yield the pre-theoretically expected truth conditions.

The black box worry means business. If it is founded, then *PM* is an example of what will henceforth be called *magic*. Consider the following magical semantics for the quantifier *every*:

(1) $[\![\textbf{every}]\!]^{M,g} = \lambda f_{et} \,.\, \lambda g_{et} \,.\, g(E(f))$
where E is a function mapping any property f to 'the every-f'

There are several properties that make this lexical entry magical. One is that the key function it introduces, E, has no definition over and above 'the function that maps a property f to whatever object it needs to in order for that object to have the property g just in case every f is g.' Neither an account of what sort of thing an 'every-f' might be, nor of what the property of its being g might intuitively be, is forthcoming. Effectively, E does no more than restate the problem the lexical entry was supposed to solve. Related to this fact is that lexical entry (1) makes no semantic predictions. It says nothing about the monotonicity properties of *every*, makes no predictions about its tendency to contextually restrict its domains, and offers no general framework for situating it theoretically with respect to other quantifiers, like *some*, *no*, or *most*. Philosophers and linguists find the standard semantics for *every* valuable precisely because it does all three of these things. And finally, no one would think it satisfactory to respond to these worries by arguing that these basic questions about the nature of E do not fall within the purview of semantics.

[9] This is not to disparage Liebesman's view, which is of great interest. Although I would caution readers against taking its skeptical conclusion at face value, Liebesman makes many keen observations about genericity along the way to that conclusion. The analogy between part-whole inheritance and member-kind inheritance, in particular, is certainly worth pursuing.

Contrast the standard Kripke semantics for necessity operators (Kripke, 1963a,b), in which both the notions of an accessibility relation and of possible worlds do crucial theoretical work, but are not magical:

(2) $[\![\textbf{necessarily}]\!]^{M,g,w} = \lambda\phi_{s,t} \,.\, \forall w'(R(w)(w') \rightarrow \phi(w'))$

Within philosophy, there is a grand tradition of debating the intuitive plausibility of possible worlds.[10] This at least shows that many philosophers find the notion of a point in logical space intuitive: it is whatever comes to mind when we hypothesize or fantasize. The accessibility relation also captures something intuitive: namely, that there is a relation between different flavors of natural language modality and some sort of background, be it what we know in the epistemic case, the rules governing a society in the deontic case, or the desires of a person in the bouletic case. This is one prima facie indication that R and w are not playing a magical role in the above definition.

But it should be noted that having an intuitive philosophical interpretation is neither necessary nor sufficient for a set-theoretic object to serve a non-magical explanatory purpose. What ultimately makes the use of accessibility and possible worlds in Kripke's semantics non-magical is that they make further semantic predictions, and that they allow logicians to understand the relation between different flavors of modality at an illuminating level of generality. It becomes possible to prove facts about the expressive capacities of formal systems resulting from different constraints on the accessibility relation. This leads to a deeper understanding of axioms, insofar as it allows logicians to discover maps between different axioms and different properties of frames (van Benthem, 1984). It paves the way for a new algebraic understanding of duality in logic (Conradie *et al.* , 2014). And much more.

This is not to say that being able to provide an intuitive philosophical interpretation for a formal operator is irrelevant to whether it is magical. The correct conclusion to draw from the fact that points of evaluation and accessibility were initially intuitively interpretable, respectively, as possible circumstances and modal flavor, is that this initial interpretation made it plausible that there might in fact be mathematical objects with the relevant properties—not that these are what points of evaluation and accessibility relations 'actually are.' Rather, initially understanding them as modal logicians originally did is precisely what opened the door to the new philosophical interpretations we have today, such as the interpretation of points as states of a computer and of accessibility as courses of action in a computer program (Pratt, 1979; Harel, 1984). In the future, the relational structures on which modal logic has now been shown to offer a local perspective (van Benthem, 2010)

[10] See Lewis (1973); Kripke (1980); Stalnaker (1976); Armstrong (1989).

will no doubt be discovered to be helpfully philosophically interpreted as models of other interesting phenomena as well.

These are some of the reasons for which the black box worry is to be taken seriously. But although the original version of the sophisticated kind theory was vulnerable to the black box worry, the purpose of this chapter is to show that in spite of appearances, there is indeed independent philosophical motivation for the sophisticated kind theory, and that with just a small extension to the sophisticated kind theory, the black box worry is easily addressed. That extension will consist of an intensional semantics for *PM*, which it will be argued has a natural and intuitive philosophical interpretation. In other words, it will be argued that in spite of its initial appearance as magical, *PM* represents a step forward. An argument that shows *PM* not to be magical will satisfy the following desiderata:

- lay the groundwork for making predictions about the inference patterns a generic statement licenses
- explain why certain exceptions to a generic statement are permissible, and what the difference between a permissible and an impermissible exception is
- take a stand on what sort of property an object-level property that has been lifted to a kind-level property is

The next section begins by outlining an intuitive metaphysical picture. After motivating that metaphysical picture on philosophical grounds, it then presents a lexical entry for *PM* for which that intuitive metaphysical picture serves as a viable philosophical interpretation. Subsequently, we turn to objections.

3.2 Ontological Preliminaries

Before considering various conceptions of what kinds might be, a word about how to interpret these remarks is in order. The problem of universals, some 2500 years of age, is not about to be definitively put to rest in a dissertation chapter, let alone in a few paragraphs. But every theory must begin with an of abductive process whereby the various possibilities in logical space are briefly vetted and at least provisionally ruled out.[11] The goal of this section is not to definitively persuade advocates of any particular view in metaphysics. As discussed in Chapter 1, it is not the purpose of this text to put forth any argument about what kinds in fact are. Rather, this chapter will begin with an intuitive exploration of what kinds might be thought to be, using the semantics of generic statements as a rough guide. If

[11] The term *abductive* should be understood in the sense of Peirce (1940), as the procedure whereby, due to time considerations, one selects which of all logically possible hypotheses are the most plausible, and tests them first. See Peirce (1940), pg. 151.

kinds were the sorts of things referred to by the NP in subject position of a generic sentence, what would they have to be?

There are any number of philosophical views one might have about what a kind is. For ease of exposition, we will divide them into four broad categories: nominalism, platonism, epistemicism, and a view that will be referred to hereafter as particularism. Nominalism about kinds holds that all talk of kinds is reducible to talk of individuals. Platonism about kinds has it that there are kinds, that there are objective, observer-independent facts about them, and that they are universals. Assume *universal* to mean: an entity that exists outside of the spatiotemporal world, that is the way it is necessarily, and that is the cause of resemblance among all objects that instantiate it.[12] For the purposes of this discussion, the distinction between the position that kinds are Aristotelian universals and platonism will not matter.[13] Aristotelian realism about kinds dispenses with the assumption that they exist outside of the spatiotemporal world, but otherwise is similar to Platonist realism. The important feature they have in common is that they take kinds to be changeless.

Epistemicism, broadly construed, understands a kind as a person's set of expectations about members of a particular category.[14] On this view, there isn't exactly any one such thing as cat-kind. There are a person's expectations about what she is entitled to presume about a given creature, given an initial state of knowing nothing about it beyond the fact that it is a cat. There are general expectations about cats. Stating that cats are furry is making the move of endorsing certain defeasible patterns of inference about cats in a conversational game. According to epistemicism, kinds are clusters of expectations.

Since nominalism purports to show that all talk of kinds could be dispensed with, it is really a variety of skepticism about kinds. As skepticism about kinds was addressed in Chapter 1, no further response is necessary here.

As regards platonism, a platonist analysis would understand statement (3-a) to be roughly synonymous with statement (3-b):

(3) a. Cats are furry.
    b. The ideal cat is furry.

The NP 'the ideal cat' can be taken literally, at face value, rather than as a manner of speaking. One might then explain why (3-a) tolerates certain exceptions on the grounds that cats

[12] This is the approach pursued in Thomspon (2004; 2009).

[13] Of course, in many other settings, it matters a great deal. It is perhaps worth noting that many contemporary scholars advocate an interpretation of Plato according to which not even he thought that forms were universals in any sense of the term (Fine, 1978; Annas, 1981; Gonzalez, 1996). Readers who favor that interpretation of Plato may take *platonist* to mean 'in the spirit of the traditional reading of Plato.'

[14] This is the approach pursued in Veltman (1996) with respect to normality statements, a close cousin of generic statements. Epistemicism is the background behind all theories that use update semantics, be it for generics, normality statements, conditionals, or imperatives.

which lack fur are just imperfect copies of the ideal cat. They may technically count as cats, but they are defective cats.

Platonism faces a serious prima facie challenge, which is shared by any view that takes kinds to be universals. This is that kinds (in the sense relevant to generic sentences) can change. Recall the assumption from above that if something is a universal, then it is the way it is necessarily. If generic sentences are about kinds, and they can differ in truth value at different moments of evaluation, then kinds cannot be universals. Although it is sometimes assumed that generic statements are somehow semantically atemporal or tenseless,[15] that assumption would seem to be contradicted by (4-a), which was false 200 years ago but is true now, and (4-b), which is perhaps a bit less true now than 200 years ago:

(4) a. Falcons live in skyscrapers.
b. Falcons live in the mountains.

A platonist might reply by saying that statements (5-a) and (5-b) are actually false, and that statement (5-c) would be the closest true statement:

(5) a. The ideal falcon lives in a skyscraper.
b. The ideal falcon lives in the mountains.
c. The ideal falcon lives at a great height.

But if that is the case, the platonist must withdraw their commitment to (4-a) and (4-b) being about universals. Sentence (5-c) is clearly not synonymous with either. The problem breaks down in the following way. Perhaps living in a skyscraper is a deep fact about the life of falcons, in which case the ideal falcon does live in a skyscraper now but didn't 200 years ago. Or perhaps living in a skyscraper is too specific a description of the falcon's life to represent any deep truths about it, in which case whatever it is that sentences (4-a) and (4-b) are about must be something other than the ideal falcon. Either way, platonism is under some initial pressure as a starting point for fleshing out a semantics for generic sentences.

Epistemicism and platonism face the same general problem that all normality theories face, which is that generic sentences sometimes ascribe incompatible properties to the same kind. Recall the following pair of gendered generic sentences from Chapter 2:

(6) a. Ducks have brightly-colored feathers.
b. Ducks are oviparious.

Having brightly-colored feathers is biologically incompatible with oviparity; the two traits are, rare gynandromorphic cases aside, in complementary distribution. If duck-kind is an ideal, how can that ideal bear two incompatible properties? It would be counterintuitive to

[15] See, for instance, Thompson (2004), pp. 3-4.

picture the ideal duck as gynandromorphic. Likewise, if duck-kind is a set of some speaker's expectations, what reasonable speaker would expect a duck, knowing nothing further about it, to have both bright feathers and egg-laying capacities? No sensible disposition to infer would have a speaker draw contradictory or otherwise improbable inferences about any particular animal.

Lastly, there is particularism, a view that is in part inspired by the account of natural kind terms in Kripke (1980), lecture 3. The driving idea behind that lecture is that natural kind terms, though predicates, are in certain respects namelike. Specifically, they are rigid designators: they designate the same object in all circumstances of evaluation. But if they are rigid designators, then they are in some sense referential expressions. This chapter will pursue the particularist route. Like platonism, particularism is committed to the view that there are objective, observer-independent facts about kinds. But unlike platonism, it steers away from understanding them as changeless abstractions, either of the sort that inhabit the third realm or of the sort that inhere in their members. The assumption that kinds are particulars lays the groundwork for neatly explaining how they can bear what seem to be incompatible properties. We will return to that explanation in section 3.4.

Assuming that kinds are particulars also makes it simple to allow that generic sentences are tensed. If kinds are particulars, they have no reason to be debarred from possible change. But the most compelling reason to suppose that kinds are particulars is methodological. Any story on which kinds are particulars can with minimal effort be repackaged into an epistemicist account, should other considerations push in that direction. An epistemicist can hold that uttering a generic sentence is an indication that the speaker will behave as though the generic sentence were true, according to the truth conditions yielded by a particularist analysis. It is straightforward to take an account that departs from particularist intuitions and, as it were, put it into the minds of a speaker and an addressee. It is considerably more difficult to adapt an account that departs from epistemicist intuitions, extruding it out of a person's mind into a particularist analysis.

Those are the basic ontological assumptions that underlie the intuitive metaphysical picture from which we will depart. Having begun with the assumption kinds are particulars, the next step is to say what sort of particular they are.

3.3 Kinds as Production Processes

The core concept behind the intuitive metaphysical picture will be that of a *production process*.[16] The first steps will be to give a basic sketch of what a production process is, then to

[16] No confusion with the term *process* from the literature on tense and aspect (Vendler (1957), Mourelatos (1978), Bach (1980), Rothstein (2008)) intended. A *process* (sometimes also called an *activity*), in that context,

specify some of the interesting features that such things possess. Eventually, the goal will be to argue that these are the very features that a kind semantics for generic sentences ought to reflect. The final step in this initial sketch will be to articulate a conception of kindhood on which kinds are production processes.

The Basic Case: Manufacturing Processes

At the most general level, a production process may be thought of as something that makes things. To illustrate the idea, picture the site of a literal manufacturing process, such as a glass bottle factory. Imagine that it is entirely automated, as many auto manufacturing plants now are. At this factory, new glass bottles are created and released into the world at regular intervals. Although it will make no difference to the eventual formalism, it is philosophically helpful to draw a distinction between the factory itself and the process of bottle manufacture taking place there. A good analogy for that distinction is the distinction between agents and actions—between a person who does things and the things she does.[17] Though it is uncontroversial to recognize a distinction between agents and actions, it so happens that mentally picturing a person's actions in our minds typically involves picturing a person. The same holds for manufacturing processes: even though it makes sense to draw a distinction between a glass bottle factory and the process of manufacturing bottles itself, mentally visualizing the process typically means mentally visualizing a factory. This is unproblematic, as long as visualizing the factory does not involve mistaking the factory for the process.

A semantic theory expresses what it takes a given phenomenon to be by setting the parameters according to which that phenomenon can vary. Any semantic theory that seeks to model manufacturing processes, then, needs to indicate what it takes manufacturing processes to be by answering the following sorts of questions. In what relations can manufacturing processes stand to one another? In what relations can they stand to other things? What properties are they capable of bearing or failing to bear? What constitutes a boundary in the logical space of manufacturing processes?

There are any number of properties that a manufacturing process might be thought to possess, but one particular class of such properties is worth singling out for its importance to this line of inquiry. If glass bottles are what Bottles Inc. manufactures, then autos are probably not what it manufactures. So Bottles Inc. is a process that manufactures bottles

---

is a temporal structure with certain formal properties, such as the subinterval property. The semantics proposed in this chapter will be atemporal, so that tensed generic sentences may be handled using the standard framework for tense. It can easily be temporalized, should that prove necessary.

[17] Though of course, the analogy only goes so far. It is odd to speak of a person's actions taking place 'at that person,' for instance.

rather than a process that manufactures cars. Being a process that manufactures cars is another way Bottles Inc. might have been, as opposed to the way it actually is. Let the name *product property* be a term for these complicated properties that regard the product of a production process. Examples of product properties will include: **process that manufactures bottles**, **process that manufactures cars**, or **process that manufactures computer chips**. Properties that ordinary individual objects can bear, such as **bottle**, **car**, or **computer chip**, will be referred to as *individual properties*.

The crucial feature of product properties is that they are systematically derived from individual properties. For instance, contrast being a transparent bottle with being a process that makes transparent bottles. Those are obviously not the same property. Being transparent is an individual property, applicable only to individual bottles. But it makes no sense to say that a process that makes transparent bottles is transparent (except perhaps in the coerced political sense of the term). Nonetheless, these properties are clearly related to one another, as can be discerned from the fact that all of the above examples involved appending the words 'process that manufactures' to the expression for an individual property. We may call the relation that the property **transparent** bears to such a manufacturing process—whatever that relation may be—the *characterizing* relation. So saying that this hypothetical manufacturing process is *characterized* by the property **transparent** will be shorthand for saying that it is a process which makes transparent things. In order to be able to talk about the relation between individual properties like **transparent** and product properties like **process that makes transparent things**, it may be stipulated that for every product property G that holds of a process p, there is an individual property F such that F characterizes p.

Another noteworthy characteristic of product properties is that they exhibit far fewer incompatibilities than their individual counterparts. Being a bottle is incompatible with being a car, but being a process that makes bottles is not incompatible with being a process that makes cars—you might imagine a factory where both are produced. Or, in case that has an improbable ring to it, imagine a factory that makes both blue and red bottles. Perhaps there are two reservoirs of molten glass, and the mechanism alternates which reservoir it draws upon when injecting the mold with the material for each new bottle, with the end result that half the bottles which emerge from it red, and half the bottles which emerge from it are blue. In that case, it is correct to say both that the factory has the property of producing red bottles and that it has the property of producing blue bottles, even though no individual bottle can be both blue and red. This makes intuitive sense, given that at most actual factories, a variety of products are manufactured.

The notion of a product property captures the following set of intuitions. A bottle manufacturing process can be one that makes transparent things even if every once in a while, an accident intervenes to prevent the bottle it happens to be making from coming out transparent. An impurity could sneak into the glass, someone could sabotage the apparatus,

or a forklift could accidentally crash into the factory, shutting down all operations. There is no guarantee that fate will hold off from intervening to prevent the occasional bottle from coming out transparent. This is precisely what makes product properties so fascinating: accidents can befall the things a process manufactures without affecting whether it bears the relevant product property.

Here is a first attempt to describe everything said so far slightly more precisely, though ultimately still informally. Let F be any individual property, p be any production process, and C be the characterizing relation—so that, for example, $C(F)$ denotes the product property that holds of p just in case it is characterized by the property F. Then:

(7) $C(F)(p)$ iff p is a process which makes things that are F

As definitions go, (7) is not particularly instructive. In fact, it is tautologous, given that the characterizing relation was defined above as the relation that an individual property F holds to a production process p just in case p is a process which makes things that are F. However, it may nonetheless serve as a starting point or blueprint for what will eventually be the definition of the characterizing relation. Once a small number of further primitive relations that can obtain between production processes have been specified, the goal will be to replace the English right-hand side of the biconditional in (7) with something in a formal metalanguage, thus making the definition something that can be used for compositional semantics. Specifically, the final version of the definition will state what it is to be a process which makes things that are F. The present purposes only require the informal idea that there are different ways a manufacturing process can be: it can be the sort of process that makes red things, the sort of process that makes blue things, and so on.

The recommended semantics for manufacturing processes will posit three relations in which one process can stand to something else. First, there is the relation in which a production process stands to the individual objects it produced. This may be called the *progenitor relation*. A computer chip manufacturing process stands in the progenitor relation to every chip it produced. Likewise for the auto manufacturing process at the local plant and the cars it produced. Think of the progenitor relation as something roughly like what is captured when, for instance, a soda bottle lists the address of the factory where it was made. Everything manufactured by a production process has a certain origin, and having that origin is a matter of standing in a certain relation to the process that created it.

It is helpful to think of the progenitor relation as a casual relation, though it must be borne in mind that being the progenitor of a product is a way of being an *efficient* cause specifically. The idea may be reformulated in the following alternative way: to ask what produced this glass bottle amounts to asking what is responsible for its existence. The kind of responsibility operative in the ancient notion of cause or aition (which originally meant 'responsible party') is exactly the kind of responsibility to have in mind when considering

the relationship between a process and the things it makes.[18] Any manufacturing process is the efficient cause of the products it manufactures.

The progenitor relation must at least be irreflexive and antisymmetric. It is highly unclear what it might mean for a manufacturing process to manufacture itself, and it is equally difficult to make sense of e.g. a bottle manufacturing the process that manufactured it. It will also prove useful to assume that the domain and range of the progenitor relation are disjoint. The reasons for this are fairly clear: no manufacturing process manufactures further manufacturing processes; nothing can be both a manufacturing process and the product thereof.

Secondly, there is a certain relation of dependency that can obtain among manufacturing processes. However, it will simplify the formulation of future definitions to work with the converse of said dependency relation, rather than the dependency relation itself. We may refer to this converse relation as the *feeding* relation. A process p feeds another process r when, in order for r to create something, p must also create something—when r cannot go into production unless p also goes into production. As an example, consider the process of manufacturing desks at one factory and the process of manufacturing screws at another, and suppose that all the desks under consideration are assembled using these screws. Then the screw manufacturing feeds the desk manufacturing process in the relevant sense: no desks can be manufactured unless screws are also manufactured.

It should be noted that the feeding relation makes some additional nuance in the previous characterization of manufacturing processes as efficient causes necessary. A manufacturing process is responsible for the existence of the things it creates. It is natural, in that case, to conclude from that the fact that process p feeds process r means that the items produced by p would not exist, if not for process r. Does it follow, then, that process r is responsible for the existence of what process p created? It does, in a certain sense. If the only reason some kind of screw is in production is so that it can go into the making of some kind of desk, then it could be said that the desk manufacturing process is derivatively responsible for the existence of some screws. However, it would not be responsible in the sense of being their progenitor. This stipulation is meant to capture the intuition that the purpose of a desk manufacturing process is to make desks, not to make screws. Screws are, at most, a byproduct. This will be important to keep in mind vis-à-vis examples of manufacturing processes which feed one another.

Requiring the feeding relation to be antisymmetric will have the effect of ruling out the possibility of two processes that feed one another. Most pairs of manufacturing processes do not stand in this relation—if they did, they would involve a perpetual exchange of

[18] See *Physics*, Bk. 2.

products between two factories. Nonetheless, it is somewhat unclear whether the linguistic data require our semantics to take a stand one way or the other on the matter. Therefore, although the analysis under consideration is consistent with the assumption that the feeding relation is antisymmetric, the version presented here will not make it. The feeding relation must be reflexive, for the trivial reason every process must make something in order to make something. And finally, the feeding relation also needs to be transitive: if it is impossible to make zippers without making zipper teeth and impossible to make luggage without making zippers, then it is also impossible to make luggage without making zipper teeth.

Finally, there is a certain parthood relation that can obtain among manufacturing processes. We may say that a manufacturing process p is part of a manufacturing process r just in case whenever p creates something, process r has thereby also created that same thing. As an example, consider the large-scale process of Ford automobile production that is currently taking place across the US, as compared with the smaller process of Ford automobile manufacture taking place at the southeastern edge of Chicago. The Chicago Ford production process is a part of the Ford production process, because whenever it makes a new car, the Ford production process has thereby also made a new car. And not just any car—it has thereby made the very same car. Since the parthood relation is best thought of as a kind of inclusion relation (one process is a part of another if and only if all of its progeny are also the other's progeny), it will behave exactly in the way that inclusion relations are known to behave, which is to say that it will be a partial order.

Modal Semantics for Manufacturing Processes

The next step is to give the definition of the characterizing relation in (7) more substance. This will involve answering two related questions:

- What is it to be e.g. a process that makes red things?
- What does the property of being such a process have to do with the property of being red?

Our basic approach to these questions is to give manufacturing processes a teleological read. The semantics will take as a primitive notion the idea that for each process, there is a way it is supposed to go. An automobile manufacturing process is supposed to succeed in making automobiles. For the process, that is a desirable outcome. If it manages to produce automobiles, then things are going well as far as that process is concerned: it is successfully performing its function. If the process fails to produce automobiles, then things are not going well for it. Talk of 'things going well' for a process might sound somewhat spooky to modern scientific ears, as though it were anthropormorphizing what takes place at a factory. But there is really nothing spooky about it. None of this talk is intended to suggest that mechanical processes somehow have desires, or that they are 'trying to' succeed in

manufacturing their products. It is simply to say that we evaluate them, in some manner or another, in the light of certain outcomes. Formally, the kind of modality in question is very much akin to bouletic modality. Manufacturing processes are a kind of activity, and as an activity, they are geared toward particular outcome. Whether or not any particular agent or intention underlies the activity is, from a formal point of view, irrelevant.

The association of manufacturing processes with ideal outcomes can take the form of a simple Kratzerian semantics for modals.[19] Define a teleological accessibility function h which, given a world of evaluation, associates every manufacturing process with a set of ideal outcomes at that world. The ideal outcomes, as per the tradition in possible worlds semantics, are represented as sets of possible worlds. So informally, h will be the following function from worlds and processes to sets of worlds:

(8) $h(w)(p) = \{w' \mid p$ does what it is supposed to in w without interruption$\}$

For example, at the actual world, h will map the process of manufacturing Volkswagens to all those worlds at which it successfully produces vehicles of the intended sort without interruption. In general, h determines what is supposed to happen whenever a process is allowed to proceed to completion. In this regard, it is parallel to a bouletic conversational background, which maps a world of evaluation and an agent to the set of worlds at which that agent's desires are realized.

With h and the progenitor relation (represented below as P) in place, a more interesting informal semantics for the characterization relation is now available:

(9) $C(F)(p)$ iff $\forall w' \in h(w)(p)$: $w' \vDash \exists x(P(p)(x) \wedge F(x))$

A property F characterizes a manufacturing process p just in case at all of p's ideal worlds, some of p's progeny are F. So the property **transparent** characterizes Bottles Inc. (i.e. Bottles Inc. is a process that makes transparent things) just in case at all the worlds where Bottles Inc. does what is supposed to without interruption, it has succeeded in producing some transparent things. Why *some*? One might expect the truth conditions to state that the ideal outcomes for Bottles Inc. involve its producing *nothing but* transparent things. But certain factors weigh in favor of an existential quantification. Note, as previously observed, that product properties admit far fewer incompatibilities than their corresponding individual properties. A manufacturing process can have more than one kind of product. Just because red bottles are produced at some factory, there is nothing preventing blue bottles from being produced at the same factory. So it should not follow from the fact that Bottles Inc. has the property of making blue things that it cannot also have the property of making red things.

[19] See Kratzer (1977). The more sophisticated Kratzerian apparatus, which makes a distinction between a modal base and an ordering source, is not necessary here.

To say that it is a process which makes blue things is only to say that blue things are one of the things it makes, potentially among several.

Extending the Account to Kinds

The next step is to argue that the logical structure we have been attributing to manufacturing processes is precisely the logical structure that kinds have. The term *production process* will, hereafter, be used to indicate anything with this logical structure, regardless of whether it is literally a manufacturing process. Kinds are typically not manufacturing processes of the sort which take place at a factory. A production process, in this new special sense of the term, does not necessarily assemble things out of a set of mechanical parts. Nonetheless, what production processes have in common with manufacturing processes is that they exhibit all of the features under discussion up to this point. They create things. As efficient causes, they are are responsible for the existence of the things they create. They can be interrupted in various ways. And for each of them, there is a way it is supposed to proceed. Production processes have a teleological character, in the sense that they are evaluated in the light of certain outcomes. The fact that for every production process, there is a way it is supposed to go allows it to have certain derived properties—to be characterized by properties of individuals. Production processes can stand to one another in the relation of (converse) dependence called *feeding*: one process may depend on another for going into production. And finally, they are organized into hierarchies of parts and wholes: there are small-scale production processes related to large-scale processes in such a way that whenever the former produces something, the latter has thereby produced the very same thing.

In case the more general idea of a production process seems difficult to visualize, the following examples ought to help stimulate the flow of intuitions. Consider cat-kind. On the picture being developed, cat-kind is the process that makes cats. The idea might initially sound odd, as though it sought to represent cat-kind as some sort of manufacturing process. But rather than thinking of the cat production process as an Industrial Revolution-era factory, as it were, gluing cat bits together, it is far more helpful to think of cat-kind as an evolutionary process. There is a stable system of alimentary, respiratory, and reproductive transactions by which the cat population is presently sustaining itself, subject to selection pressures and the rest. *That* is what created all the cats that are presently alive. Every cat on the planet currently exists thanks to the process of cat production.

Thinking of kinds along these lines leads directly to a satisfying explanation for why generic statements, even though they superficially sound like strict generalizations, are able to have exceptions without being thereby falsified. Recall that manufacturing processes can be interrupted: Bottles Inc. can be one which makes things that are leak-proof, even if some mistakes make it through every once in a while. The fact that the occasional bottle

comes out cracked due to an unforeseen accident is entirely compatible with Bottles Inc. bearing the property of making things that don't leak. A similar story holds for cats: the cat production process is one that produces furry things. When left to its own devices, that process creates nothing but furry cats. However, the cat production process exists in a world where all sorts of unexpected contingencies interfere with things. So a particular cat might have been shaved, or exposed to radiation, or engineered not to have fur by way of selective breeding. None of the above scenarios are incompatible with the fact that the evolutionary process responsible for the existence of cats is a process which makes furry animals. They are merely indications that from time to time, that evolutionary process can be interfered with in such a way as to yield unexpected outputs.

A further advantage to such a conception of kinds is that it explains the intuition, shared by nearly everyone who writes on generic statements, that it is in principle impossible to enumerate their exception categories in advance. For instance, the prevailing feeling in the literature on generics is that there is no way to translate a statement of the form in (10-a) into a statement of the form in (10-b):

(10) a. Cats have fur.

b. Every cat who:
has not been shaved ∧
has never been exposed to radiation ∧
is not a sphinx ∧
has no genetic anomalies ∧ ... has fur.

The impossibility of such a translation is nicely explained if we assume that there is no way to predict all the forms that interruption might take in advance. Fate is endlessly creative. This assumption may or may not be justified—it all turns on whether determinism is true, and an argument for or against determinism is beyond the scope of this work. The point is only that taking the exception categories for a generic statement to be unspecifiable in advance goes along with taking every conceivable accident that might keep a production process from coming to completion to be similarly unspecifiable. To view kinds as production processes is to see a strong connection between these two doctrines.

One final benefit to the processual perspective on kinds is that it has the potential to be as broad as generic sentences would seem to require. For generic statements are not only about biological kinds, like cats or birds. As observed in Chapter 1, there is no difficulty whatsoever in making generic statements about artifacts, people, careers, social conventions, political institutions, or even philosophy. From a certain vantage point, it is possible to view all of these putatively unnatural kinds as production processes in their own right. Consider the following generic sentence:

(11) Brides wear white.

At first, it seems like a stretch to imagine that there is a process that makes brides. But upon further examination, the idea is not without its plausibility. There is a social-institutional process responsible for creating marriages—the activity in which the social institution of marriage engages, whatever that may be—and it has multiple wings: the wing dedicated to the production of brides, the wing dedicated to the production of grooms, the wing dedicated to the production of bachelorette parties, the wing dedicated to the production of elaborate wedding cakes, and so forth. The relation between the marriage production process and its various 'wings' is best understood as the relation of feeding, in the sense described above. The bride production process feeds the wedding production process, in the sense that it is impossible to create a (traditional) wedding without creating a bride.

One might object that this leads to the bizarre consequence that brides are weddings, on the grounds that for every bride, a wedding production process is responsible for its existence. However, this bizarre consequence does not in fact follow from adopting the processual perspective on social kinds, any more than the consequence that screws are desks follows from the fact that a screw production process feeds a desk production process. Suppose that screws would not exist without the process of making desks which the process responsible for making them feeds;[20]. Even then, the desk production process would not be directly responsible for the existence of screws. It would only directly be responsible for the existence of desks, in the sense that screws are a byproduct of the process—not the main attraction.

To assert (11) is to claim that the bride production process is a process which produces things in white garments. Assuming (11) is true—which is potentially up for debate—exceptions to the social norm regarding brides will be permissible in exactly the same circumstances in which exceptions to the norm regarding cat fur were permissible. There may be brides who opt for a different mode of attire. But if (11) is indeed true, then whenever this happens, it is because something came along and interfered with the process that gives rise to the existence of brides—this particular bride's desire to have a nontraditional wedding, perhaps. So a similar pattern of reasoning applies to biological and social kinds. Evolutionary processes and cultural institutions, though profoundly different in all sorts of important respects, have in common that they create things in a way that is susceptible to interruption.

The most vivid way to draw the formal analogy between manufacturing processes and production processes is to think of schools. There is an activity of producing people with certain social-institutional qualifications, such as having a degree, which is currently

[20] This is an additional relationship of dependency that would be easy to model in the semantics under consideration.

taking place at most schools. Indeed, that is why in popular culture, schools are often metaphorically referred to as factories for producing job candidates. Viewing other social kinds in a similar light only requires the assumption that some social institutions can have a de facto status, rather than being legally accredited bureaucratic entities, and that certain institutionally-conferred qualifications can have a comparable de facto status, instead of being legally recognized credentials. This further assumption is made all the time in day to day reasoning.

We turn next to more challenging cases, such as non-living natural kinds. One such kind makes its appearance in Tarski's celebrated Convention T sentence:

(12) Snow is white.

From the processual perspective, sentence (12) would be stating (roughly) that the process responsible for producing snow is a process which produces things that are white. Folk meteorology provides an excellent candidate for this process: namely, the atmospheric component of the H_2O cycle. A similar story could be told about this sentence:

(13) Water is refreshing.

Here, the thought would be that the water-manufacturing component of the H_2O cycle is the sort of process that makes refreshing things. A similar folk-scientific backstory would apply to other meteorological kinds. The same strategy might even be pursued for generic statements about chemical kinds:

(14) Gold is radioactively stable.

These cases are more challenging for the intuitive metaphysical picture under discussion, because many authors have the intuition that being gold is an accidental property: the accidental property of being comprised of atoms with atomic number 79. But even here, the processual perspective on kinds is a natural fit. On that picture, there indeed is a process responsible for creating gold: in this case, the process of supernova nucleosynthesis that created and scattered the gold that exists all over the galaxy. Exceptions to (14) arise when something extraordinary interferes with that process—as, for example, when physicists synthesize radioactive isotopes of gold.

Now, one might grant that all hitherto observed gold derives from a single process of nucleosynthesis while still being hesitant to claim that the *purpose* of that process was to create gold. Talk of gold having a purpose may initially sound like a regression into Aristotelian science, which is obviously well-known to be empirically false. That is not necessarily a problem for the picture under consideration—after all, as discussed in Chapter 1, it is entirely consistent with the project of natural language metaphysics that the commitments it uncovers are false. Nonetheless, it is not at all clear that the picture under consideration does in fact involve a regression into pre-Galilean science. As long as *purpose* is understood

in a sufficiently weak, non-anthropomorphized sense, the idea that gold has a purpose (in a manner of speaking) is arguably very much a part of our folk conception of nature. One indication that commonsense reasoning does indeed work under the assumption that there is a way gold is supposed to be is that when physicists do synthesize radioactive isotopes of gold, it is not uncommon to describe such activity as 'tampering with nature.'

Such is the general idea behind the notion of kinds as production processes. Many of the things that humans encounter and create were made by processes which, though not manufacturing processes taking place at a factory, have the formal features thereof. The next step will be to apply the semantics just developed for production processes to a compositional analysis of generic sentences.

Modal Semantics for the Predicate Modifier

Thus far, the intuitive metaphysical picture of kinds as production processes has proceeded informally. Before we state the formal version of the analysis, a few broad methodological remarks are in order. As stated in section 2.1, this analysis will adopt the framework for compositional natural language semantics articulated in Heim & Kratzer (1998), which maps syntactic structures (here understood as binary branching trees, generated by a finite set of syntactic operations on lexical items) to truth conditions, on the basis of lexical entries for every terminal node and a small set of composition rules that explain how to derive the denotation of each parent node from the denotation of its child nodes. The only composition rule required for present purposes is function application. However, rather than using English as a metalanguage for stating denotations, this analysis will use an intensional, typed, higher-order lambda calculus, more or less of the sort presented in Gamut (1991) §5.8, but with an additional atomic type p specifically for production processes, with relation symbols P, N, and $\sqsubseteq$ for the progenitor, feeding, and parthood relations, respectively, and an additional symbol Base for the teleological accessibility function.[21]

The syntax and semantics for our metalanguage are as follows:[22]

Syntax:

- **Types** - $\tau ::= e \mid p \mid t \mid s \mid \langle \sigma, \tau \rangle$
  Where this doesn't cause confusion, $\langle \sigma, \tau \rangle$ will be abbreviated as $\sigma\tau$.
- **Constants** - $c ::= a_\tau \mid R_{\sigma\tau} \mid P_{pe} \mid N_{pp} \mid \sqsubseteq_{pp} \mid \textsc{Base}_{\langle s,\langle p,\langle s,t\rangle\rangle\rangle}$
- **Variables** - $v ::= x_\tau \mid f_{\sigma\tau}$

[21] To economize on alphabet letters, the analysis will use p both to refer to the type itself and to variables of that type, in addition to using the letters $q, r, \ldots$ for variables of type p.

[22] Syntax stated in Backus-Naur form.

- **Terms** - $\alpha ::= c \mid v \mid \alpha_{\sigma\tau}(\beta_\sigma) \mid \lambda x_\tau \alpha_\sigma$
- **Formulas** - $\phi ::= R_{\sigma t}(\alpha_\sigma) \mid f_{\sigma t}(\alpha_\sigma) \mid \neg\phi \mid \phi \wedge \psi \mid \phi \rightarrow \psi \mid \forall x_\tau \phi$

Semantics:

- Domains:
  - $\mathcal{D} = \{D_\tau \mid \tau \in \textbf{Types}\}$
  - For all $\alpha_\tau \in$ **Terms**: $[\![\alpha_\tau]\!] \subseteq D_\tau$
  - For all $\phi \in$ **Formulas**: $[\![\phi]\!] \subseteq D_t$
  - For any D_ρ and D_υ, $D_{\rho\upsilon} = D_\upsilon^{D_\rho}$ (the set of all functions from D_ρ into D_υ)
- Models: $M = \langle \mathcal{D}, I \rangle$
- $[\![x_\tau]\!]^{M,g,w} = g(x_\tau)$
- $[\![a_\tau]\!]^{M,g,w} = I(a_\tau)$
- $[\![\alpha_{\sigma\tau}(\beta_\sigma)]\!]^{M,g,w} = [\![\alpha_{\sigma\tau}]\!]^{M,g,w}([\![\beta_\sigma]\!]^{M,g,w})$
- The propositional connectives $\neg$, $\wedge$, and $\rightarrow$ are defined in the usual way.
- for any variable x_τ and formula ϕ, $[\![\forall x_\tau \phi]\!]^{M,g,w} = 1$ iff for all $d \in D_\tau$, $[\![\phi]\!]^{M,g[x_\tau \rightarrow d],w} = 1$
- for any variable x_τ and formula ϕ, $[\![\exists x_\tau \phi]\!]^{M,g,w} = 1$ iff $[\![\neg\forall x_\tau \neg\phi]\!]^{M,g,w} = 1$
- for any variable x_τ and expression α_σ,
  - $[\![\lambda x_\tau \alpha_\sigma]\!]^{M,g,w}$ = the function $h \in D_{\sigma\tau}$ s.t. for all $d \in D_\tau$, $h(d) = [\![\alpha]\!]^{M,g[x_\tau \rightarrow d],w}$

This semantics indicates the bedrock model-theoretic assumptions of the logic employed by the analysis. The logic will be interpreted via many-sorted higher-order models.[23] Some of the ideas discussed informally in previous sections can then be modelled by introducing constraints on these models. Following Malink (2006), these constraints may be referred to as *axioms*. (However, they are not to be mistake for axioms in the sense of sentences that form the basis of a proof system.)

(15) ***General Axioms:***
For all $D_\sigma, D_\tau \in \mathcal{D}$: $D_\sigma \cap D_\tau = \varnothing$

[23] For the reasons indicated in Gamut (1991), §5.8, we leave the specification of these models as indefinitely higher-order. For the fragment of natural language under present consideration, nothing beyond first-order models will be necessary. And nothing beyond second-order models should be required for larger fragments of natural language.

Axiom (15), along with the initial constraints on the domains, suffices to yield the features of the progenitor relation that were sought after in section 3.3. Since all domains are disjoint and the progenitor relation is of type $\langle p, e\rangle$, it will have disjoint domain and range, and consequently be irreflexive and antisymmetric. But the desired constraint on the feeding relation, that it be a preorder, must be specified explicitly:

(16) ***Feeding Axioms:***

a. For all $a \in D_k$: $\langle a, a\rangle \in I(N)$

b. For all $a, b, c \in D_k$: if $\langle a, b\rangle \in I(N)$ and $\langle b, c\rangle \in I(N)$ then $\langle a, c\rangle \in I(N)$

Finally, some axiom is needed to capture the intuitive meaning of the parthood relation. Pursuing the above suggestion that it is a kind of inclusion:

(17) ***Parthood Axioms:***

For all a, $b \in D_k$ and $x \in D_e$: if $\langle x, a\rangle \in I(P)$ and $a, b \in I(\sqsubseteq)$, then $\langle x, b\rangle \in I(P)$

This axiom states that whenever one process feeds another, all progeny of the former are progeny of the latter. A formulation on these lines seems to be the simplest way of formally spelling out the intuition that if one production process is part of another, then whenever it makes something, its parent process thereby also makes that same thing. As mentioned earlier, since this conception of parthood makes it into an inclusion relation, it is a conception on which parthood is a partial order.

We now turn from metalanguage to object language. For the purposes of exposition, the analysis pursued in this chapter will posit an unpronounced predicate modifier node in the syntax, as in (18):

(18)

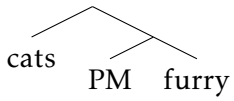

At some point, it will of course be necessary to consider whether there is independent evidence for such a syntactic analysis. If there is not, the predicate modifying operation will need to be implemented in some other way, which is easy enough. Making it into a type-shifting rule, for instance, is formally trivial. For the moment, however, the focus will be on devising an analysis that works. Since the teleological accessibility function was a function from *pairs of* worlds and kinds to sets of worlds (rather than a function from worlds to sets of worlds), the operator Base will be a function of type $\langle s, \langle p, \langle s, t\rangle\rangle\rangle$, and *PM* will denote a function of type $\langle\langle e, \langle s, t\rangle\rangle, \langle p, \langle s, t\rangle\rangle\rangle$. If Base maps the world of evaluation and a kind to the characteristic function of those worlds at which the ideal outcomes for the kind come to pass, then *PM* will denote a function with the following definition:

(19) $[\![\mathbf{PM}]\!]^{w,g} = \lambda f_{est}\ .\ \lambda p\ .\ \ldots$

$\ldots\ \forall w'(\textsc{Base}(w)(p)(w') \rightarrow \exists x(P(w')(x)(p) \wedge f(w')(x)))$

Given that this is the lexical entry that drives the entire analysis, it is worth taking a moment to think it through intuitively. (19) is a Montagovian rendering of an operator meant to produce the truth conditions in (9). *PM* will denote a function from object predicates to kind predicates, mapping an individual predicate f to a new predicate which applies to a kind p just in case at all of p's ideal worlds, some of p's progeny are f. So in the case of (18), it takes the property **furry** to a new property that holds of a kind p just in case at all of p's ideal worlds, some of p's progeny are furry. So sentence (18) will be true at a world w just in case the denotation of *cats* (i.e. cat-kind) has that property in w.

A modal semantics for *PM* along these lines has the potential to yield some promising results. Of particular interest are 'unintended effect' generic sentences:

(20) Clowns are creepy.

For the current point it is helpful to bracket out the relativist worries that this statement raises and proceed as though *creepy* were not a predicate of personal taste.[24] Even with that simplifying abstraction in place, the example is worth dwelling on. Suppose the speaker to be a coulrophobic who utters (20). What she seems to be saying is not that clowns strive to be creepy, but that being creepy is an unintended side effect of clowns comporting themselves in the way clowns do. According to the semantics under consideration, statement (20) is true just in case at all worlds where clown-kind (the set of social-institutional processes responsible for creating people who are competent in the clown arts) does what it is supposed to, some of its progeny are creepy.

At first, these might sound like the wrong truth conditions. Nowhere in the clown charter is it stated that clowns are supposed to be creepy. On the contrary; they are supposed to be funny. But luckily, that is not what the analysis predicts. To see why, consider an analogy to bouletic modality. Imagine a driver who wants to make a U-turn on the highway, but who is under the mistaken impression that the exit for the U-turn is on the right. A passenger then says:

(21) You should be in the left lane.

According to the standard semantics for modals (Kratzer, 1977, 1981), the above claim amounts to saying that the driver's desire coming true entails her being in the left lane. But that is not the same thing as saying that she *desires* to be in the left lane. Arguably, in that situation, the driver does not desire to be in the left lane, except perhaps in some special Socratic sense of the term.[25] Something similar applies to disturbing clowns. A speaker

[24] Incidentally, all approaches to predicates of personal taste, whether they be realist or relativist, are compatible with the semantics for PM given in (19). See Lasersohn (2005).

[25] In the early Platonic dialogues, Socrates' skepticism about akrasia is driven by the belief that it is impossible for any person to want what is bad for her. See *Apology*, *Euthyphro*.

who states that clowns are creepy is not saying that they are supposed to be creepy. Rather, she is saying that their being creepy is a (possibly unintended) consequence of their being the way they are supposed to be. And that is precisely the right prediction about sentences like this. There is a way clown-kind is supposed to manufacture clowns,[26] and if it succeeds in manufacturing clowns in that way, then some of the things it produces will be creepy.

That example ought to provide some sense of what the definition for *PM* given in (19) can do. An important upshot of this definition is that although it attributes a certain teleological character to generic statements, there is a limit to how teleological it makes them out to be. They do not always directly describe how production processes are assessed in the light of different outcomes. Sometimes they describe what follows from those evaluations. This subtlety in the meaning of generic sentences has, to the author's knowledge, so far gone unrecognized in the literature.

3.4 Conjunctive Inference Patterns

How inference patterns relate to semantics is a matter of ongoing debate. Brandom (1994) argues that the meaning of a sentence *is nothing more* than the set of inferences it sanctions, that the meaning of a word is nothing more than the equivalence class consisting of all the substitution inferences in whose premises it figures, and that truth and reference can be defined in terms of the consequence relations that obtain between sentences. The approach advocated there has been given the name *inferentialism*. Williamson (2012) expresses reservations about inferentialism, favoring more traditional *referentialist* semantic frameworks, which define logical consequence in terms of truth and reference (or related concepts, like satisfaction). The core worry is version of problem of logical omniscience: it doesn't seem to be the case that one has to be able to draw all the substitution inferences in which a concept figures in order to count as possessing the concept. On the other hand, it does seem clear that one needs to be competent at drawing *some* of those inferences in order to possess a given concept. The question is: where should the inferentialist draw the line?

This chapter will not take a firm stance in the inferentialist/referentialist debate. Although the inferentialist project is fascinating, it is still too early to know exactly what its philosophical payoff is—especially given that it has not yet been applied to any concrete problems in natural language semantics.[27] But despite their points of discord, inferentialists and referentialists can both agree to a more modest conceit: regardless of whether the

[26] Some approximation of: silly in an outlandish and self-deprecating way, perhaps?

[27] Interestingly, Greg Restall's forthcoming *Meaning in Action* research project, which was just given funding by the Australian Research Council, proposes to do just this. This project will no doubt constitute a major development in the field.

inferences a sentence sanctions are part of its meaning or part of something else, a semantic analysis of that sentence is obliged to make predictions that are at least *consistent* with what we know about the inferences it sanctions.

Quantificational theories are typically thought to have two strong points: 1) they explain how the truth of generic statements supervenes on truths about the individuals being generalized over (or, to rephrase the same idea using kind terminology, how kinds inherit properties from their members), and 2) they make precise predictions about the inferences that generic sentences validate, whether they be classical or default.[28] The criticisms of the simple kind theory from section (12) show that any kind theory worth its salt should do these two things as well. One of the purposes of this text is to show that unlike the simple kind theory, the sophisticated kind theory has all of the logical flexibility of a quantificational theory, but has the added advantage of achieving the distributional coverage of earlier kind theories. As will become evident in section 3.7, most criticisms of the kind theory apply only to simple kind theories, which the analysis here also rejects. A consequence of answering the black box worry, as this analysis does, and providing a lexical entry for PM, is that the kind theory is now able to help itself to all the advantages of a quantificational theory.

Most of the work on generic inference patterns has been in the context of default reasoning, the branch of artificial intelligence that studies reasoning in the absence of complete information.[29] Scholarship in default reasoning researches the logic of 'normally' statements, such as 'Bats normally have wings' or 'Policemen normally have badges'—a class of sentences that intuitively share many features with generic statements. Default reasoning is a rich area of research, and a number of frameworks have been put forth to cover a wide range of default inference patterns. Providing a default logic that uses the analysis put forth in this chapter to validate the standard inventory of default reasoning patterns is an exciting opportunity for future research.

We now revisit an important inference pattern that holds for generic sentences, first raised in section 1.1, which is not only valid by default, but deductively valid as well:

(22) Fs are G.
Fs are H.
$\therefore$ Fs are both G and H.

Here is an instance of the argument pattern in (22):

[28] This helpful framing of the debate is due to Nicholas Asher.

[29] McCarthy (1980) is the foundational text in this area. For some examples of how it can be applied to the study of generic sentences, see Geurts (1988), Morreau (1988), Blutner (1988), Krifka (1988), Veltman (1996), Asher & Pelletier (1997), and Cohen (1997).

(23) Cats are solitary.
Cats are playful.
∴ Cats are solitary and playful.

This argument is intuitively valid. Intuitions about the gendered version of this argument pattern are difficult to pin down, but it arguably at least has an interpretation on which it is valid:

(24) Lions have manes.
Lions are lactiferous.
∴ Lions have manes and are lactiferous.

Since it is examples like this that pose problems for normality theories, and quantificational theories tend to be normality theories, it should not be surprising that they post problems for many quantificational theories. Suppose we adopt the simple normality theory given in section 2.2. The simple normality theory incorrectly predicts the premises and conclusion in (24) to be false, because a normal female lion lacks a mane and a normal male lion cannot produce milk. The conclusion is more than false—it is necessarily false! Not only is abnormal for any lion both to have a mane and produce milk—it is biologically impossible. So the simple normality theory is stuck with the problem of how to assign non-contradictory truth conditions to the conclusion. Furthermore, in order to deal with the premises, the simple normality theory is driven to find some way of restricting the domain of quantification to males in the one case and females in the other case. And indeed, this is what Asher & Pelletier (1997) suggest. But there are substantial problems with that approach. Contextual domain restriction is, if anything, a form of *context sensitivity*. The interpretation of the first premise in (24) as being about males has nothing whatsoever to do with context—there are no contexts in which it is understood to be about normal female lions. Also, as Asher & Pelletier (1997) themselves observe, gendered generic sentences can be uttered discourse-initially without any risk of infelicity, unlike many examples of contextually domain-restricted quantified statements. So the gendered interpretations of such sentences cannot plausibly be argued to arise as a result of quantifier domain restriction.

On the other hand, it would be inaccurate to claim that these problems are insuperable. Nickel (2010b, forthcoming) responds to them by developing a more sophisticated theory of normality, which varies according to two parameters: way of being normal and respect of normality. Cohen (1999b) suggests that sentences like the premises in (24) are evaluated with respect to a set of salient alternatives. So there are many interesting avenues to be explored when it comes to dealing with these gendered examples quantificationally. The only point to make about the solutions on offer is that they come at the cost of incurring further theoretical complications. The most obvious complication is: if gendered generic sentences involve some sort of domain restriction, how does the mechanism responsible for

this domain restriction differ from/interact with the mechanism responsible for contextual domain restriction?

One advantage of this chapter's approach to these gendered examples is that it avoids all such complications, while remaining a normality theory at heart. As observed in section 3.3, product properties fail to retain the incompatibilities of their individual counterparts. The analysis under consideration, including definition (19), renders the argument pattern in (24) as a simple validity of propositional logic:

(25) $PM(maned)(lionKind)$
$PM(lactiferous)(lionKind)$
$\therefore PM(lactiferous)(lionKind) \wedge PM(maned)(lionKind)$

When thinking about generic sentences as generalizations over individuals, it is possible to ask: how could one single lion both have a mane and be able to produce milk? A single lion cannot have both of these properties. Thus, viewing gendered generic sentences as generalizations about a range of individuals raises a basic problem. But when generic statements are viewed as making claims about production processes, the problem never so much as arises, because there is nothing contradictory about a single production process producing things with differing properties. Though it so happens that being a maned feline and being a lactiferous feline are incompatible properties, being a *process* that creates maned felines and being a *process* that creates lactiferous felines are perfectly compatible properties.

3.5 The Kripke-Putnam Semantics for Kind Predicates

As mentioned in section 3.2, the account developed in this chapter shares some affinities with the semantics for natural kind terms proposed in Kripke (1980), lecture 3, which holds that kind predicates exhibit certain namelike properties. Kripke's view has two interrelated components: first, that kind predicates are in some sense rigid designators, and second, that they are irreducible to definition. To run these ideas through the standard example, the predicate *cat* has the same denotation across all circumstances of evaluation; if its denotation could vary across circumstances of evaluation in the manner of ordinary predicates, then (26-a) would be conceptually contradictory in the way (26-b) is:

(26) a. Imagine a situation where cats turned out to be robots rather than animals.
b. Imagine a situation where circles turned out to be straight, like a line.

Furthermore, *cat* is irreducible to any particular definition. There is no way to define *cat* as 'small, furry animal with whiskers that mews,' because it is always possible to imagine a counterfactual scenario in which the consensus among biologists was fundamentally deluded and these creatures really lacked these features. More generally, for any non-sortal predicate F, it always makes sense to ask:

(27) What if cats were discovered not to be F?

This shows, according to Kripke, that *cats* cannot be shorthand for 'things that have property *F*.' Its extension is not determined by a definition, but by reference to paradigmatic instances, somewhat in the manner of demonstrative noun phrases. To be a cat is to be one of *these* animals—not to be a creature with properties *X*, *Y*, and *Z*. Almog (2011) refers to this philosophical outlook as 'nature without essence'—where *essence* is taken to mean something like 'definition.'

Although Kripke arrived at this view through his discussions with Hilary Putnam, and the view is standardly referred to as the Kripke-Putnam view, there are several noteworthy differences between the view just summarized and what is actually argued for in Putnam (1962; 1970; 1975). On Putnam's picture, the extension of a natural kind predicate is determined by reference to paradigmatic instances. But whereas Kripke never specifies exactly how its extension is determined, Putnam provides the following narrative. First, the scientifically naive generation identifies samples of some substance, which they intuitively recognize to have something important in common. Although the best science of the epoch is not up to the task of precisely identifying this property, the extension of the kind predicate denoting the relevant substance is stipulated to contain whatever is the same as the paradigmatic samples with respect to this as yet unspecified property. Eventually, it is assumed, future generations will evolve the scientific sophistication that is required to uncover what the property is. In the case of chemical kinds, the property might be an atomic number or chemical formula. In the case of biological kinds, the property might be a certain genetic code.

For instance, one might imagine that an earlier generation first noticed a herd of giraffes wandering about in the wilderness. To refer to these creatures, citizens of this generation coined the term *giraffe*. What determines whether something is a giraffe is whether it shares some property (to be discovered in the future) with the initially identified specimens—not whether it has a long neck, eats leaves, is about fourteen tall, has two to three horns, or is yellow with brown spots. The earlier generation know that there is some such property, even though they lack the resources to identify it. Fast forward thousands of years, in the post-Watson and Crick era, and it becomes possible to say that biology has finally discovered what a giraffe is: namely, a creature with such and such genome. So the phenomenon of natural kind terms is a result of what is sometimes called *semantic deference*, the behavior whereby a community of speakers guiltlessly use terms whose exact meaning eludes them on the grounds that somewhere, a community of experts is conversant (or will later be conversant) in that exact meaning.

One rather tangential point here is that it is doubtful that having such and such genome really does capture what it is to be a giraffe, or a cat, or a person, in the sense

that Putnam requires. If having a human genome really is what it means to be a human, then it should be true that someone is a human if and only if they have that genome. And yet, it is possible for something to have that genome without being a human. For example, consider a beaker full of HeLa cells: these are cells cultivated some sixty years ago from the now deceased cancer patient Henrietta Lacks, but which are still among the most common human cells presently used for cancer research. They are endowed with the relevant kind of genetic spiral, and yet no one would claim that the beaker was a person.[30]

More relevant, however, is the fact that Putnam's story about how the extensions of natural kind predicates are determined stands in tension with the spirit of Kripke's view. Although Putnam is not committed to saying that *cats* just means 'animals with such and such genome'—he can claim that having such and such genome is feature possessed necessarily and a posteriori by cats—it still suggests that being a cat really amounts to nothing more than having the relevant genome. And that is substantially different from saying that being a cat amounts to being one of *these creatures*.[31]

The processual perspective on kinds developed in this chapter can be brought in to dissipate this tension. It can hold onto Putnam's narrative, in broad outline. The extension of the predicate *giraffe* will be determined by reference to paradigmatic giraffes, and further giraffes will fall under it if they stand in the relevant equivalence relation to the paradigmatic giraffes. But instead of taking that equivalence relation to be **share property F**, the new approach will take that equivalence relation to be **were produced by the same process**. It then becomes possible to retain the rest of Putnam's narrative, including the notion of semantic deference. But rather than deferring to the moment when future generations uncover the property that makes giraffes giraffes, the earlier generation will be deferring to the moment when future generations uncover what production process it is that gave rise to those initially observed giraffes. They will have known that it was some production process without yet being equipped to identify which one it was. This modified version of Putnam's account is no longer in tension with its Kripkean starting point, because it in no way impinges on the demonstrative character of the predicate *giraffe*. On such a view, the when our ancestors identify key specimens of giraffe-kind, they are pointing not to individual specimens, but to the production process that gave rise to those specimens, whatever that production process may turn out to be.

The intuitive metaphysical picture outlined in this chapter also addresses a well-known tension in Kripke's own view. If kind predicates are rigid designators, rigid designators have the same denotation in all circumstances of evaluation, and predicates denote

[30] This example is adapted from Thompson (2009), pg. 55.
[31] For further elaboration of this worry, see Almog (2011).

their extensions, then Kripke's view makes the unseemly prediction that kind predicates have the same extension across all circumstances of evaluation.[32] But obviously, predicates like *cat*, *giraffe*, or *human* vary in extension across different circumstances of evaluation. For example, in the nearby possible world in which Tina Fey has a sister, the extension of the predicate *human* differs from its extension in the actual world. Thus, although many have shared Kripke's intuitions about the rigidity of kind predicates, it has never been fully clear what such rigidity might actually amount to. And the processual perspective neatly resolves this difficulty, offering a simple way to make sense of the idea that kind predicates are rigid designators. The extension of the predicate *human*—the members of humankind—may vary across circumstances of evaluation. But the fact that it refers to the human production process will not.[33]

A small part of the philosophical payoff yielded by the processual perspective on kinds, then, is that it fits neatly into the Kripke-Putnam semantics for kind terms, and can even be brought in to resolve certain tensions in their accounts.

3.6 Taxonomic Generics

To round out the proposal, something needs to be said about an important subcategory of generic sentences which has received almost no attention in the literature thus far. This class of generics will hereafter be referred to as *taxonomic generics*. Taxonomic generics behave differently than the examples that are standardly discussed, insofar as they do not tolerate exceptions. Indeed, they may even express metaphysical necessities.[34] To get a sense of the distinction, contrast these two sentences:

(28) a. Jaguars are spotted.
b. Jaguars are mammals.

The first is a characterizing sentence of a familiar kind: it seems to express a generalization that can retain its truth in the face of (certain) exceptions. For instance, sentence (28-a) is true even though some melanistic jaguars have a pure black coat with no spots. (28-b), on the other hand, exhibits no such tolerance to exceptions. Not only are there no jaguars that fail to be mammals; it is difficult even to imagine how a jaguar could fail to be a mammal while remaining a jaguar. This was one of the original bases for Aristotle's distinction be-

[32] This problem is raised in Abbott (1989).
[33] Of course, drawing a distinction between the denotation of a kind predicate and the members of the kind it denotes will require some additional formalism. However, it will only require additional formalism that is already needed for independent reasons.
[34] Whether they ultimately do express metaphysical necessities is a complicated question that lies beyond the scope of this chapter.

tween attribute predication and substance predication—or, in contemporary terminology, between sortal and non-sortal predication.[35]

As observed in section (55), a rough grammatical diagnostic for this distinction is whether the predicate is an adjective or a noun: adjectival predicates are non-sortal predicates and nominal predicates are sortal predicates.[36] Rough though this diagnostic may be, it is not quite as rough as one might think. The second grammatical diagnostic from section (55) is also useful for distinguishing sortal from non-sortal predicates: only sortal predicates can serve as the answer to a 'what is it' question.

One of the ideas behind Aristotle's *Categories* is that predicating one kind of another (as opposed to predicating a non-sortal property of either a kind or another non-sortal property) yields an exceptionless truth.[37] Interestingly, generic statements exhibit sensitivity to a very similar distinction. Divide all predicates into sortal categories, based on whether or not they pass the grammatical diagnostics from section (55). The result will be that generic sentences with a sortal predicate in object position (such as (28-b)) have one set of truth conditions and generic sentences with non-sortal predicate in predicate position (such as (28-a)) have the characterizing truth conditions given in (19). There are a number of ways to construe the truth conditions of a sentence like (28-b), but the following two alternatives seem like the best two candidates for a starting point:

(29) a. **Option 1:** The jaguar production process $\sqsubseteq$ the mammal production process.
b. **Option 2:** Every jaguar is a mammal.

The best way to arrive at the truth conditions in (29-b) compositionally is probably not to use *PM* at all, but rather to interpret (28-b) as a distributive plural predication. The challenge there is to explain why bare plurals in subject position only seem to appear in these taxonomic constructions. Alternatively, the cost of exploring the first seems to be that the definition of *PM* becomes strangely disjunctive, having one meaning when it operates on an ordinary predicate, and an entirely different meaning when it operates on a kind predicate. Since that is a substantial, this chapter will pursue the second option. First, we stipulate that *PM* only has non-sortal predicates in its domain. Next, following Chierchia (1998), we define a type shifting rule that moves a NP denoting kind p to an NP denoting the complete atomic join semilattice generated by the set of p's progeny:

(30) $[\![\mathbf{NP}_p]\!] \rightsquigarrow [\![\mathbf{NP}_e]\!]$ = the maximal plurality of p's progeny

[35] See *Categories*, §1.1.

[36] It should be noted that this way of putting things is anachronistic, given that the category **adjective** didn't really exist until the Roman empire. See Luhtala (2005).

[37] This idea has been explored in detail by recent reconstructions of Aristotle's modal syllogistic, such as Malink (2006, 2013) and Rini (1998, 2011).

A benefit to either approach, provided the result can be achieved compositionally, is that it will offer a novel take on a traditional problem in the literature on generics: namely, the problem of explaining why the following pattern of inference is invalid:

(31) Chickens lay eggs.
Every egg-layer is female.
∴ Chickens are female.

Some have concluded from this example that no principle of deductive closure applies to generic sentences:

(32) **The Principle of Deductive Closure:**
If $\phi(x) \Rightarrow \psi(x)$, then $PM(\phi)(x) \Rightarrow PM(\psi)(x)$

Although the question whether generic sentences really do validate this pattern of inference is to a certain extent open, there are various indications that this pattern of inference does generally hold. For instance, the following instantiation of it is intuitively deductively valid:

(33) Cardinals have red feathers.
Everything with red feathers has feathers.
∴ Cardinals have feathers.

The more one explores this case, the more difficult it becomes to formulate an exception to deductive closure that differs substantially from (31). A better alternative, perhaps, is to understand the conclusion of (31) as a taxonomic generic statement, as in: 'chickens are hens.' Assuming that those truth conditions can be arrived at compositionally, the semantics for *PM* given in (19) predicts (31) to be invalid while still holding onto the principle of deductive closure. That seems like a welcome result.

3.7 Objections

The beginning of this chapter was focused on foundational and philosophical criticisms of the sophisticated kind theory. It is now time to turn to some of the standard *semantic* objections that have been levelled against the kind theory. It will be argued that the standard semantic objections to kind theories are not really objections to kind theories as such, so much as objections to the simple kind theory specifically. The lack of attention the sophisticated kind theory has received so far in the literature has caused the simple kind theory take center stage. But the existing semantic objections to kind theories leave the sophisticated kind theory untouched, and are nearly all addressed in one of two ways: either by adopting the sophisticated kind theory, or by acknowledging the distinction between habituality and genericity.

Scope Ambiguity

The presence of scope ambiguity in generics containing an indefinite noun phrase in object position is standardly assumed to pose a problem for the kind theory.[38] If generic sentences are nothing more than monadic predications, the indefinite noun phrase should have nothing to interact with scopally. The example typically invoked is the following:

(34) Swans have a favorite nesting spot.[39]

There are two readings of this sentence. On one, there is a single nesting area that swans prefer. On the other, swans tend to have a preferred nesting area, but different swans prefer different nesting areas. If generic sentences were to contain an unpronounced quantifier at logical form, this scope ambiguity could be explained in the usual way.[40] An analysis according to which generic statements are simple monadic predications lacks the resources to predict this ambiguity.

But of course, since it is only the simple kind theory that views generic statements as simple monadic predications, it is only the simple kind theory that lacks those resources. The sophisticated kind theory need only assume that *PM* is the sort of operator that exhibits scope effects with indefinites. Even ignoring the definition of *PM* in (19), the fact that *PM* is intensional already makes it extremely likely that generic sentences give rise to scope effects with indefinites, given that it would be odd for an intensional operator not to exhibit *de re*/*de dicto* ambiguities. Definition (19) confirms that *PM* indeed behaves as expected. Since *PM* is a well-defined logical operator like any other, the sophisticated kind theory on offer in this chapter can easily account for the ambiguity in sentence (34) using a standard theory of quantifier scope ambiguity, such as flexible types.

The flexible types approach to quantifiers is a method for dealing with two problems in one fell swoop: a) that natural language quantifiers give rise to scope ambiguity, and b) that they are uninterpretable in object position. Consider the following example:

(35) Every magician owns a rabbit.

Analogously to the generic example, sentence (35) has two available interpretations: one in which the existential quantifier phrase *a rabbit* scopes over the universal quantifier phrase *every magician*, and another in which the universal takes wide scope. *A rabbit* is uninterpretable in object position because dyadic predicates such as *own* are meant to take objects as arguments, not second-level functions. To deal with both of these problems, the flexible

[38] See Cohen (2001), pg. 193.

[39] This example originates from Schubert & Pelletier (1987), pg. 407. Interestingly, they do not seem to think it poses a problem for the kind theory.

[40] That is, either by type-raising (Hendriks, 1988) or quantifier raising (May, 1977).

types approach allows quantifiers to shift their denotation in one of the following two ways when in object position:

(36) a. Ordinary quantifier phrase: $\langle\langle e,t\rangle,t\rangle$
$[\![\textbf{a rabbit}]\!] = \lambda f_{\langle e,t\rangle} \,.\, \exists x(rabbit(x) \wedge f(x))$
b. Wide scope quantifier phrase: $\langle\langle e,t\rangle,t\rangle \rightsquigarrow \langle\langle e,\langle e,t\rangle\rangle,\langle\langle\langle e,t\rangle,t\rangle,t\rangle\rangle$
$[\![\textbf{a rabbit}]\!] \rightsquigarrow \lambda f_{\langle e,\langle e,t\rangle\rangle} \,.\, \lambda Q_{\langle\langle e,t\rangle,t\rangle} \,.\, \exists x(rabbit(x) \wedge Q(f(x)))$
c. Narrow scope quantifier phrase: $\langle\langle e,t\rangle,t\rangle \rightsquigarrow \langle\langle e,\langle e,t\rangle\rangle,\langle\langle\langle e,t\rangle,t\rangle,t\rangle\rangle$
$[\![\textbf{a rabbit}]\!] \rightsquigarrow \lambda f_{\langle e,\langle e,t\rangle\rangle} \,.\, \lambda Q_{\langle\langle e,t\rangle,t\rangle} \,.\, Q(\lambda x \,.\, \exists y(rabbit(y) \wedge f(x,y)))$

The expressive power of the lambda calculus makes it possible to systematically vary the definition of a single expression like *every* or *a* so that it can 'pass arguments' up the syntactic tree if need be. But if the lambda calculus provides the resources to account for scope ambiguity between quantifiers, negation, attitude verbs, and modals, why not employ the very same strategy to deal with scope ambiguity between quantifiers and the predicate modifier? An analogous solution for generics would allow the denotation of *a favorite resting spot* in sentence (34) to shift in either of the following two ways:

(37) a. Wide scope quantifier phrase: $\langle\langle e,t\rangle,t\rangle \rightsquigarrow \langle\langle e,\langle e,t\rangle\rangle,\langle e,t\rangle\rangle$
$[\![\textbf{a f.r.s.}]\!] \rightsquigarrow \lambda f_{\langle e,\langle e,t\rangle\rangle} \,.\, \lambda x \,.\, \exists y(restSpot(y) \wedge f(x,y))$
b. Narrow scope quantifier phrase: $\langle\langle e,t\rangle,t\rangle \rightsquigarrow \langle\langle e,\langle e,t\rangle\rangle,\langle\langle\langle e,t\rangle,\langle p,t\rangle\rangle,\langle p,t\rangle\rangle\rangle$
$[\![\textbf{a f.r.s.}]\!] \rightsquigarrow \lambda f_{\langle e,\langle e,t\rangle\rangle} \,.\, \lambda h_{\langle\langle e,t\rangle,\langle p,t\rangle\rangle} \,.\, \lambda p \,.\, h(\lambda x \,.\, \exists y(restSpot(y) \wedge f(x,y)))(p)$

(37-a) and (37-b) correspond exactly to the desired pair of readings for sentence (34):

(38) a. For some nesting spot x, swan-kind is a process that makes entities with a preference for x.
b. Swan-kind is a process that makes entities with a preference for some nesting spot or other.

This case is an excellent example of the kind of logical flexibility possessed by quantificational theories but lacking in the simple kind theory. Take-home point: the sophisticated kind theory has the same degree of logical flexibility.

Mosquitoes

Sarah-Jane Leslie's famous 'mosquito sentence' poses problems for every theory of genericity on the market, be it quantificational or kind-theoretic (Leslie, 2007, 2008):

(39) Mosquitoes carry the West Nile Virus.

Sentence (39) is particularly interesting to think about in connection with normality theories, because although native speakers robustly judge it to be true, not only do most mosquitoes not carry the West Nile Virus—it is arguably not even normal for a mosquito to carry the

virus. Nickel (forthcoming) suggests that sentence (39) is taken to be true because it has an ability reading. Thus, it poses no problem even for the simple normality theory of generics, because even though actually carrying the West Nile is not normal for a mosquito, *being able* to carry the West Nile Virus is. Mosquitoes are a vector for the virus.

That suggestion seems plausible, and the approach pursued by this chapter is to build on it with the following observation. The ability reading is not an artifact of generic statements; it is an artifact of habitual statements. To see this, consider the following non-generic habituals:

(40) a. Matt plays chess.
b. Matt eats meat.

Neither of those statements is a generic, in the sense outlined in section 1.3. However, both of them give rise to habitual/ability ambiguities. That is, they give rise to the following pairs of interpretations:

(41) a. Matt is able to play chess.
b. Matt makes a habit of playing chess at regular intervals.

(42) a. Matt can eat meat, if necessary.
b. Matt makes a habit of eating meat at regular intervals.

To hear reading (41-a), imagine that someone is needed to play chess at a party in order to entertain a seven-year-old who loves the game and is waiting to be picked up. They do not need to play it well; they only need to play it well enough to keep a child occupied for the duration of the evening. Unfortunately, no one at the party so much as knows the rules of the game. But at the eleventh hour, Matt shows up. In that context, interpretation (41-a) is indeed available for sentence (40-a). This can be seen from the fact that sentence (40-a) can still be true even if Matt rarely or never plays the game.

Contrast a non-habitual generic:

(43) Snakes are cold-blooded.

Sentence (43) does not give rise to the following pair of interpretations:

(44) a. Snakes are characterized by the ability to be cold-blooded.
b. Snakes are characterized by actually being cold-blooded.

This seems like a clear indication, then, that the ability interpretation is an artifact of habituals rather than generics. And so a fully compositional explanation of these ambiguities will have to wait for a fully compositional theory of habituals. For present purposes, it suffices to observe that this example offers useful lessons on the importance of treating genericity and habituality separately.

The Typhoons Example

The most influential critique of the kind theory was put forth by Carlson himself in a later paper (Carlson, 1989). In this text, Carlson drew attention to another kind of scope ambiguity involving prepositional phrase complements, arguing that his earlier theory failed to predict it:

(45) Typhoons arise in this part of the Pacific.[41]

This example is typically taken to have two readings, one on which it is saying something about this part of the Pacific, and one on which it is saying something about typhoons:

(46) a. It is characteristic of this part of the Pacific that typhoons arise here.
b. It is characteristic of typhoons that they arise in this part of the Pacific.

Whereas the scope effects in sentence (34) could be explained by the quantificational NP, there is no clear cause for the scope effects in this sentence, especially under a kind analysis. Therefore, this example led Carlson to recommend a disjunctive analysis of generic sentences, according to which they are sometimes statements about kinds and sometimes generalizations over individuals.

However, the standard understanding of this example's import glosses over the important distinctions drawn in section 1.3, and stands to benefit from being put under their lens. The first observation to make in this connection is that although generic, sentence (45) is also habitual. So strictly speaking, it lies outside the scope of the semantic analysis being put forth here. In addition, it is not entirely clear that there even are two distinct readings. If (46-a) and (46-b) indeed have different truth conditions, then it should be possible to describe a situation that verifies the one but falsifies the other. The literature on the typhoons sentence has not been forthcoming with such descriptions. Finally, it should also be noted that only the first of these two readings is generic—the second is merely habitual.

In order to get around these difficulties, we may vary the example to see whether the same phenomenon can be seen to obtain with a non-habitual generic. Milsark's typhoons sentence introduces unnecessary complications such as habituality and prepositional complements into the picture. But Carlson (1989)'s analysis predicts that a relational statement with one kind expression in subject position and the other in object position ought to give rise to the same kind of scope ambiguity. In principle, if we were to observe scope effects in this case, it would be the kind of counterexample the more self-critical incarnation of Carlson was after. Here is an example of such a sentence:

(47) Capuchin monkeys live with squirrel monkeys.

[41] This example comes from Milsark (1974).

If there were scope ambiguity in this sentence, then the two readings would be:

(48) a. It belongs to capuchin monkeys to live with certain squirrel monkeys.
(*Still true when squirrel monkeys generally do not live with capuchin monkeys.*)
b. It belongs to squirrel monkeys to live with certain capuchin monkeys.
(*Still true when capuchin monkeys generally do not live with squirrel monkeys.*)

Strikingly, the two readings in (48) are not present in sentence (47)—sentence (47) is simply false, in the event that capuchin monkeys generally do not live with squirrel monkeys. That makes it doubtful that any example in the vein of (45) will fall within the scope of the theory developed in this chapter, which concerns only non-habitual generics.

So in the end, more work is required to demonstrate that this sort of counterexample has traction. First, more of an argument that sentence (45) does indeed give rise to the kind of ambiguity alleged needs to be given, especially in light of Liebesman (2011)'s persuasive counterarguments against this.[42] Second, there needs to be some understanding of what more general phenomenon sentence (45) is an instance of. Meteorological generic sentences? Generic sentences involving locations? Generic sentences with prepositional phrases? That can be clarified by investigating a wider range of data. And finally, supposing that some form of ambiguity in this sentence is eventually isolated, it still remains to be determined whether such ambiguity only arises in habitual generic sentences, whether it occurs in ordinary generic sentences, and also whether it occurs in non-generic habitual sentences. Until then, sentence (47) casts serious doubt on whether such ambiguity is indeed a property of ordinary generic statements. So this counterexample is not particularly problematic for the analysis under consideration.

Context-Sensitivity

Sterken (2015a,b) argues that a kind-theoretic analysis of generics offers no obvious source for their context-sensitivity. That would seem to be at odds with one of the principal claim of Chapter 2, which is is that a kind-theoretic approach does the best job of accounting for the particular brand of context-sensitivity we see in generic sentences.

Thankfully, Sterken's view is not in fact at odds with what was argued in Chapter 2. The reason is that Sterken only considers the simple kind theory, rather than the sophisticated kind theory under discussion here. Indeed, it is at first difficult to see what source for context-sensitivity there could be in a simple kind theory along the lines described in section (12):

[42] See Liebesman (2011), §4.1.

(49) *F*s are *G* : *true iff G(f)*
(where f is an individual constant denoting F-kind)

Nonetheless, it should be noted, in anticipation of the proposal to be made later on in section 4.1, that with the following addition, even the simple kind theory can accommodate a certain amount of context sensitivity:

(50) *F*s are *G* : *true iff* $G(\iota k\colon \forall y(F(y) \rightarrow x \in k))$

That is, rather than taking the bare plural *Fs* to be the *proper name* of a kind, why not take it to be a *definite description* referring to a kind? As will be argued in section 4.1, a workable fully compositional kind theory will likely have to go that route anyway, for compositionality reasons.[43] In the above formulation, the bare plural *Fs* would refer to that kind of which everything in the extension of the predicate *F* is a member. Depending on how the iota operator is defined, it may very well exhibit some form of context sensitivity—for instance, it could have a similar semantics to the definite article, according to which it picks out the unique contextually salient object that fits the description in its scope.

Furthermore, the sophisticated kind theory outlined in this chapter provides an additional source for context sensitivity in the predicate modifying operation. According to definition (19), *PM* is a modal operator. Although definition (19) is not context sensitive, it would be easy enough to replace Base, which is defined within each model of the type-logical metalanguage, with a contextual background function of the sort familiar from Kratzer (1977, 1981). Were further evidence to support the hypothesis that *PM* is more like a modal auxiliary verb than it was originally made out to be, this would be an additional source of context sensitivity in generics.

So even the simple kind theory furnishes us with one potential source of context sensitivity, and the kind theory defended here could furnish us with as many as two, should that prove necessary.

Comparatives

Another criticism of the kind theory relates to its predictions vis-à-vis comparatives and equatives:[44]

(51) a. Horses are taller than cows. *comparative*
b. Cows are taller than horses. *comparative*
c. Cows are (exactly) as tall as horses. *equative*

[43] This is because anyone who understands the sentence 'Squirrels are black' and the words *from* and *Poland*, should have everything they need to understand the sentence 'Squirrels from Poland are black.'
[44] See Nickel (2010a).

Cows and horses are about 5 feet tall, on average. Horses exhibit greater variation: the shortest horses are shorter than any cow and the tallest horses are taller than any cow. Nickel claims that (51-a) and (51-b) are both false. But interestingly, he would also like to say that (51-c) is false. Horses and cows may have the same *average* height, but their heights have different statistical distributions: the histograms tabulating the respective heights of the two populations would not line up. Perhaps, in that case, (51-c) is false.[45] As Nickel observes, this assignment of truth values is logically ruled out in ordinary comparatives:

(52) a. Evelyn is taller than Vivian.
b. Vivian is taller than Evelyn.
c. Evelyn is (exactly) as tall as Vivian.

Unlike sentences (51-a)-(51-c), if (52-a) and (52-b) are both false, then (52-c) must be true, given that *taller than* is a linear order.

How is this a problem for the kind theory? To the author's knowledge, no kind theorist has yet proposed a fully compositional semantics for comparatives.[46] However, Nickel conjectures that a kind theory of comparatives would be required to take the following shape. Since bare plural noun phrases are proper names of kinds, the logical form of a generic comparative should be akin to that of an ordinary comparative:

(53) a. *a* is *G* : *true iff* ***G(a)***
b. *F*s are *G* : *true iff* ***G(F-kind)***
c. *a* is bigger than *b* : *true iff* ***size(a)*** > ***size(b)***
d. *F*s are bigger than *G*s : *true iff* ***size(F-kind)*** > ***size(G-kind)***

According to Nickel, the difficulty with such a treatment of generic comparatives is that it predicts (51-c) to be necessarily true if (51-a) and (51-b) are both false. But if Nickel's intuitions are correct, this inference pattern doesn't hold for generics, because two kinds can 'differ' in height without 'tying' in height.

Without making any definitive judgment about robustness of these data, which have yet to be fully tested, it may be observed that even if Nickel's judgments are vindicated by future experiments, they will only pose a problem for the simple kind theory—not for the sophisticated kind theory. As argued in section (12), there is ample independent reason to

[45] It is not easy to have clear intuitions about these cases, which probably means that more experiments need to be done on these data. Before coming to share Nickel's intuitions, it is first necessary to think one's way into the scenario for a while.

[46] Nickel attributes the account below to Krifka *et al.* (1995), but that article only discusses generic statements like (i-a), not generic comparatives like (i-b):

(i) a. *F*s get bigger as you head north.
b. *F*s are bigger than *G*s.

reject a simple kind theory anyway. The basic reason these observations are unproblematic for the sophisticated kind theory is that given its additional logical structure, it has no commitments one way or the other regarding the entailment from the falsity of (51-a) and (51-b) to the truth of (51-c). If that is discovered to be a correct entailment, it can be modelled by making one set of assumptions about the predicate modifier, and if it is discovered not to be, that fact can be modelled by making a different set of assumptions about the predicate modifier.

The sophisticated kind theory under consideration in this chapter posits a predicate modifier for monadic predicates, much like its earlier cousin in Carlson (1977a). To accommodate constructions featuring bare plural noun phrases in object position, the analysis will require either a new predicate modifier that shifts dyadic object relations to dyadic kind relations, or a type shifting rule for the original predicate modifier. There are many ways that strategy might be pursued in detail, and it is beyond the scope of this response to give a fully compositional analysis of generic comparisons. Since the current goal is only to show why even the most obvious extension of the sophisticated kind theory to transitive predicates is not committed to the truth conditions in (53), we may begin with the latter approach, for ease of exposition.

One straightforward form that type shifting rule could take would be the following:

(54) $PM(f_{\langle e,\langle e,t\rangle\rangle})(p)(q) \rightsquigarrow PM(\lambda y\ .\ PM(\lambda x\ .\ f(y)(x))(p))(q)$

Assuming such a rule, the truth conditions for a sentence like (51-b) would be as in (55-b), rather than (55-a):

(55) a. *size(F-kind) > size(G-kind)*
b. $PM(\lambda y\ .\ PM(\lambda x\ .\ x > y))(\textit{cow-kind})(\textit{horse-kind})$

Very roughly, the formula in (51-b) is true just in case horse-kind is characterized by the property of being shorter than the characteristic height of cows. And sentence (51-a) will be true just in case cow-kind is characterized by the property of being shorter than the characteristic height of horses. Here are the logical forms of sentences (51-b) through (51-c), respectively:

(56) a. $PM(\lambda y\ .\ PM(\lambda x\ .\ x > y))(\textit{cow-kind})(\textit{horse-kind})$
b. $PM(\lambda y\ .\ PM(\lambda x\ .\ x > y))(\textit{horse-kind})(\textit{cow-kind})$
c. $PM(\lambda y\ .\ PM(\lambda x\ .\ x \sim y))(\textit{cow-kind})(\textit{horse-kind})$

In order for (56-c) to be entailed by the falsity of (56-a) and (56-b), *PM* would have to have no scope effects with linear relations. But *PM*, as defined in (19), clearly does. Suppose it is false that at every world in the modal base associated with horse-kind, it is the case that for some horses, at every world in the modal base associated with cow-kind, some cows are taller than those horses. And suppose it is false that at every world in the modal base

associated with cow-kind, it is the case that for some cows, at every world in the modal base associated with horse-kind, some horses are taller than those cows. It doesn't follow that at every world in the modal base associated with cow-kind, it is the case that for some cows, at every world in the modal base associated with horse-kind, some horses are the same height as those cows. So adopting the type shifting rule in (54) is one way the theory put forth in this chapter could make Nickel's prediction.

A second option would be to have a more specialized type-shifting rule that only quantified over the accessible worlds associated with the kind in subject position:

(57) $PM(f_{\langle e,\langle e,\langle s,t\rangle\rangle\rangle})(p)(q) \rightsquigarrow \lambda f_{e,est} \,.\, \lambda q \,.\, \lambda p \,.\, ...$
$... \forall w'(\textsc{Base}(w)(p)(w') \rightarrow \exists x, y(P(x)(p)(w') \wedge P(y)(q)(w') \wedge f(x)(y)(w')))$

The lack of scope ambiguity in sentence (47) is a point in favor of this second approach. And the greater generality of the first approach is a point in its favor. Either way, we have found a way to expand the theory presented in this chapter so as to make the desired predictions. Therefore, there is no particular reason to think a quantificational theory is required to give a semantics for generic comparatives.

As for a positive story, this is not the place to articulate and defend a fully compositional analysis of generic comparatives. That would be a topic for a long paper unto itself. As always, there are many options, but for the time being it may at least be said that Nickel's own rather promising proposal can be implemented either in a quantificational framework or in a kind-theoretic framework. The proposal is to generic comparatives very closely on the model of definite plural comparatives, as in:

(58) The destroyers are bigger than the frigates.

Interestingly, this sentence is true even though the biggest frigates are bigger than the smallest destroyers. So its truth conditions must be weaker than 'every destroyer is bigger than every frigate.' Nickel proposes that in sentences like (58), conversational context supplies:

(59) a. a partition over the set of destroyers
b. a partition over the set of frigates
c. a comparability relation mapping every cell in one partition to a cell in the other

The cells of each partition correspond to relevant subcategories of destroyers and frigates, respectively. Suppose that there are three kinds of each, and call them sergeant, lieutenant, and colonel. Nickel suggests that sentence (58) is true just in case all the sergeant destroyers are bigger than all of the sergeant frigates, all the lieutenant destroyers are bigger than all of the lieutenant frigates, and so on. This seems intuitive. And furthermore, the destroyers and frigates can easily turn out not to have equal heights even though neither is is bigger than the other. Why? Nickel's truth conditions predict that this situation will arise if e.g.

some of the sergeant battleships are bigger than some of the sergeant frigates but some of them are smaller. If, for any pair of partition cells, it fails to be the case that every member of the one cell is bigger than ever member of the other, then the situation in which neither is bigger than the other has come to pass.

Nickel would like to say something similar about generic comparatives:

(60) Battleships are bigger than frigates.

Once again, conversational context supplies information about which relevant kinds of battleship are comparable to which relevant kinds of frigate. This time, however, rather than checking to see whether everything in the one kind of battleship is bigger than everything in the corresponding kind of frigate, sentence (60) checks to see whether every *normal* battleship of each subkind is bigger than every *normal* frigate of the corresponding subkind. The semantics works this way because Nickel endorses a quantificational normality theory, on which the truth conditions of '*F*s are *G*' is (very, very roughly) that every normal *F* be *G*.[47]

Nickel's positive proposal fits rather well with the analysis presented in this chapter. On that picture, definite plural comparatives contain a distributive operator that interacts with the contextually supplied information about how to partition the extensions of each predicate, and generic comparatives are more or less the same, except that they have a generic quantifier in place of that distributive operator. But of course, that operator is located exactly where the predicate modifier would be in a kind-theoretic analysis. Instead of verifying whether all normal battleships of each subkind were bigger than all normal frigates of the corresponding subkind, it could verify whether the characteristic height of each battleship subkind were larger or smaller than the characteristic height of the corresponding frigate subkind.

What this example demonstrates is that the sophisticated kind theory is no less expressively flexible than a quantificational theory. As with many of the arguments against the kind theory that have been considered so far, the upshot of the argument from generic comparatives is that generic sentences have more logical structure to them than the simple kind theory posits. Such arguments can be added to the continually growing body of evidence against the simple kind theory. These arguments in favor of quantificational theories favor the sophisticated kind theory equally strongly.

[47] Here, the simple normality theory is going proxy for Nickel's multidimensional normality theory, in order to make for an easier illustration of how generic comparatives work. For Nickel's full normality theory of generics, see the summary in section (8).

3.8 Kind Semantics vs. Quantificational Semantics

Now that the sophisticated kind theory has been laid out, it is worth turning to the matter of why it should count as non-quantificational. Given that the definition of the predicate modifier has quantifiers in it—in particular, modal universal quantifiers over possible worlds—it is reasonable to wonder whether the sophisticated kind theory ultimately collapses back into the quantificational theory. A second reason to think the semantics for *PM* ought to count as quantificational is that it is a normality theory. As seen in section 2.2, the quantificational theory of Asher & Pelletier (1997) achieves its results by quantifying over normal possible worlds, and the quantificational theory of Nickel (2010b, forthcoming) achieves its results by quantifying over normal individuals. Given that *PM* also quantifies over normal worlds, in virtue of what does the analysis presented here merit the label *kind theory*?

On the first point, that the analysis put forth in this chapter quantifies over possible worlds is an inevitable symptom of the fact that generic statements have intensional purport. Everyone in the literature agrees that generic statements either themselves have intensional purport or directly entail statements with strong intensional purport.[48] To evaluate their truth, a speaker must consider not only what actually is the case, but what is the case in other non-actual situations. And it so happens that the standard semantics for intensional statements interprets them as quantifying over possible worlds. This is of course not to deny that there exist several formal theories of intensionality which make no reference to possible worlds.[49] There is no a priori proof that the analysis put forth in this chapter could not in principle be translated into one of these alternative, non-mainstream frameworks. That qualification aside, to the extent that generic statements have intensional purport, and to the extent that the standard theoretical apparatus for intensionality involves quantifying over possible worlds, quantification over possible worlds (or some equivalent thereof) is inevitable in a semantic analysis of generic statements.

There is even a case to be made that some notion of normality is inescapable in a theory of generic statements. The most significant competitor to normality semantics for generic statements would be the probabilistic semantics put forth in Cohen (1999b; 1999a). But the debate about how to interpret probability philosophically is ongoing, and it is still very much up for negotiation whether probability statements are a kind of modal statement, whether modal statements are a kind of probability statement, or whether the two

[48] Noting, as always, Nickel (forthcoming) as the lone exception to this consensus view. And even Nickel (forthcoming) argues that generic statements 'go intensional' in cases where the NP in subject position has a null extension.

[49] One interesting example that gives a proof-theoretic analysis of modality is Restall (2012). Another is the semantics articulated in Warmke (forthcoming), called *ersatzism*, which analyzes modality in terms of propositional properties.

are distinct. Cohen (2013) offers an interesting argument that normality and probability theories make different empirical predictions, but the examples provided there, though of great philosophical interest, are subtle and inconclusive. If it turns out that philosophically interpreting probability requires appeal to some notion of normality, then there is a sense in which even probabilistic theories are normality theories.

More broadly, there is a distinction between what kind theories and quantificational theories take generic statements to be about. Quantificational theories view generic statements as general statements about individuals, and kind theories view generic statements as particular statements about kinds. According to a quantificational theory, what is the case vis-à-vis a kind's members form part of the generic sentence's truth conditions. According to a kind theory, the truth of a generic statement does not directly depend on what is the case vis-à-vis the members of a kind; it only depends on whether the kind has the relevant property. This is not to say that what is the case vis-à-vis the actual members of a kind is irrelevant to the truth of a generic statement about that kind. It is only to say that what is the case vis-à-vis the members of a kind is best understood as *evidence* for or against the truth of a generic, rather than as a part of its truth conditions. The relation between facts about individual members and facts about the kind itself is best understood under the rubric of confirmation theory in epistemology (Hawthorne, 2011).

Finally and most importantly, on the second point, there is a subtle but crucial distinction between the respective notions of normality at play in quantificational normality theories and in the normality theory developed in this chapter. The key point here is that the conversational background function Base takes as input not an individual, but a kind. The result is that the notion of normality to which this semantics appeals is kind-level normality, rather than individual-level normality. And the set of ideal outcomes for an individual is simply not the same as the set of ideal outcomes for a kind. For instance, the ideal outcomes for an individual person are much more specific than the ideal outcomes for human-kind: they might include outcomes in which she successfully installs the bleeding-edge version of an operating system on recalcitrant hardware, in which she runs a marathon in record time, or in which she learns how to play the guitar solo from *Beat It*. None of these would be ideal outcomes for the evolutionary process that produces people. Ideal outcomes for that process would include creating some reasonable number of reproductively fit individuals, sustaining itself rather than coming to an end, and so on.

Considering what is normal at the level of a production process makes certain philosophical tasks simpler because many of the problem cases for normality never arise at the level of the production process. In what sense it is normal for *this male lion* to give birth to live young? One could try to develop a notion of normality that is philosophically flexible enough to predict, somehow, that something an animal could never do was normal for it. But a simpler option is to consider what is normal for a process that makes things. By tak-

ing the processual perspective not just on kinds but on normality, the philosopher avoids having to make pre-theoretically intuitive notions like that of the normal more theoretically supple than they need to be.

That individual-level notions of normality are forced to take on additional theoretical complexity can be seen from the fact that, for example, Nickel (2010b, forthcoming) cannot get by merely with brute normality. Instead, the operative notion of normality varies according to two parameters. Alternative normality theories, such as the default logic presented in Bastiaanse & Veltman (2015), are pushed in exactly the same direction. But let us focus on the first example, since it is more explicit about the motivation for representing normality as a dual-parameter notion.

According to the analysis in Nickel (forthcoming), there is no such thing as normality *tout court*. There is only such a thing as being F-normal with respect to a determinable of G. To see why the individual-level notion of normality requires bulking the pre-theoretical notion of normality up in this way, suppose that *normal* is a standard extensional intersective adjective. Then the following statements come out inconsistent:

(61) a. Pengo is a normal penguin.
$normal(pengo) \wedge penguin(pengo)$
b. No normal penguin flies.
$\forall x((normal(x) \wedge penguin(x)) \rightarrow \neg flies(x))$
c. Every penguin is a bird.
$\forall x(penguin(x) \rightarrow bird(x))$
d. Every normal bird flies.
$\forall x((normal(x) \wedge bird(x)) \rightarrow flies(x))$

That is an unhappy result. Intuitively, normal penguins have different properties from normal birds, even though all penguins are birds. To get around that difficulty, then, quantificational normality theories introduce the first parameter of variation. Instead of working with the notion what is normal, they work with as many notions of F-normality as there are different properties that can be substituted for F: penguin-normality, bird-normality, whale-normality, mammal-normality, and so forth.

So suppose now that *normal* is a subsective adjective, like *skillful*, and consider the following pair of default inferences:

(62) a. Trish is a normal turtle $\rightsquigarrow$ Trish has a long lifespan
b. Trish is a normal turtle $\rightsquigarrow$ Trish is eaten by predators at a young age

There is a strong pre-theoretical intuition that it is normal both for a turtle to be eaten by predators at a young age and for a turtle to grow old. It is metaphysically impossible for a turtle to be normal in both of these ways simultaneously, and yet both seem to be ways

of being a normal turtle.[50] Nickel resolves this dilemma by saying that there are many different respects not just of normality, including turtle normality and reptile normality, but many different respects of turtle normality, many different respects of reptile normality, and so on for any property *F*. It is normal with respect to a turtle's encounters with its ecological environment to be eaten at a young age, and it is normal with respect to a turtle's ontogeny to grow old. To generalize, for every property *G*, there is a way of being normal with respect to a determinable of *G*.

Whether normality is, in fact, multidimensionally variable in this way is a fascinating philosophical question. There may turn out to be further considerations which push philosophers of science, metaphysicians, or semanticists in that direction. But even if that turns out to be the case, it is still worthy of note that viewing normality at the level of the production process avoids all of these complexities. A philosophical picture that never faces the task of imputing conflicting properties to a single individual can go back to working with a notion of normality *tout court*. Although no turtle can enjoy both the outcome of living a long life and the outcome of being eaten early on by a predator, there is no contradiction whatsoever in supposing that a normal outcome for the evolutionary process responsible for producing turtles is one in which most of them are eaten at a young age, and in which it creates more turtles than are necessary for it to sustain itself. And since penguin-kind and bird-kind are different production processes (related, of course, as part to whole), there are no constraints on how the ideal outcomes associated with each ought to be logically related.

There is a fundamental philosophical trade-off here. One can hold onto the more intuitive unidimensional notion of normality while introducing a new entity type that is a bit less likely to occur in explicit folk reasoning—the kind. Or one can hold onto the standard inventory of entity types while introducing a level of complexity into the operative notion of normality that goes well beyond anything that is likely to occur in explicit folk reasoning. Were the question strictly a priori, the trade-off might be argued to be even. It is the author's view that the empirical observations of Chapter 2 make the former alternative preferable, all things considered. But irrespective of which alternative is to be preferred, the very existence of such a trade-off establishes that there is a substantive choice underlying the decision to go with a quantificational analysis or a sophisticated kind-theoretic analysis, and is the real reason why the kind theory is ultimately worthy of the name.

[50] To feel the force of this intuition, consider how neither consequent in (62) would be surprising, given the truth of the antecedent.

3.9 Summary

It is time to take stock. Chapter 2 drew attention to some linguistic data that pose problems for the quantificational approach to generics. The kind theory seemed as though it had the potential to fare better with respect to these data, but the simple kind theory was ruled out on the basis of difficulties already raised in the literature. The remaining alternative was the sophisticated kind theory, which, in spite of its promise, has received almost no discussion in the literature on generics since its original formulation in the 1970s.

But the sophisticated kind theory, in its original formulation, was vulnerable to the worry that it might be based upon an operator that magically yields the desired result, with no explanation offered as to how. This chapter aimed to show that although the sophisticated kind theory might have the prima facie appearance of a black box theory, the core operator behind the analysis in fact has an intuitive intensional meaning which is easy to spell out in the idiom of model-theoretic semantics.

The next step was to spend some time explaining this intensional meaning, which involved sketching out the intuitive metaphysical picture underlying it. This intuitive metaphysical picture was discovered to have philosophically advantageous features. It offers a simple explanation for generic sentences involving sexual dimorphism by offering a general explanation of how incompatible properties can characterize the same kind. It understands the kind of normality that is operative in generic sentences as akin to bouletic modality, which allows it to make the correct predictions about the truth conditions of 'unintended consequences' generics. It resolves a major tension in the Kripke-Putnam view on natural kind predicates, reconciling widespread intuitions that such predicates have an element of rigidity to them with the unavoidable fact that their extensions must vary from possible world to possible world. And most interestingly, it makes available a notion of kind-level as opposed to individual-level normality, which is capable of doing the philosophical work of more baroque multidimensional normality theories in a simpler way.

At the same time, it is not the intention of this text to defend the intuitive metaphysical picture sketched out in this chapter. The intuitive metaphysical picture may very well be mistaken, and it may very well be correct. That is a question to be taken up on another occasion. All that is required for the purposes of this chapter is to establish that the metaphysical picture sketched out has enough intuitive plausibility as a folk metaphysics to serve as a helpful interpretation of the model-theoretic analysis of PM. If this metaphysical picture can be of use in making sense of what sort of thing the atomic type p is supposed to be, what the conversational background function Base maps to what, and how the extensions of the logical $\sqsubseteq$, N, and P relation symbols are to be assigned, then it will have achieved its purpose, which is to show that the sophisticated kind theory is not a black box theory. Furthermore, even though the goal of this chapter is not to defend the processual

perspective on kinds as a view in metaphysics proper, it retains its philosophical interest as a case study in how undertaking a fully compositional semantic analysis of a fragment of natural language can spur exploration into philosophical territory that would otherwise have remained unexplored.

Recall the desiderata laid out at the beginning of the chapter. The lexical entry for *PM*, put together with this philosophical interpretation a) is the sort of static semantic analysis that can easily have some further dynamic default reasoning framework built on top of it,[51] b) explains the difference between a permissible and an impermissible exception to a generic, and c) offers a general explanation of the principle in virtue of which object-level properties are mapped to the relevant kind-level properties. Carlson (1977a)'s second principal motivation for positing *PM* (the first was the need to make the correct predictions about bound variable interpretations) was that in I-generics, there seemed to be some systematic relation between object-level properties and kind-level properties. But in spite of this, Carlson never explained what that systematic relation was. This chapter fills in that important gap, vindicating the sophisticated kind theory and freeing it up to do the work set out for it in Chapter 2.

[51] Asher & Pelletier (1997)'s static conditional logic for generics serves exactly this purpose.

Chapter 4

Accounting for the Three Contrasts

Chapter 2 examined three distinctions between quantified sentences and generic sentences. These distinctions pose problems both for quantificational theories and kind theories. They pose problems for quantificational theories in the sense that they bring out how explicitly quantified sentences differ from generic sentences in their interpretation. If generic sentences truly did contain an unpronounced quantifier, we would expect them to pattern with the explicitly quantified sentences. Generic sentences pattern differently from quantified sentences in this regard, which sheds doubt on the hypothesis that they contain a quantifier. But these observations also raise questions about kind theories, insofar as presently available kind theories have no explanation for them. It's important to recognize that the quantificational theory has no explanation for them either, which at this stage means that the two approaches are on equally uncertain footing.

The aim of the coming chapter is to tip that balance. The extended kind theory proposed in Chapter 3 makes it straightforward to provide a compelling explanation of these phenomena, where analogous extensions to the quantificational theory are unsatisfactory at best, and unfeasible at worst.

The strategy breaks down as follows. First, we will consider the cohesiveness presupposition on the subject of generic sentences, explaining what it means for a predicate to be cohesive. The definition of cohesiveness articulated in Chapter 2 was a sort of placeholder, meant to draw on some basic intuitions about which predicates have random, miscellaneous, or otherwise haphazard extensions and which do not. This chapter replaces that skeletal account with something more substantial, employing the foundations laid down in Chapter 3 to provide a simple explanation for these intuition patterns. The cohesiveness presupposition is potentially problematic for a quantificational theory, given that (as far as I am aware) no natural language quantifier is known to place any kind of cohesiveness condition on its restrictor predicate. The classical kind theory fares better, since it correctly predicts that generic sentences will be infelicitous if the NP in subject position fails to refer.

However, the classical kind theory has nothing whatsoever to say about what a kind is or what it means for a category to have kindlike cohesion, which means that it can only explain when kind-referring terms fail to refer by adverting to pre-theoretical intuitions about when members of a predicate's extension have something interesting in common and when they do not. This chapter aims to tell a story about what those intuitions are based on.

After having shown how the extended kind theory from Chapter 3 can shed light on the cohesiveness presupposition, the next step is to explain why generic sentences give rise to artifactual/non-artifactual ambiguity, and what general constraints there are on artifactual interpretations: when they arise, when they fail to arise, whether there are further distinctions to be drawn between multiple artifactual interpretations, and so forth. Artifactual interpretations also pose problems for the quantificational theory, given that quantified sentences tend not to allow for them. Similarly to the cohesiveness presupposition, while this phenomenon poses no a direct problem for the classical kind theory, the classical kind theory will need to be expanded in order to have anything to say about it. However, as we shall see, the proposed expansion is simple and intuitive.

Finally, there is the phenomenon of quantifier domain restriction. If there were such a thing as a generic quantifier, one would expect generic sentences to undergo the same kind of contextual domain restriction as quantified sentences. Strikingly, they do not. The data here are subtle. As we saw in Chapter 2, the status of counterexamples in the vein of Condoravdi (1992, 1997) is not entirely clear, and native speakers vary in their judgments about them. But whatever their status, the result poses problems for the quantificational theory. The lack of clarity regarding these data suffices to rule these out as clear-cut counterexamples. A second alternative would be to recall that Condoravdi's own view about such sentences, which represents the present consensus within natural language semantics, is that they are strict universal (rather than generic) statements. Assuming that standard wisdom to be correct, they are once again not counterexamples to the claim that generic sentences do not contextually domain restrict.

But of course, both lines of response would come across as evasive if they ended there. It would be more of an achievement to show that even if these were indeed cases of generic statements interacting with conversational context we would still have strong reasons to doubt that these generic sentences involve anything like *domain restriction*. One reason is that domain restriction comes in two varieties: deictic and anaphoric. As we saw, contextual narrowing in these examples only occurs anaphorically—not deictically. Some authors have proposed that there are 'anaphoric-only' quantifiers, but those cases seem to be few and far between. A second reason for doubting that Condoravdi-style cases involve contextual domain restriction is that quantified sentences domain restrict as a kind of default: it is only in the exceptional case that quantified sentences are interpreted unrestrictedly. By contrast, the default for generic sentences is to be interpreted in the broadest way possible.

The final section of this chapter will argue that the best analysis of these cases, even if they turn out to be admissible as data, is as involving the very same mechanism that makes artifactual interpretation possible, plus a process of accommodation whereby the existence of a new kind is added to the common ground as a last resort pragmatic repair procedure. A kind-theoretic analysis is ideally positioned to account for the meanings of these cases as well as their limited distribution. Furthermore, only a kind-theoretic analysis can offer a unified explanation of the three data points under discussed. A quantificational account of these cases would take them to involve saturation of free variables, which is anything but a last-resort pragmatic repair mechanism. Generic sentences only receive their narrower interpretations in very particular circumstances, and only an analysis that sees these interpretations as arising out of accommodation can do justice to the particularity of these circumstances.

The core approach of this chapter can be summed up as follows: explain the three phenomena from Chapter 2 in terms of the well-known semantic properties of definiteness. The cohesiveness presupposition will be explained as an existence condition, the availability of artifactual interpretations will be explained as a familiarity condition, and the lack of contextual domain restriction will follow naturally from the fact that although the iota operator used in the analysis is quantificational,[1] it quantifies over kinds, rather than individual objects.

4.1 The Cohesiveness Presupposition

Chapter 2 observed that generic sentences are more selective than quantified sentences about which predicates they will accept in subject position. To make a general statement using a quantifier, any predicate will suffice; whether the members in the extension of said predicate have anything in common is irrelevant. Generic statements, by contrast, are most naturally made about a class of things whose members have something significant in common. The name *cohesive* was set aside for predicates that denote such a class, and the name *haphazard* was set aside for predicates that do not.

Thus far, we have remained silent about what exactly it means for a class of things to have something meaningful in common. The following can be a helpful starting point for easing into that difficult question. We all share a basic set of intuitions that green things, for example, are all green for different reasons.[2] A tennis ball may have been dyed green,

[1] According to an inclusive conception of quantification, at least. It should be acknowledged that there is a debate in the literature as to whether the definite article should count as a quantifier, and more generally as to whether there are determiners that are not quantifiers. See Fara (2001); Glanzberg (2007).

[2] These intuitions surely go back to the opening of Aristotle's *Categories*.

a parrot may be biologically disposed to grow feathers with green pigment, a car may have been painted green, an Apple IIc display may have been designed to emit green light, and so on. There is a temptation to say that this group of objects is only 'accidentally' green, not in the sense that their nature disposes them to be something other than green (it is surely in the parrot's nature, after all, to be green), but in the sense that each is green for a different reason. Since there is a different explanation for why each of these objects is green, no additional information about further features that one of them may have follows from learning that any single one of them is green. Or, more precisely, any further features which follow from any one of them being green will of necessity pertain to color. The set of giraffes, for instance, is different. One single cause is responsible for making all of them giraffes. An interesting ramification of this fact is that once we discover some animal to be a giraffe, we can expect it to have further features that this single cause may also have brought about—features which pertain to topics other than biological species.

Strikingly, the distinction between cohesive and haphazard predicates seems to be encoded in the more colloquial registers of English. Consider the expression *a thing*:

(1) a. Hipsters who can code are, like, a thing now.
b. What? Since when are bilingual Americans a thing?
c. Whoa. Did parking lanes, like, suddenly become a thing while I was away?

What sorts of things merit the label 'a thing'? Kinds about which it is possible to make a generic statement. If a predicate is cohesive, then when it occurs in a bare plural noun phrase, it can be used in a generic sentence, and it can be described as 'a thing.' Recall the observation from Chapter 2 that a predicate which is haphazard in a particular conversational context sounds awkward in subject position of a generic sentence. This *a thing* construction can be used to register that awkwardness in conversation, and even to help massage the conversational background into one in which it is cohesive:

(2) **Jean:** People hacking into my website right now are from France.
Joan: What do you mean, 'people hacking into your website right now'?
It's not like people hacking into your website right now are a thing.

Finally, although it would be fanciful to make too much of the following, it is surely at least suggestive that *a thing* can be directly predicated of a bare plural noun phrase. Specifically, the suggestion is that although it features a plural suffix, a bare plural noun phrase is something which, at some level, we treat as denoting a particular. Somehow, classes whose members have something meaningful in common seem to be connected with 'things.'

Packaging this set of intuitions into a more precise definition presents a set of challenges which warrant particular caution. According the account sketched out in Chapter 3, the factor differentiating predicates like *giraffe* from predicates like green is that they are *sortal* predicates. Sortal predicates differ from non-sortal predicates, on the one hand, in

what it takes to be competent at using them, and, on the other hand, in which grammatical constructions comfortably admit them. In order to be competent at using a non-sortal predicate P, one needs (at least) to know how to sort those things that fall into its extension from those things that do not. But being competent at using a sortal predicate requires a bit more. In order to be competent at using a sortal predicate S, one needs (at least) to be able to know what it takes for one S to be the same S as another, and how to count Ss.

That was one way of explaining the difference between predicates we pre-theoretically characterized as cohesive and predicates we pre-theoretically characterized as having neither rhyme nor reason to them: cohesive predicates tend to be sortal predicates, and haphazard predicates tend to be non-sortal predicates. But ultimately, that was a transitional picture, undertaken without the resources made available by the account in Chapter 3. With those additional resources in hand, it is possible to give a simpler and more attractive account of cohesiveness. The observations from Chapter 2 still apply—it is accurate to say that generic sentences select for sortal NPs in their subject position. However, it is now possible to supplement that sketch with a richer story about what it means for the members of a class to have something significant in common.

The semantics from Chapter 3 leads naturally to the following definition of cohesion: a predicate is cohesive just in case everything that falls under its extension was brought into existence by the same production process. Where p is a variable ranging over production processes, and P is the progenitor relation from Chapter 3:

(3) **Cohesive Categories**
A predicate F denotes a *cohesive* category iff:
$\exists p : \forall x(F(x) \leftrightarrow P(p, x))$

I believe this is the best way to cash out the intuitive idea, invoked in much of the literature on genericity, that certain noun phrases refer to 'well-established kinds.'[3] Strictly speaking, according to the present account, it is not quite the case that the predicates under consideration literally denote kinds. What a predicate denotes is its extension (or the characteristic function thereof). But what semanticists are trying to capture when they consider whether an NP is well-established kind-referring is, I believe, something like the following. The predicate *giraffe* is cohesive, because everything in its extension was produced by a single production process: the giraffe production process. The predicate *green* (by contrast) is haphazard, because not everything in its extension was produced by any particular production process. Tennis balls are produced in factories, parrots are produced through sexual repro-

[3] See especially Krifka *et al.* (1995), pp. 11-13, 67-70, and 93-110 for the original use of the term, and Oosterhof (2008), Dobrovie-Sorin & Pires de Oliveira (2008), and Pelletier (2009) for examples of more recent usage. Note that Pelletier is a bit more careful with his terminology, preferring to talk of NPs being 'semantically connected' to well-established kinds.

duction, and so on. In other words, there is no particular 'green thing production process.'[4] Rather, the different green things that there are were created by different production processes.

Pelletier talks of a 'semantic connection' between certain NPs and well-established kinds, and suggests that certain constructions select for NPs which exhibit this special connection.[5] That way of talking is more on the right track, and it allows us to ask: what is a well-established kind, and what does it mean for an NP to be semantically connected to it in the relevant way? Krifka *et al.* (1995) state that although we have strong intuitions about which kinds are well-established and which are not, it is exceedingly difficult to specify what underlies those intuitions.[6]

However, the prospects for a theory of well-established kinds are perhaps not as bleak as the received wisdom on the matter would suggest. The theory of kinds as production processes can be used to formulate elegant answers to the two questions raised at the beginning of this section. On the matter of what a well-established kind is, it would define a well-established kind as a production process that has been admitted into the common ground of the conversation. This makes possible a straightforward explanation of what it would mean to establish a kind in the course of a conversation: namely, establishing a kind in a particular context would mean introducing a production process into discourse. And on the nature of the 'semantic connection,' the theory of kinds as production processes would say that at the level of the whole bare plural noun phrase, the connection in question is indeed one of reference. But at the level of the predicate within the bare plural noun phrase, the connection is the relation borne by any cohesive predicate to a production process that created everything falling under its extension. The following example may be used to illustrate this terminological distinction:

(4) Bears have fur.

In this sentence, the predicate *bear* denotes the set of all bears (or the characteristic function thereof), and the bare plural noun phrase *bears* denotes the production process responsible for producing everything in that set: bear-kind. This represents a slight departure from Carlson (1977a)'s original proposal, which was that bare plural noun phrases are *names* of kinds. But it is more faithful both to the spirit of Carlson's original proposal and to what we know about compositionality to construe bare plural noun phrases as something more like definite descriptions referring to kinds. The assumption that they literally are proper

[4] Note that the definition in (3) allows for the possibility that everything in a predicate's extension was produced by *more than one* production process. All that's required for a predicate to be cohesive is that everything in its extension be the progeny of *at least one* production process.

[5] Pelletier (2009), pg. 5.

[6] See Krifka *et al.* (1995), pp. 11, 13.

names that refer to kinds leads into compositionality worries. To see this, consider a generic sentence with a complex noun phrase in subject position:

(5) Bears from North America have fur.

Supposing that bare plural noun phrases are names of kinds leads directly to one of two conclusions about sentence (5). The first alternative is to argue that sentence (5) is not really a generic statement, and the second alternative is to argue that *bears from North America* is really a name, despite its surface appearance as a complex NP. Both options are counterintuitive. It is certainly the case that understanding the meaning of certain NPs, such as *The Bay of Pigs*, requires a native speaker to add a new word to her lexicon.[7] Nonetheless, such NPs are clearly the exception that proves the rule. In general, it must be possible for a native speaker to understand what a sentence with any arbitrary complex NP means without having to learn any new words. So in order to respect the possibility of generic sentences about categories of things that have never before been explicitly referred to, it is necessary to grant that cohesive predicates are indeed predicates, and are made into expressions that denote kinds in the course of becoming noun phrases.[8]

The view that bare plural NPs are proper names also makes unseemly predictions about simple bare plural NPs. For example, it would suggest that the NP *bears* has no special syntactic relation to the predicate *bear*; that people learning English would have to learn each of these words separately. This is highly counterintuitive, for similar reasons.

A better alternative would be to suppose that *bear* is a predicate which gets mapped to a kind-denoting expression when it is in a bare plural noun phrase. There are many ways to achieve this result, but one straightforward option is to adopt an analysis on which *-s*

[7] Or, at least, something like that is required under the orthodox analysis of proper names in natural language as behaving like individual constants in logic. Although the theory presented here will work under the orthodox assumption, it should at least be mentioned that a growing constituent of semanticists have raised some serious challenges for the traditional post-Kripkean analysis of proper names. See Burge (1973), Geurts (1997), Elbourne (2005), Gray (2012), Muskens (2011), and Fara (forthcoming). Note that since adopting the predicate view of names is a way of reducing proper names to descriptions, adopting the predicate view in effect grants that bare plural noun phrases, if they indeed are referential expressions, must be something more like singular terms.

[8] Of course, this is one of two possible approaches. Another option, explored in Chierchia (1998) and Zamparelli (2000), is to analyze common nouns as names of kinds, then as being mapped into predicates either via type shifting operations or via operators that are syntactically higher up in the NP. Cohen (2004) proceeds in the other direction, analyzing common nouns as underlyingly predicative, with the ability to be mapped to singular terms referring to kinds via a complementary set of type shifting operations. There are many factors that might play a role in adjudicating between these two approaches, some of which are addressed in McNally (n.d.). But for the purposes of the discussion here, although both are on equal footing insofar as they respect the principle of compositionality, the former approach involves an additional type shifting step with no immediate payoff. Acknowledging, then, that there is still a debate to be had on this issue, this text will pursue the latter approach.

denotes a function mapping cohesive predicates to the production process responsible for creating everything that falls under their extension:

(6) $[\![\textbf{-s}]\!] = \lambda f_{e,t} \,.\, \iota p : \forall x(f(x) \rightarrow P(p,x))$

The following example will illustrate how this analysis works:[9]

(7) Hackers are security-conscious.

(8) a. $[\![\textbf{hacker}]\!] = \lambda x \,.\, hacker(x)$
b. $[\![\textbf{-s}]\!]([\![\textbf{hacker}]\!]) = \iota p : \forall x(hacker(x) \rightarrow P(p,x))$
c. $[\![\textbf{hackers are security conscious}]\!] =$
1 iff $[\![\textbf{PM}]\!](\lambda x \,.\, securityConscious(x))(\iota p : \forall x(hacker(x) \rightarrow P(p,x))) =$
true just in case whenever the hacker production process succeeds in doing what it is supposed to, it creates some things which are security conscious

Intuitively, this sentence is felicitous because the hackers in the world were indeed created via a common process: namely, the process of initiating people into a particular social community, where one is expected to have certain values and have special technical skills. Since it is generally acknowledged that a single process is responsible for creating all the hackers in the world—the technology world forms a single, global community—the existence of that process will be part of the common ground of most standard conversations. Hackers are 'a thing,' which means that hacker-kind does not typically need to be introduced into discourse. Semantically, these facts manifest themselves in the following way: the iota operator comes with an existence presupposition, which searches for a discourse referent satisfying the description 'the hacker production process'; if the context were somehow not to contain such a discourse referent, then the derivation would crash and the sentence would be infelicitous. In contexts where the relevant production process is not present in the common ground of the conversation, that is precisely what happens:

(9) ??People hacking into my website right now are from France.

There may very well be people hacking into my website right now, and many of them may very well be from France. Nonetheless, ordinary contexts do not contain the 'process that makes people who are hacking into Matt's website right now' discourse referent. In order to make sentences like (9) felicitous, that process will have to be introduced into discourse, perhaps via such fanciful scenarios as were contemplated in Chapter 2.

[9] Except in places where intensionality is absolutely crucial, as in definitions (10)–(13), this chapter will observe the convention of presenting the formal definitions of lexical items will in their extensional guises to enhance readability. Since *PM* is an intensional operator (and therefore has no extensional variant), these semi-derivations will stop at *PM*.

Of course, there are many questions as to whether (6) is the best way to formally implement these intuitive ideas. The analysis just broached locates the kind-forming operator in the plural affix *-s*. But why pursue that option? The three most salient alternatives are to make it into a type-shifting operation, in the manner of Chierchia (1998), Zamparelli (2000), and Cohen (2004), to locate it in some other operator on the syntactic tree, or to build it into the definition of *PM*. The first option is sufficiently similar to (6) that it will not be pursued here. Which of those two alternatives is preferable will depend on whether independent evidence comes down in its favor; if large-scale crosslinguistic considerations about how expansive the list of possible type-shifting rules should be (Partee, 1989) gives us reason to go with a type shifting approach, then the above proposal is straightforward to adapt. If additional independent syntactic evidence for either a second lexical entry for *-s* or an unpronounced operator that does the same thing is uncovered, then that gives us reason not to prefer the type shifting approach. Either way, (6) and a type-shifting variant thereupon are close philosophical cousins.

The second option is also sufficiently similar to (6) as not to warrant special treatment. One consideration in favor of placing the definition in (6) in an unpronounced operator is that generic statements about mass nouns such as snow or water seem to behave identically to bare plural generic statements. However, given that apart from their lack of a plural suffix, mass nouns take plural morphology in other respects (Lasersohn, 2011), there may be reason to posit an unpronounced plural affix in their case as well. Further syntactic evidence will rule one way or the other, and either option fits naturally with the approach recommended in this chapter.

By contrast, the remaining option, which would be to move the definition in (6) into the lexical entry for *PM*, does represent a substantially significant alternative, replete with different predictions. First, here is the definition of *PM* from Chapter 3:

(10) $[\![\textbf{P-M}]\!]^{w,g} = \lambda f_{est} \,.\, \lambda p \,.\, ...$
$... \forall w'(\textsc{Base}(w)(p)(w') \rightarrow \exists x(P(w')(x)(p) \wedge f(w')(x)))$

The alternative under consideration would do a number of things. First, if the kind-forming iota operator were moved into the lexical entry for *PM*, noun phrases would no longer be taken to denote kinds, and *PM* would effectively be made into a generalized quantifier, in the sense that it would be function of type $\langle\langle e, \langle s, t\rangle\rangle, \langle s, t\rangle\rangle$:

(11) $[\![\textbf{P-M}]\!]^{w,g} = \lambda f_{est} \,.\, \exists! p : \forall x(f(x)(w) \rightarrow P(x)(p)(w)) \wedge ...$
$... \forall w'(\textsc{Base}(w)(p)(w') \rightarrow \exists x(P(w')(x)(p) \wedge f(w')(x)))$

One immediate problem with this option is that it fails to predict generic sentences to have a cohesiveness presupposition. Generic sentences with haphazard NPs in subject position come out false, rather than anomalous. There is no semantic mechanism in this definition that can make it a felicity condition on any generic sentence that there exist a single pro-

duction process responsible for producing every object in the extension of the NP in subject position.

A better way to move the definition in (6) into into the lexical entry for *P-M* would be to make *P-M* into a quantifier:

(12) $[\![\textbf{P-M}]\!]^{w,g} = \lambda f_{est} \,.\, \lambda g_{est} \,.\, ...$
$... \forall w'(\textsc{Base}(w)(\iota p(\forall x(f(x) \leftrightarrow P(x)(p)))(w') \rightarrow ...$
$... \exists y(P(w')(y)(\iota p(\forall x(f(x) \leftrightarrow P(x)(p)(w'))) \wedge f(w')(y)))))$

This remains a reasonable option for those who have independent reservations about the analysis of the noun phrase in (6). It should be clear there is nothing at the level of the logical metalanguage that prevents this meaning from being expressible. It isn't as though the ideas put forth in this and the previous chapter are in principle inexpressible in quantificational form. Nonetheless, there is all-things-considered argument to be made that in view of the available empirical evidence, the analysis in (10) is preferable. After all, we still need some way of explaining the data from Chapter 2 regarding how bare pluaral noun phrases in subject position pattern gramamtically with name.

Perhaps, then, option (11) should be amended so as to replace the standard existential quantifier $\exists$ in (12) with a Strawsonian existential quantifier, $\exists_S$, which makes any formula in which it appears truth valueless if the open formula it operates on has a null extension:[10]

(13) $[\![\textbf{P-M}]\!]^{w,g} = \lambda f_{est} \,.\, \exists_S !p : \forall x(f(x)(w) \rightarrow P(x)(p)(w)) \wedge ...$
$... \forall w'(\textsc{Base}(w)(p)(w') \rightarrow \exists x(P(w')(x)(p) \wedge f(w')(x)))$

Option (13) can be used to account for the cohesiveness presupposition, but it raises a further nontrivial problem. As argued in section (50), in order to apply to generic comparatives, *PM* needs to be extensible to transitive predicates. The direct Schönfinkelization of *PM* contemplated in (50) will not work, because each time *PM* is applied, an additional production process will have to exist in order for the generic comparative to be felicitous. Repeating that option here:

(14) $PM(f_{\langle e,\langle e,t\rangle\rangle})(p)(q) \rightsquigarrow PM(\lambda y \,.\, PM(\lambda x \,.\, f(y)(x))(p))(q)$

For the following sentence to be true, then, there would have to be single a production process responsible for making all the tigers that had the property (a property which, it should be noted, is exceedingly awkward to state in English) of being such that lion-kind was characterized by being tougher than them:

(15) Lions are tougher than tigers.

[10] This is a close cousin of the presuppositional δ operator from Beaver (2001).

That is obviously a bad prediction. It is true that the alternative contemplated in section (50) would avoid the problem, because it would not be required to introduce an additional existence condition for each level of Schönfinkelization. Repeating the alternative transitive type-shifting rule here:

(16) $PM(f_{\langle e,\langle e,\langle s,t\rangle\rangle\rangle})(p)(q) \rightsquigarrow \lambda f_{e,est}\ .\ \lambda q\ .\ \lambda p\ .\ ...$
$... \forall w'(\textsc{Base}(w)(p)(w') \rightarrow \exists x, y(P(x)(p)(w') \wedge P(y)(q)(w') \wedge f(x)(y)(w')))$

As observed previously, this transitive rule has the advantage of correctly predicting relational generic statements not to exhibit scope ambiguity. But it has the disadvantage of being particular to this example, and an important desideratum on type-shifting rules is that they be as general as possible (Hendriks, 1988). Given that it is still up for negotiation how many varieties of type shifting there are, it is important for the kind theorist to keep her options open. Moving the existence condition into PM, then, has the undesirable effect of narrowing the available options for theorizing about generic sentences with bare plural noun phrases in object position, and is thus a possibility best left unexplored for the time being.

These are some of the reasons why the analysis of nouns in subject position of a generic as predicates which are lifted to descriptions via a kind-forming operator, as in (6), combined with the original definition of PM from Chapter 3, makes for a preferable choice. It will be assumed henceforth.

4.2 Artifactual Interpretations

The first contrast between generic statements and quantified statements, observed in Chapter 2, was that generic statements seemed to come along with a cohesiveness presupposition on the predicate in subject position. The previous section examined how the philosophical picture of kinds as production processes, along with the rendition of that picture into a formal semantics for PM, could account for that feature. This section will turn to the second such feature—generic sentences' susceptibility to artifactual interpretations—and show that this feature also lends itself to straightforward explanation via this semantic theory.

It is natural to make generic statements about biological species such as cats or whales, and equally natural to make generic statements about artifacts such as computers or cars. But some generic sentences allow for both interpretations, and interestingly, those interpretations have the potential to be mutually incompatible. Given a statement about, for example, some kind of animal, conflicting properties may seem to characterize that kind, depending on whether, in a given context, the conversational participants are regarding it

naturally or artifactually.[11] Here are some examples to illustrate the possibilities, extrapolated from Nickel (2008):

(17) a. Birds have feathers. *no artifactual interpretation; true biologically*
b. Birds have fur. *no artifactual interpretation; false biologically*

(18) a. Knives are sharp. *no biological interpretation; true artifactually*
b. Knives are dull. *no biological interpretation; false artifactually*

(19) a. Dobermans have pointy ears. *artifactually true; biologically false*
b. Dobermans have floppy ears. *biologically true; artifactually false*

(20) a. Every doberman has pointy ears. *simply true*
b. Every doberman has floppy ears. *simply false*
c. No doberman has pointy ears. *simply false*
d. No doberman has floppy ears. *simply true*

In what sense does sentence (19-a) have an *artifactual* interpretation? This is a convenient extension of the term *artifact* from the more familiar case of e.g. screwdrivers, overcoats, or lampposts. Clearly, this is not to say that dobermans literally are artifacts in the sense of something crafted out of inert, raw materials. But there is a sense of *dobermans* in which dobermans are indeed human creations.[12] One respect in which dobermans are artifactlike is that dog breeds are themselves the result of unnatural selection on the part of breeders. It is possible to invent a breed with something roughly like a predetermined set of phenotypic characteristics over just a few generations of selective breeding. More importantly, though, dobermans are bred in order to play a certain role in our society: the role of the guard dog.[13] In order to play that role, they are required to have good hearing, to appear threatening to potential intruders, and to be recognizable as dobermans.[14] Having their ears cropped at a young age helps with all three of these things. These are the facts that a truthful utterance of sentence (19-a) captures. Sentence (19-b), when truthfully uttered, captures the fact that dobermans are born with the phenotypic characteristic of having floppy ears. In contrast to Huskies or German Shepherds, they are not biologically disposed to grow pointy ears.

[11] Additionally, there may be more than one way of doing either.

[12] Searle (2009), chap. 5 provides a helpful framework for thinking about these cases. Roughly, the model is that we imbue certain artifacts with special social-institutional status by jointly agreeing to count them as special social-institutional objects. For instance, we jointly confer the status of 'thing with the power to purchase one single-origin macchiato' on a five dollar bill, by all agreeing to treat it as having that status. A similar story could apply to the archetype of the doberman.

[13] See Leslie (forthcoming) for an extraordinarily interesting account of how generic statements are used to describe different social roles that a single group of things (or people) can play.

[14] This third aim contributes to the second, insofar as dobermans have a reputation.

The next step is to begin thinking about how to model these facts. Chapter 2 argued that since quantified sentences do not give rise to artifactual interpretations, whatever is responsible for this element of context sensitivity must be something that appears in generic sentences but not quantified sentences. The theoretical picture of kinds as production process turns out to be particularly well-suited for explaining what this responsible party might be.

Before proceeding to that explanation, a quick detour into the problem of *qua*-objects is in order. What is a *qua*-object? Over the years, a number of contributors to the literature on metaphysics have been tempted to argue that there are cases in which an object can possess a property only under a certain description, rather than simply possessing the property, period.[15] Consider the following examples, adapted from Fine (1999) and Asher & Luo (2012):

(21) a. *Qua* physical acoustic pattern, the third movement of the Moonlight Sonata makes my silverware rattle.
b. *Qua* piece of music, the third movement of the Moonlight Sonata is in the key of C# minor.

(22) a. *Qua* lexical item, this is an adjective.
b. *Qua* graphic logo, this is in Helvetica.

(23) a. *Qua* legal act, Oswald's killing of Kennedy was murder.
b. *Qua* bodily movement, Oswald's assassination of Kennedy involved coiling his finger.

(24) a. *Qua* police officer, you've been suspended.
b. *Qua* undercover agent, you're my top operative in the field.

(25) a. *Qua* physical object, the math book was stolen.
b. *Qua* informational object, the math book is easily understood.

Some philosophers have wanted to argue that the above cases are examples of properties only holding of their bearers under a certain description. The driving intuition behind this view is that it only makes sense to apply the predicate *adjective* to a word insofar as it is a lexical item, not insofar as it is a set of written graphemes on a page. It only makes sense to apply the predicate *in C# minor* to a piece of music insofar as it is music, not insofar as it is a set of sound waves. And it would be strange to say that legally speaking, Kennedy's assassination was a finger-crooking. In the undercover scenario, the addressee is off duty under the

[15] See Anscombe (1957), Geach (1967), Wiggins (2002), Fine (1999), and Lewis (2003). For an ambitious attempt to remodel Montagovian semantics and type theory so as to accommodate these examples, see Asher (2011) and Asher & Luo (2012).

one description, but on-duty under the other. Finally, only physical objects can be stolen, only abstract informational objects can be learned, absorbed, mastered, or understood, and objects like books seem intelligible under either of these two guises.[16]

Wiggins (2002) famously argued against this principle, on the grounds that it leads to a rejection of the principle of substitutivity, which is a cornerstone of modern logic. The reason that allowing for *qua*-objects leads to a rejection of the substitutivity principle is as follows: suppose sentences (24-a) and (24-b) are both true. Then the police officer is identical to the undercover agent—this is part of what it means to say that those are different descriptions which pick out the same person—but they can differ in which properties they possess, thus violating the substitutivity principle. Surely, insists Wiggins, first-order logic with identity captures something deep about the human reasoning faculty, and there is simply no way to dispense with the principle of substitutivity without wreaking havoc on first-order logic with identity. So how could it be that the principle of substitutivity fails to play a comparably important role in our reasoning? For these reasons, the possibility of *qua*-objects (along with the rejection of the substitutivity principle) is standardly assumed to be a radical view, and is not accepted by the majority of contemporary philosophers.

When first confronted by examples such as those in (19), it is initially tempting to assimilate them to the cases we have been considering. Just as, from a certain point of view, there is no such thing as 'mastering' a book *qua* physical hunk of matter, it is likewise impossible for a doberman *qua* social-institutional artifact to be born with floppy ears, because the notion of social-institutional artifacts simply gets no grip in that case. Every doberman has two aspects: a biological aspect and an artifactual aspect. Certain properties hold of it under the former, and certain properties hold of it under the latter.

However, in spite of its initial appeal, the ultimate prospects for this option are weak. Suppose that a quantificational analysis of generic statements were somehow correct, despite the problem that examples (20) raise for that hypothesis, which were discussed in Chapter 2. Even then, explaining the artifactual interpretations in (19) would require adopting the view that one of those two sentences was generalizing over every doberman-*qua*-creature and the other was generalizing over every doberman-*qua*-artifact. But allowing that there can be *qua* objects, even if it has something intuitive going for it, as we just acknowledged, is controversial. There may indeed be reasons to revisit Peter Geach's relative identity thesis. Nonetheless, it speaks against the quantificational theory of generics, at least to a certain degree, when the only obvious way to implement it requires endorsing a radical philosophical thesis.

[16] Perhaps somewhat surprisingly, Asher (2011) argues that these guises are finite, enumerable, and part of the information we associate with many common nouns in their lexical entries.

How does the kind theory compare? Where a quantificational theory of the phenomenon in (19) is forced to choose one option—the option of allowing that each individual doberman has certain properties under certain aspects and other properties under others—the kind theory has two additional options to choose between. The first such option is to assume that doberman-kind itself (rather than any individual doberman) is a *qua*-object to which different properties can apply under different aspects. Of course, going that route would involve all the same difficulties as allowing that individual dobermans were *qua*-objects. But it would involve others as well, including the following. Suppose the biological doberman production process and the artifactual doberman process were one and the same. It would follow that making a biological doberman and making an artifactual doberman would be the same thing, and that is an undesirable prediction.[17] *Of course* it is possible to create a biological doberman without creating an artifactual doberman; that would seem to be the whole point of calling something a cultural product. Given any particular doberman, it could theoretically have grown up in the wild, holding onto all of its biological properties but lacking all of its artifactual properties.

The remaining option, which is the one this text will pursue, is to say that biological doberman-kind and artifactual doberman-kind are two separate kinds. They are separate, but related to one another in an important way. Biological doberman-kind is the evolutionary process responsible for making the organisms we refer to as dobermans, and artifactual doberman-kind is the process responsible for making those organisms *into* animals that have been trained in the doberman arts, so to speak. The question is what exact phenomenon this talk of *making one thing into* another might be tracking.

It is tempting to think of the biological kind and the artifactual kind as the same thing, because (nearly) all the dobermans that one is likely to encounter have the cultural characteristics of a doberman, in addition to the fact that dobermans were engineered as a breed by way of artificial selection. This means that any given doberman a person comes across will be both an artifactual doberman and a biological doberman. But even if every biological doberman is identical to every artifactual doberman, it does not follow that the two kinds are identical. A simple thought experiment should suffice to establish this. Imagine that every doberman puppy in the world was released onto an island and left to its own devices, while the remaining adult doberman population was allowed to die off. In that case, it would be safe to say that even though the artifactual doberman production process had met its demise, the biological production process continued to exist. This possibility, I would argue, is an indication that two processes can have the same progeny and yet be

[17] In terms of the model-theoretic semantics given in the previous chapter, this follows from the reflexivity of the parthood relation. If kind p is identical to kind q, then whenever p makes something, q has thereby also made something.

distinct. How? Such a scenario arises whenever one process makes its products exclusively out what another process created.

Luckily, the conception of kinds as production processes offers a convenient way to capture this notion. To say that one kind of thing p is made out of another kind of thing q is to say that p's progeny are made out of q's progeny. Or, to borrow some basic notions from the previous chapter, simply that p is a part of q, and that q feeds p. The following terminology will serve as shorthand for the idea. In the situation where one kind p is made out of another kind q, say that p comprises q (written $p \ni q$) just in case q feeds p and p is a part of q. Recall the definitions of parthood and feeding from Chapter 3, stated in English to reflect the fact that these are primitive relations in the type-logical metalanguage used by the semantics:

(26) **Feeding: $\mathbf{N(p,q)}$**
p feeds q iff p requires q to create something in order to itself create something

(27) **Parthood: $\mathbf{p \sqsubseteq q}$**
p is a part of q iff whenever p creates something, q has thereby created something

It is then possible to define the further notion of what it means for one process to comprise the other as follows:[18]

(28) **Comprisal: $\mathbf{p \ni q}$**
p comprises q iff $p \sqsubseteq q \wedge N(q,p)$

(29) **Same Comprisal Chain Relation: $\mathbf{p \approx q}$**
p is on the same comprisal chain as q iff $p \ni q \vee q \ni p$

An example to illustrate what it means for one kind to comprise another is in order. First, recall the examples invoked in Chapter 3 to illustrate parthood and feeding. Suppose there is a process of making desks, and a process of making screws, and that the desks are assembled using these screws. Then the screw production process feeds the desk production process, because the desks cannot go into production without the screws. The screw production process is not a part of the desk production process, however, because in making a desk one has not thereby made a screw. (Nor the other way around.) Next, suppose that there is a whale production process and a mammal production process. Then the whale production process is part of the mammal production process, because in making a whale it has thereby also made a mammal. But it does not feed the mammal production process, because whales are in no way required in order to make mammals. (Nor the other way around,

[18] Unlike the relations defined in (27) and (26), comprisal is not primitive, and thus is easily defined within the logic set forth in Chapter 3 as described here.

of course.) The mammal production process would continue to produce mammals without any problem, were all the whales in the world to vanish. Thus, parthood and feeding are logically independent.

Packaging these two notions into the above definition of comprisal yields interesting results. One process p comprises another process q just in case q feeds p and p is a part of q. This definition is useful not only for capturing what it means for one kind of thing to be made out of another kind of thing in general, but also for understanding statements about artifactual doberman-kind more specifically. Artifactual doberman-kind is characterized by different properties than biological doberman-kind, but it isn't a completely unrelated kind. Rather, the two bear a special connection: artifactual dobermans are *made out of* biological dobermans. They are made out of biological dobermans in the sense that in order to create an artifactual doberman, one raises a biological doberman as an artifactual doberman. Being made up of biological dobermans, according to the above definition, breaks up into two parts. Biological dobermans are necessary for the creation of artifactual dobermans: the former can exist without the latter, but not the other way around. So the biological doberman production process feeds the artifactual doberman production process. But in addition, in creating an artifactual doberman one has thereby also created a biological doberman—it isn't as though an animal ceases to be (biologically) a doberman when it is enculturated with the relevant set of traits. (Contrast this with the table example, where table-kind is fed by screw-kind even though tables are not screws.) So the artifactual doberman production process is also part of the biological doberman production process.

Chapter 2 made use, once again transitionally, of a subkind relation between certain pairs of kinds. This was prior to any particular proposal about what kinds might be. Just as the sortal-inspired account of kinds from Chapter 2 shall now be replaced with the conception of kinds as production processes from Chapter 3, the sortal-inspired account of the subkind relation shall now be replaced with a notion more suitable for explaining the linguistic data: namely, the comprisal relation. Where before we spoke of moving up and down subkind chains, we can now speak of moving up and down comprisal chains. For current purposes, a comprisal chain can be thought of as a set of kinds ordered under the comprisal relation. The distinctive features of sentences like (19) can be accounted for by positing a mechanism which, in any given context, allows a generic sentence to be reinterpreted such that it is about a kind either up or down in the comprisal chain from the kind explicitly mentioned, provided that that related kind is admitted by the common ground. Consider Nickel's example once again:

(30) Dobermans have floppy ears.

Rather than being interpreted as a statement about artifactual dobermans, it is interpreted as a statement about a kind admitted by the common ground that lies either up or down

the comprisal chain from artifactual dobermans—i.e. as a kind that either comprises or is comprised by artifactual dobermans. In the biological context, biological dobermans are admitted by the common ground, since it is generally known that dobermans have to be groomed into their proverbial role.[19] Biological dobermans are comprised (in this special sense of the term) by artifactual dobermans, and they do have floppy ears. Therefore, in this context, the sentence comes out true.

The next question is how to incorporate these intuitive ideas into the semantic theory put forward in (6). As always, there is an embarrassment of options, and as always, there will be an initial temptation to account for this phenomenon via some form of pragmatic coercion mechanism. However, this temptation is probably best resisted, given that the theory outlined in the two preceding chapters offers simple and elegant resources for accounting for artifactual interpretations directly within the semantics. The general motivation behind accounting for artifactual interpretations within the semantics is that generic statements systematically allow for them. The only contexts in which they do not are contexts in which there is no contextually salient artifactual kind that the kind explicitly mentioned comprises. Gricean strategies are a better fit for local, purpose-oriented context sensitivity than for broad, empirical regularities of this sort.

The new proposal will be to introduce a second iota operator into definition (6). Rather than mapping the nominal predicate to the production process responsible for creating everything in its extension, this revision of (6) maps the nominal predicate to *the contextually salient production process up or down in the comprisal chain* from the production process responsible for creating everything in its extension:

(31) $[\![\textbf{-s}]\!] = \lambda f_{e,t} \,.\, \iota q : q \approx \iota p(\forall x(f(x) \rightarrow P(p,x)))$

With this new apparatus, a simple explanation for the contrast in (19) is now ready to hand. In a conversational context that is geared towards morphological traits that were selected for, biological doberman-kind is the contextually salient kind in the same equivalence class as the kind responsible for making everything in the extension of the predicate *doberman*. Biological doberman-kind is indeed characterized by the property **floppy-eared**, so sentence (30) comes out true. In a conversational context that is geared towards establishing what it takes to perform the social role of the doberman well, artifactual doberman-kind is the contextually salient kind in the same equivalence class as the kind responsible for making everything in the extension of the predicate *doberman*.[20]

[19] In contexts where this fact about dobermans is not known, artifactual doberman-kind is not available to the common ground as a discourse referent, sentence (30) is thought to pertain to biological doberman-kind, and thus thought to be false.

[20] Though this is not the place to make the argument in full, this framework for understanding artifactual interpretations is also a promising way to explain why generic sentences with dual-character concepts in

In more pedestrian cases, the definition in (31) collapses back into the definition in (6), because every production process is on the same comprisal chain as itself. (This follows from the fact that the *same comprisal chain as* relation is an equivalence relation, which itself follows from the fact that the comprisal relation is the conjunction of a preorder and a partial order.) So:

(32) Bears are furry.

Sentence (32) is felicitous because there is a contextually salient kind in the same equivalence class as the kind responsible for making thing everything in the extension of the predicate *bear*—namely that kind itself, or bear-kind—and bear-kind is indeed characterized by furriness.

Thus, the availability of artifactual interpretations in generic sentences is readily explicable using the kind theory, with the added benefit of correctly predicting such interpretations not to arise in quantified sentences.

4.3 Contextual Domain Restriction

The third contrast between generic sentences and quantified sentences discussed in Chapter 2 pertained to the absence of domain restriction in generic sentences. This is the most troublesome contrast for the quantificational theory, given that domain restriction is a core feature of natural language quantification. The theory defended in this chapter correctly predicts that generic sentences do not contextually domain restrict, because they contain no quantificational determiners, and thus contain no domains to be restricted. The kind-forming operator from definition (31) does contain an iota operator which is context sensitive, but not in the same way as a quantifier. It checks for a contextually salient *kind*; not for a contextually salient subset of what was explicitly mentioned.

As mentioned in Chapter 2, a class of potential counterexamples to the claim that generic sentences do not contextually domain restrict, first observed in Condoravdi (1992, 1997), has also been discussed in the literature. To review the argument from Chapter 2, either these are generic sentences or they are not (as Condoravdi argued). If they are not generic sentences, then they are not counterexamples. And if they are generic sentences, there are still compelling reasons not to think they involve quantifier domain restriction. Quantifier domain restriction occurs by default, whenever the speaker utters a quantified sentence. When quantifiers allow for unrestricted interpretations, it is only because the

subject position give rise to deontic interpretations, as argued in Leslie (forthcoming). In these cases, the deontic interpretation arises when there is an artifactual interpretation of the human social group under discussion. This analysis of deontic generics will be pursued in future work.

context allows for it. Generics, by contrast, typically pertain to the kind picked out by the NP in subject position, because in the ordinary case, that is the kind that is most contextually salient. Artifactual interpretations are available only in circumstances where another kind on the same comprisal chain is contextually salient.

In the event that further evidence suggesting that Condoravdi-style counterexamples are indeed generic statements comes in, the apparatus from the previous section will be able to account for them in a simple and elegant way. They will be analyzed as cases in which a new kind is introduced into discourse via an accommodation procedure.[21] The fact that accommodation is a last-resort pragmatic repair mechanism offers the ideal explanation for why in the majority of cases, quantified sentences contextually domain restrict where generic statements do not, while still allowing that in certain special cases, a new kind can be retroactively introduced into discourse. To revisit the contrast, consider the first sentence in the pair from Chapter 2:

(33) Squirrels are friendly to people.

In ordinary contexts, this sentence is false, because the contextually salient production process responsible for making everything in the extension of the predicate *squirrel* (namely, biological squirrel-kind) is not characterized by friendliness. But Washington Square Park squirrel-kind, which biological squirrel-kind comprises, is characterized by friendliness. Similarly to artifactual dobermans, which are in a sense made out of biological dobermans, Washington Square Park squirrels are made out of biological squirrels. Washington Square Park squirrel-kind, as a fairly specific production process, is typically not contextually salient. But noting that the locale Washington Square Park exerts causal influence over its denizens, which could in principle lead a distinctive set of behavioral patterns to emerge there, makes it easier to retroactively introduce this related kind into discourse as a last-resort option. Thus—assuming these data check out—sentence (34) comes out true:

(34) Washington Square Park is quite a place. Squirrels are friendly to people.

This approach is a straightforward application of the general observations from the previous section about a distinctive, independently observable feature that generic sentences seem to have—that they give rise to artifactual interpretations. If indeed there are data here that require explanation, there are strong reasons to prefer this approach to an analysis that understands these cases as involving a generic quantifier which contextually restricts its domains.

[21] This is a standard move in the literature on the semantics-pragmatics interface. See Condoravdi & Gawron (1996); Aloni (2005).

4.4 Summary

Chapter 2 drew three distinctions between generic sentences and quantified sentences. Generic sentences come with a cohesiveness presupposition, are susceptible to artifactual interpretations, and do not contextually domain restrict. Quantified sentences have no semantic constraint on what predicate can appear in restrictor position, do not give rise to artifactual interpretations, and contextually domain restrict.

This chapter sets forth kind-theoretic analyses of the first two phenomena. As to the third, a kind-theoretic analysis predicts that generic sentences do not domain restrict, insofar as a sentence with no quantifier has no domain. Furthermore, should further investigation of the examples from Condoravdi reveal that the wrong prediction, this can easily be set right by combining the kind theory with a pragmatic accommodation mechanism. Quantificational theories offer no comparably straightforward way to capture the limited scope of these examples. Thus, even in the worst-case scenario, the kind theory does a better job of capturing the characteristic semantic features of generic sentences than do its competitors.

Chapter 5

Metaphysical Horizons

Having explored what it might take to give a semantic theory of generic sentences, we now return to the metaphysical themes of Chapter 1. There, it was argued that native speakers of a natural language can take on metaphysical commitments merely in virtue of their competence in that language. More specifically, the best available semantic theory for any given language can tell us what speakers of that language implicitly assume to be possible existents, what relations they assume to be possible among those existents, and what they assume to be the fundamental constraints on those relations. The best available semantic theory of a given language tells us, by way of the structure behind the models used to interpret the logic it is written in, what speakers of a language assume there to be, how the things they assume there to be can be configured, and what laws of necessity govern them.

Chapter 1 concluded by discussing the example of temporal adverbs like *twice*, and considered whether an event semantics was required in order to capture their truth conditions:

(1) Matt was punched twice.

Intuitively, sentence (1) is true both in the circumstance in which Matt is punched simultaneously in two different places, and in the circumstance in which Matt is at two different times in two different places. And, obviously, it is false in the circumstance in which he is punched only once at one time. A semantics that understands adverbs like *twice* as quantifying over times and evaluating propositions at those times, it was argued by Bach (1986), does not suffice to capture all three of those data points.[1]

The preceding chapters have argued that a kind-theoretic semantics for generic statements does the best job of accounting for three of their distinctive features. These features were presented in Chapter 2 as contrasts with quantified sentences, but are also helpfully

[1] See Chapter 1 for the full argument as to why.

thought of as distinctive features in their own right—as data points that any semantic theory of generic statements ought to capture. Assume that by now, we have at least nontrivial reason to think that such an approach is theoretically warranted. The next question, in light of the considerations adduced during Chapter 1, is whether these observations tell us anything of interest about what speakers of English implicitly presume kinds to be.

This concluding chapter will argue that they do. If the analysis put forth in the preceding chapters is correct, and if the philosophical interpretation of fully compositional semantic theories put forth in Chapter 1 is correct, then generic sentences give us reason to think that speakers of English are committed to the existence of kinds, at least in a certain sense of the term *kind*. The next step is to pin down exactly what notion of kinds speakers of English are committed to, assuming the truth of the above two premises.

5.1 What Are Kinds?

We saw in Chapter 1 that the philosophical issues here are fairly subtle. It would be going too far to claim that the search for the best semantic theory can reveal a commitment to any rich, substantive conception of what kinds are, of the sort that readers of e.g. Boyd (1991), Hacking (1991), or Rheins (2012) might be used to. Rather, whatever conception of kindhood the best available semantic theory of a language will reveal its speakers to be implicitly committed to is of necessity austere, formal, abstract, and mathematical. The best available semantic theory of temporal adverbs cannot tell us whether speakers of a language are committed to the existence of events, if by *events* we mean concrete, four-dimensional spacetime volumes.[2] But if by *event* we mean 'member of a set of entities ordered under an interval ordering,' or 'member of a set of entities in the domain of the relations denoted (respectively) by the predicates *agent*, *patient*, and *theme* in our logical metalanguage, which are subject to the following constraints' or something of that sort, then a semantic theory can reveal a commitment to events in that sense. It is less of a commitment than certain philosophical traditions may have hoped a semantic theory to be able to reveal, but it is also substantially more than nothing.

The intuitive starting point for the theory of kinds developed in Chapter 3 was the thought that kinds might helpfully be envisioned as processes which make things. A process that makes things is exactly the sort of thing that seems to correspond intuitively to an object to which properties modified by *PM* could apply. One reason is that to every property applicable to individual objects, there would seem to correspond a further property of

[2] Unless, of course, that semantic theory itself has some way of representing concreteness or four-dimensionality. But it won't be able to tell us that events are things which have these features, under our pre-theoretical conception thereof.

being a process which makes things that have the first property. Another reason is that it is fairly intuitive to imagine such a production process being interrupted, thus causing the things it makes to be created, occasionally, with various unpredictable abnormalities. This intuitive starting point then led to a story about three different primitive relations that can obtain among kinds and their members. Two production processes can stand in the feeding relation, which is what happens when the output of one production process is a necessary input to another. They can stand in the parthood relation, which is what happens just in case whenever one kind makes something, another kind has thereby also made something. And finally, all kinds bear the progenitor relation to their members, which is to say that every process is an efficient cause of what it makes, in something like the sense in which a parent is an efficient cause of her child.

This intuitive starting point enabled us to construct further derived notions that were useful in accounting for the data in Chapter 2—particularly the availability of artifactual interpretations, but also the possibility of accommodating the existence of new kinds into the common ground. One such derived notion was the *comprisal* relation on kinds, which captures the intuition that some kinds have artifactual uses. Horses are a biological species with biological features that were selected for, both naturally and unnaturally, but they are also an animal that plays a particular role in our society. Likewise for dobermans: as animals, dobermans have one set of species-sustaining functions, and as cultural products, they have another set of functions, which typically revolve around guarding privately-owned territory. The comprisal relation provides a means of spelling out what artifactual doberman-kind has to do with cultural doberman-kind. Artifactual dobermans are a kind of biological dobermans, given that whenever an artifactual doberman is created, a biological doberman has thereby also been created. In addition, there is the intuition that biological dobermans are *shaped into* artifactual dobermans through a regimen of diet, training, and bodily modification. Understanding artifactual dobermans and biological dobermans as standing in this relation is a way of spelling out this intuition: that artifactual dobermans are, in a sense, *made out of* biological dobermans. So we will say that artifactual doberman-kind *comprises* biological doberman-kind. More generally, one kind p comprises another kind q just in case q feeds p and p is a part of q.

This comprisal relation ended up being of some theoretical utility because it enabled us to define *comprisal chains*, and account for the availability of artifactual interpretations in a straightforward way directly within the semantics. The operator that maps predicates to kinds was rendered, very roughly, as a definite operator mapping the explicitly mentioned kind to the contextually salient kind up or down on the comprisal chain from it.

But these derived relations, according to the model of natural language metaphysics sketched out in Chapter 1, do not play a role in establishing a speaker's metaphysical commitments. If the semantic theory defended here were to provide the best explanation of

the behavior of generic statements in English, then the metaphysical commitments revealed thereby would be limited to the following:

(2) a. There are two atomic domains of objects (call them D_p and D_e).
b. There is an atomic domain of worlds (call it D_s).
c. There is a preorder $N \in D_p \times D_p$, a partial order $\sqsubseteq \in D_p \times D_p$, and an irreflexive relation $P \in D_p \times D_e$.
d. There is an accessibility function (call it BASE) that maps members of D_p (at a world of evaluation) to a set of ideal outcomes.

And that is all. According to the hypothesis under consideration, speakers of English are committed to the existence of kinds, in the relevant sense of the term *kind*. That is less of a commitment than one might think, because in the relevant sense of the term kind, there is no more to being a kind than inhabiting the above mathematical structure. But inhabiting the above mathematical structure is still something. Although the commitments an investigation into natural language semantics can reveal are more modest than one might have thought, they are far from trivial.

5.2 The Processual Perspective vs. The Standard View

Going back to Carlson (1977a), Schubert & Pelletier (1987), Chierchia (1998), and Zamparelli (2000), the standard assumption in natural language semantics about kinds has been that they ought to be formally modelled as mereological lattices constructed out of the individuals that belong to them. A natural question to ask, in the midst of this discussion of natural language metaphysics, is whether that semantic theory has an alternative metaphysical picture built into it.

The standard theory indeed presents a different metaphysical picture, but not an incompatible one. If it is correct, then it indicates a commitment on the part of native speakers of many languages (including English) to a different sort of metaphysical entity: namely, mereological sums with the characteristic algebraic features of a join semilattice. But the theory of kinds as plural lattices has a different set of ontological commitments that can be recovered from it, precisely because that theory is intended to model different linguistic phenomena. Carlson (1977a)'s original reason for construing kinds as mereological sums constructed out of individuals is that it is possible to make generic statements about kinds.[3] These cases are difficult to gather firm data on, but one can see where Carlson's intuitions are coming from by considering the following example:

(3) Zoo animals figure heavily in the world of AI databases.

[3] See Carlson (1977a), pp. 66-69, pp. 168-171, pp. 286-293.

Pinning down the exact truth conditions of sentence (3) is no simple task. Nonetheless, it is at least reasonably clear that sentence (3) is a statement about the kinds of animal that play a prominent role in the representation of commonsense knowledge: lions, bears, or alligators, as opposed to rotifers, tardigrades, or polychaetes. Not about individual animals. If that is correct, then for any generic, there is potentially a question as to whether it should be understood as bearing on individuals that fall under the kind or subkinds that fall under the kind:

(4) **Individual Interpretation:**
Big cats are cuddly.
a. ⇝ Shere Khan is presumably cuddly.
b. ⇝ Tigger is presumably cuddly.
c. etc.

(5) **Subkind Interpretation:**
Big cats are cuddly.
a. ⇝ Jaguars are presumably cuddly.
b. ⇝ Cougars are presumably cuddly.
c. etc.

Does the sentence 'Big cats are cuddly' give rise to both interpretations? If so, then envisioning kinds as mereological lattices provides a convenient way for our semantic theory to make this prediction, because this approach treats subkinds as bearing the same relation to the main kind as individuals falling under the kind bear to it. On the other hand, it is doubtful that formally rendering kinds as mereological lattices is the best way to make that prediction. One reason is that kind readings are not restricted to the generic construction. Indeed, the following pair of continuations shows that a kind-referential interpretation is available in principle wherever there is a common noun:

(6) (*pointing at a rat*) I love that rat!
a. ...Rusty is the most wonderful rat in the world.
b. ...The Norwegian rat is the most wonderful kind of rat.

The availability of these interpretations in non-generic constructions gives us strong reason to incorporate an account of this phenomenon into our semantics for nouns, rather than into our semantic theory of kinds. It could be treated as a kind of polysemy, on the lines proposed by Leslie (2013) for deontic/dual-character interpretations, or it could be part of the lexical semantics for common nouns.

The purpose of my semantics, as distinguished from either of these two alternatives, is to model the behavior of generic sentences as regards the availability of artifactual interpretations, the cohesiveness presupposition, and their lack of contextual domain restriction.

Combining the kinds as lattices model with the kinds as production processes model into a single theory, in principle, means accounting for all of the relevant phenomena, which is precisely the intent of modern linguistic theory. Model specific features of linguistic competence with detailed theories which are compatible with one another. Eventually, in the ideal situation, those detail-oriented theories can be combined into a large-scale theory that models a substantial portion of a native speaker's full linguistic competence. Each theory models a different fragment of English, with the many fragments overlapping as little as possible. If, at the idealized end of inquiry, we assemble a sufficiently rich set of fragments, the hope is that the result will look at least very close to actual English.

One straightforward way to combine the two approaches would be to code tokens of my atomic type p as mereological lattices. We are only beginning to understand the formal properties of logics that code simple types as functions or sets, such as Montague's EFL.[4] Rather than featuring atomic types for objects, worlds, and truth values, the only two atomic types in this logic are objects and intensions. And intensions are here coded in the traditional way as functions from possible worlds to truth values.[5] And so, a natural suggestion would be to do something similar for the logic laid out in Chapter 3: code production processes as objects with complex, logically visible mathematical structure, such as join semilattices. That way, the crosslinguistic morphological phenomena that Chierchia (1998) is interested in capturing could be accounted for by precisely the same mechanisms.

Another question that naturally arises in this connexion is: why have an atomic type for production processes in the first place? Why not just use a logic with an atomic domain of mereological lattices, as in Link (1983) or Landman (1991), and have the Base function map ideal outcomes to a subset of that domain? Here certain details from the picture laid out in Chapter 1 become important: in particular, the idea that the models used to interpret a logic can tell us something about the metaphysics that logic assumes—*up to isomorphism*.[6] Most crucially, if this alternative shows itself to be equivalent to my semantics, with kinds coded as plural lattices, plus the standard semantics for plurals, then I claim that both have the same metaphysical purport: they assume there to be objects with the structure outlined in (2). One can call such objects kinds, as I have been doing, or one can call them something else, if one prefers to reserve the term *kind* for some other type of entity. The terminology is

[4] Kristina Liefke has done some fascinating research on this pseudo-intensional framework, suggesting that it can model a surprising number of intensional phenomena in natural language, despite being extensional (at least according to the criterion that defines an intensional logic as one with an atomic type for possible worlds). See Liefke (2015).

[5] See Carnap (1947), Montague (1973).

[6] As mentioned in Chapter **??**, I suspect that we can make the condition even weaker: up to elementary equivalence. If that condition is correct, then dividing the set of model-theoretic structures into equivalence classes under the relation *same ontological commitments* will result in an even smaller set of equivalence classes. Though I am not prepared to defend that further claim here, I note its prima facie plausibility.

not particularly important; far more interesting are what, if anything, we are committed to merely in virtue of our competence in a language.

I think the more ontologically austere alternative under consideration *would* be equivalent to the semantics set forth here. This is because the more ontologically austere alternative would still have to recognize a distinction between arbitrarily formed plural lattices and cohesive plural lattices, in order to account for the cohesiveness presupposition. And ultimately, that would be no different from recognizing an atomic type for kinds. Whereas my semantics checks to see whether there is a contextually salient kind to which all members in the extension of the predicate in subject position belong (or one closely related thereto via the comprisal relation), this alternative semantics would check for a contextually salient plural lattice in the domain of the Base function. The fact that in this semantics, the Base function would have to be a partial function would be going proxy for the fact that in my semantics, there is an atomic type set aside for kinds. So even if one were to attempt to do without kinds as an atomic type, the fact that obvious alternatives would be equivalent suggests that in an important sense, that attempt would preserve the metaphysical commitments of the original analysis.

5.3 Lessons

Chapter 1 concluded by suggesting that the methods of natural language semantics in the Montagovian tradition had the potential to reopen the investigation into whether studying language can be a way of doing philosophy. I think that the case made here for a kind theory, if sound, indeed holds such promise for philosophy. There are many similar discussions elsewhere in semantics, with similar implications for philosophical method. For example, there is an interesting debate as to whether facts about adjectives commit speakers of a particular language to the existence of degrees, and if so, what degrees actually are.[7] On another front, Davidsonian (and especially neo-Davidsonian) semantics has provided strong reasons for thinking that the semantic behavior of verb phrases is most suitably captured using a logic with events as an atomic type.[8] The discussion carried out over the course of this text might helpfully be thought of as an intervention in a similar discussion.

As I hope to have demonstrated, revealing a native speaker's implicit metaphysical commitments takes hard work, and can only come about as part of an ongoing conversation among a large community of researchers. The payoff might seem a little less spectacular that

[7] For semantic theories that employ a logic with degrees as an atomic type, see Bartsch & Vennemann (1975), Stechow (1984), and Kennedy (2007). For the comparison class-based approaches that take degrees to be derived rather than entities, see Klein (1980) and Burnett (2015).

[8] See Davidson (1967), Parsons (1991), and Kratzer (2015).

what some philosophers are used to. As an abstract field whose theories employ the idiom of mathematical logic, model-theoretic semantics cannot reveal concrete commitments to specific propositions. Some might argue that those are the only commitments worth thinking about. But I take the position that that is a serious mistake. In general, philosophy is at its best when it keeps an open mind about the practical upshot of any given theoretical decision, and strives to avoid uncritically adopting the standard assumptions about what does or doesn't make a difference. Although theoretical questions suffer from being pursued solely with a view to particular narrow practical applications, it is a virtue of any great theoretical innovation that it opens the door to wide-ranging practical consequences in the future that couldn't have been foreseen.

A worry one might have about natural language metaphysics is that if it is correct, and certain metaphysical commitments could in principle be discovered to come not just with competence in a single language, but with competence in any language whatsoever, then it might follow that a person could be trapped with a particular set of metaphysical commitments, without having any choice in the matter. After all, opting out of competence in any language is not a viable course of action for anyone. Here, the abstractness of the metaphysical commitments that the best model-theoretic semantic theory can reveal is potentially useful. A metaphysical commitment such as the assumption that there is a set of objects ordered under an interval ordering may be precise, but it is general enough that it is difficult to imagine it being hugely constraining on a person's behavior.

Of course, this is not to suggest that such commitments make no difference to a person's life. Commitments are norms that we are forced to respond to in one way or another—we acknowledge them when we violate them and we acknowledge them when we behave in accordance with them. Were the search for the best available semantic theory to reveal a commitment of the sort under consideration among anyone who speaks any language, such a commitment could not but have an important effect on that person's life. However, the effect would be one of influencing their general orientation/perspective, rather than ruling out and ruling in some specific set of behaviors/beliefs. These features of our general orientation may not have a deterministic role to play in our lives, but that should not be taken to imply that their role is anything but central. The opportunity to subject them to careful, collaborative scrutiny is simply too promising for philosophy to pass over.

Bibliography

Abbott, Barbara. 1989. Nondescriptionality and Natural Kind Terms. *Linguistics and Philosophy*, 12(13), 269–291.

Almog, Joseph. 2011. Nature Without Essence. *Journal of Philosophy*, 107(7), 360–383.

Aloni, Maria. 2005. A Formal Treatment of the Pragmatics of Questions and Attitudes. *Linguistics and Philosophy*, 28(5), pp. 505–539.

Aloni, Maria, Kimmelman, Vadim, Roelofsen, Floris, Sassoon, Galit Weidman, Schulz, Katrin, & Westera, Matthijs (eds). 2012. *Logic, Language and Meaning - 18th Amsterdam Colloquium, Amsterdam , The Netherlands, December 19-21, 2011, Revised Selected Papers*. Lecture Notes in Computer Science, vol. 7218. Springer.

Annas, Julia. 1981. *An Introduction to Plato's Republic*. Vol. 103. Oxford University Press.

Anscombe, G.E.M. 1957. *Intention*. Cornell University Press.

Appiah, K. Anthony, & Gutmann, Amy. 1996. *Color Conscious: The Political Morality of Race*. Princeton University Press.

Aquinas, Thomas. 1947. *Summa Theologica*. Benzinger Brothers. Translated by Fathers of the English Dominican Province.

Aristotle. 1929. *Physics*. New York: G. P. Putnam.

Aristotle. 1975. *Categories and De Interpretatione*. Clarendon Aristotle series. Clarendon Press.

Armstrong, D.M. 1989. *A Combinatorial Theory of Possibility*. Cambridge Studies in Philosophy. Cambridge University Press.

Asher, N., & Luo, Z. 2012. Formalisation of coercions in lexical semantics. *In: Sinn und Bedeutung*, vol. 17.

Asher, Nicholas. 2011. *Lexical Meaning in Context: A Web of Words*. Lexical Meaning in Context: A Web of Words. Cambridge University Press.

Asher, Nicholas, & Morreau, Michael. 1995. What Some Generic Sentences Mean. *Pages 300–338 of:* Carlson, Gregory, & Pelletier, Francis Jeffrey (eds), *The Generic Book*. Chicago University Press.

Asher, Nicholas, & Pelletier, Francis Jeffrey. 1997. Generics and Defaults. *Pages 1125–1179 of:* van Benthem, Johan, & ter Meulen, Alice (eds), *Handbook of Logic and Language*. MIT Press.

Asudeh, Ash. 2004. *Resumption as resource management*. Ph.D. thesis, Stanford University.

Bach, Emmon. 1980. On Time, Tense, and Aspect: An Essay in English Metaphysics. *In:* Cole, Peter (ed), *Radical Pragmatics*. Academic Press, Inc.

Bach, Emmon. 1986. Natural Language Metaphysics. *Pages 573–595 of:* Weingartner, R. Barcan Marcus & G.J.W. Dorn & P. (ed), *Logic, Methodology, and Philosophy of Science VII*. Elsevier.

Bartsch, Renate, & Vennemann, Theo. 1975. The Grammar of Relative Adjectives and Comparison. *Pages 168–185 of: Formal Aspects of Cognitive Processes, Interdisciplinary Conference*. London, UK, UK: Springer-Verlag.

Bastiaanse, H. A., & Veltman, F. 2015. *Making the Right Exceptions*. Under Review.

Bastiaanse, Harald Andreas. 2014. The Intensional Many - Conservativity Reclaimed. *J. Philosophical Logic*, 43(5), 883–901.

Beaney, Michael (ed). 1997. *The Frege Reader*. Malden, Masachusetts: Blackwell Publishing.

Beaver, David I. 2001. *Presupposition & Assertion in Dynamic Semantics*. Stanford: CSLI Publications.

Beitchman, J. H., Cohen, N. J., Konstantareas, M.M., & Tannoch, R. (eds). 1996. *Language, Learning and Behaviour Disorders*. New York: Cambridge University Press.

Blutner, Reinhard. 1988. Defaults and the Meaning of Generic Sentences. *In: Genericity in Natural Language: Proceedings of the 1988 Tübingen Conference*.

Borg, Emma. 2004. *Minimal Semantics*. Oxford University Press.

Borg, Emma. 2012. *Pursuing Meaning*. Oxford University Press.

Boyd, Richard. 1991. Realism, Anti-Foundationalism and the Enthusiasm for Natural Kinds. *Philosophical studies*, 61(1), 127–148.

Brandom, Robert. 1994. *Making it Explicit: Reasoning, Representing, and Discursive Commitment*. Harvard University Press.

Burge, Tyler. 1973. Reference and Proper Names. *Journal of Philosophy*, 70(14), 425–439.

Burnett, Heather. 2015. *Gradability in Natural Language: Logical and Grammatical Foundations*. Oxford Studies in Semantics and Pragmatics: Oxford University Press.

Caplan, Rochelle. 1996. Discourse deficits in childhood schizophrenia. *In:* Beitchman *et al.* (1996).

Cappelen, Herman, & Lepore, Ernie. 2005. *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Blackwell Publishing.

Cariani, Fabrizio. 2013. 'Ought' and Resolution Semantics*. *Noûs*, 47(3), 534–558.

Carlson, Gregory N. 1977a. *Reference to Kinds in English*. Ph.D. thesis, University of Massachusetts, Amherst.

Carlson, Gregory N. 1977b. A Unified Analysis of the English Bare Plural. *Linguistics and Philosophy*, 1(3), 413–458.

Carlson, Gregory N. 1982. Generic Terms and Generic Sentences. *Journal of Philosophical Logic*, 11(2), 145–181.

Carlson, Gregory N. 1989. The Semantic Composition of English Generic Sentences. *In:* Chierchia, Gennaro, Partee, Barbara, & Turner, Raymond (eds), *Property Theory, Type Theory, and Semantics*. D. Reidel Publishing.

Carnap, Rudolf. 1947. *Meaning and Necessity: A Study in Semantics and Modal Logic*. Meaning and Necessity: A Study in Semantics and Modal Logic, no. v. 3. University of Chicago Press.

Chierchia, Gennaro. 1998. Reference to Kinds Across Languages. *Natural Language Semantics*, 6(4), 339–405.

Chomsky, Noam. 1995. *The Minimalist Program*. Current studies in linguistics series. MIT Press.

Cohen, Ariel. 1997. Default Reasoning and Generics. *Computational Intelligence*, 13(4), 506–533.

Cohen, Ariel. 1999a. Generics, Frequency Adverbs, and Probability. *Linguistics and Philosophy*, 22, 221–253.

Cohen, Ariel. 1999b. *Think generic! The meaning and use of generic sentences.* Center for the Study of Language and Information.

Cohen, Ariel. 2001. On the Generic Use of Indefinite Singulars. *Journal of Semantics*, 18(3), 183–209.

Cohen, Ariel. 2004. Existential Generics. *Linguistics and Philosophy*, 27(2), 137–168.

Cohen, Ariel. 2013. Genericity. *In:* Aloni, Maria, & Dekker, Paul (eds), *Cambridge Handbook of Semantics*. Camb.

Condoravdi, Cleo. 1992. Strong and Weak Novelty and Familiarity. *In:* Barker, Chris, & Dowty, David (eds), *Proceedings of SALT*. Cornell.

Condoravdi, Cleo. 1997. *Descriptions in Context*. Outstanding Dissertations in Linguistics. Garland Publishing, Incorporated.

Condoravdi, Cleo, & Gawron, Jean Mark. 1996. The Context-Dependency of Implicit Arguments. *Pages 1–32 of:* Kanazawa, Makoto, de Swart, Henri ette, & Piñon, Christopher (eds), *Quantifiers, Deduction, and Context*. Stanford: CSLI Publications.

Conradie, W., Ghilardi, S., & Palmigiano, A. 2014. Unified Correspondence. *In: Johan van Benthem on Logic and Information Dynamics*. Outstanding Contributions to Logic. Springer.

Dalrymple, Mary. 2001. *Lexical Functional Grammar*. Lexical functional grammar, no. v. 34. Brill Academic Pub.

Davidson, Donald. 1967. The Logical Form of Action Sentences. *In: The Logic of Decision and Action*. University of Pittsburgh Press.

Dekker, P.J.E. 2012. *Dynamic Semantics*. Studies in linguistics and philosophy. Springer London, Limited.

Deo, Ashwini. 2009. Unifying the imperfective and the progressive: partitions as quantificational domains. *Linguistics and Philosophy*, 32(5), 475–521.

Dobrovie-Sorin, C., & Pires de Oliveira, Roberta. 2008. Reference to kinds in Brazilian Portuguese. *Pages 107–121 of:* Gr, Atle (ed), *Proceedings of SuB 12*University of Oslo, for Department of Literature, Area Studies and European Languages.

Elbourne, P.D. 2005. *Situations And Individuals*. Current Studies in Linguistics Series. MIT Press.

Fara, Delia Graff. 2001. Descriptions as Predicates. *Philosophical Studies*, 102, 1–42.

Fara, Delia Graff. forthcoming. Names Are Predicates. *Philosophical Review*.

Farkas, Donka F., & Sugioka, Yoko. 1983. Restrictive If/When Clauses. *Linguistics and Philosophy*, 6(2), 225–258.

Fine, Gail. 1978. Knowledge and Belief in Republic V. *Archiv für Geschichte der Philosophie*, 60(2), 121–39.

Fine, Kit. 1999. Things and Their Parts. *Midwest Studies in Philosophy*, 23(1), 61–74.

Frege, Gottlob. 1951. On Concept and Object. *Mind*, 60, 168–180.

Gallin, Daniel. 1975. *Intensional and Higher-order Modal Logic*. North Holland Mathematics Studies, vol. 19. Elsevier.

Gamut, L. T. F. 1991. *Logic, Language, and Meaning, Volume 2: Intensional Logic and Logical Grammar*. University of Chicago Press.

Geach, P. T. 1961. Aquinas. *In:* Geach, P.T. & Anscombe, G. E. M. (ed), *Three Philosophers*. Basil Blackwell.

Geach, P. T. 1967. Identity. *The Review of Metaphysics*, 21(1), 3–12.

Geach, P.T. 1980. *Reference and Generality: An Examination of Some Medieval and Modern Theories*. Contemporary philosophy series. Cornell University Press.

Geurts, Bart. 1988. The Representation of Generic Knowledge. *In:* Krifka, Manfred (ed), *Genericity in Natural Language: Proceedings of the 1988 Tübingen Conference.*

Geurts, Bart. 1997. Good News about the Description Theory of Names. *Journal of Semantics*, 14(4), 319–348.

Glanzberg, Michael. 2007. Definite Descriptions and Quantifier Scope: Some Mates Cases Reconsidered. *European Journal of Analytic Philosophy*, 3, 133–158.

Gonzalez, Francisco J. 1996. Propositions or Objects? A Critique of Gail Fine on Knowledge and Belief in "Republic" V. *Phronesis*, 41(3), pp. 245–275.

Gopnik, M., & Crago, Martha B. 1991. Familial aggregation of a developmental language disorder. *Cognition*, 39(1), 1 – 50.

Gray, Aidan. 2012. *Names and Name-Bearing: An Essay on the Predicate View of Names*. Ph.D. thesis, The University of Chicago.

Greenberg, Yael. 2003. *Manifestations of Genericity*. Outstanding Dissertations in Linguistics Series. Taylor & Francis.

Greenberg, Yael. 2007. Exceptions to Generics: Where Vagueness, Context Dependence and Modality Interact. *Journal of Semantics*, 24(2), 131–167.

Groenendijk, Jeroen, & Stokhof, Martin. 1991. Dynamic predicate logic. *Linguistics and Philosophy*, 14(1), 39–100.

Gross, Steven. 2006. Can Empirical Theories of Semantic Competence Really Help Limn the Structure of Reality? *Noûs*, 40(1), 43–81.

Gupta, Anil. 1980. *The Logic of Common Nouns: An Investigation in Quantified Modal Logic*. Yale University Press.

Hacking, Ian. 1991. A tradition of natural kinds. *Philosophical Studies*, 61(1-2), 109–126.

Hansen, Nathaniel. 2011. Color Adjectives and Radical Contextualism. *Linguistics and Philosophy*, 34(3), 201–221.

Harel, David. 1984. Dynamic Logic. *Pages 497–604 of:* Gabbay, D., & Guenther, F. (eds), *Handbook of Philosophical Logic: Extensions of Classical Logic*, vol. 2. D. Reidel Publishing.

Harris, Zellig. 1954. Distributional structure. *Word*, 10(23), 146–162.

Hawthorne, James. 2011. Confirmation Theory. *In:* Bandyopadhyay, Prasanta S., & Forster, Malcolm (eds), *Philosophy of Statistics, Handbook of the Philosophy of Science, Volume 7*. Elsevier.

Heim, Irene. 1982. *On the Semantics of Definite and Indefinite Noun Phrases*. Ph.D. thesis, University of Massachusetts at Amherst.

Heim, Irene, & Kratzer, Angelika. 1998. *Semantics in Generative Grammar*. Wiley-Blackwell.

Hendriks, Herman. 1988. Type Change in Semantics: The Scope of Quantification and Coordination. *In:* van Benthem, Johan, & Klein, Ewan (eds), *Categories, Polymorphism, and Unification*. ITLI.

Kant, Immanuel. 1998. *Critique of Pure Reason*. Cambridge University Press. Translated by Paul Guyer and Allen Wood.

Kennedy, Christopher. 2007. Vagueness and grammar: the semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy*, 30(1), 1–45.

Kennedy, Christopher, & McNally, Louise. 2010. Color, Context, and Compositionality. *Synthese*, 174(1), 79–98.

Klecha, Peter. 2011. Optional and obligatory modal subordination. *In:* Ingo Reich, Eva Horch, Dennis Pauly (ed), *Proceedings of Sinn & Bedeutung 15*. Saarbrücken, Germany: Universaar – Saarland University Press.

Klein, Ewan. 1980. A Semantics for Positive and Comparative Adjectives. *Linguistics and Philosophy*, 4, 1–45.

Kratzer, Angelika. 1977. What 'must' and 'can' must and can mean. *Linguistics and Philosophy*, 1(3), 337–355.

Kratzer, Angelika. 1981. The notional category of modality. *Pages 38–74 of:* Eikmeyer, H.-J., & Rieser, H. (eds), *Words, Worlds, and Context*. de Gruyter.

Kratzer, Angelika. 1995. Stage-Level and Individual-Level Predicates. *Pages 125–175 of:* Carlson, Gregory, & Pelletier, Francis Jeffrey (eds), *The Generic Book*. Chicago University Press.

Kratzer, Angelika. 2015. *The Event Argument and the Semantics of Verbs*.

Krifka, M. 2004. Bare NPs: Kind-referring, Indefinites, Both, or Neither? *In: Proceedings of Semantics and Linguistic Theory (SALT) XIII*.

Krifka, Manfred. 1988. The Relational Theory of Genericity. *In: Genericity in Natural Language: Proceedings of the 1988 Tübingen Conference*.

Krifka, Manfred, Pelletier, Francis Jeffry, Carlson, Gregory, ter Meulen, Alice, Chierchia, Gennaro, & Link, Godehard. 1995. Genericity: An Introduction. *In:* Gregory Carlson, Francis Jeffry Pelletier (ed), *The Generic Book*. The University of Chicago Press.

Kripke, Saul. 1980. *Naming and Necessity*. Harvard University Press.

Kripke, Saul A. 1963a. Semantical Analysis of Modal Logic I. Normal Propositional Calculi. *Zeitschrift fur mathematische Logik und Grundlagen der Mathematik*, 9(56), 67–96.

Kripke, Saul A. 1963b. Semantical Considerations on Modal Logic. *Acta Philosophica Fennica*, 16(1963), 83–94.

Lambek, Joachim. 1958. The Mathematics of Sentence Structure. *The American Mathematical Monthly*, 65(3), pp. 154–170.

Landman, Fred. 1991. *Structures for Semantics*. Studies in Linguistics and Philosophy. Springer.

Langdon, Robyn, Davies, Martin, & Coltheart, Max. 2002. Understanding Minds and Understanding Communicated Meanings in Schizophrenia. *Mind & Language*, 17(1-2), 68–104.

Lasersohn, Peter. 2005. Context Dependence, Disagreement, and Predicates of Personal Taste*. *Linguistics and Philosophy*, 28(6), 643–686.

Lasersohn, Peter. 2011. Mass Nouns and Plurals. *In:* Maienborn, Claudia, Portner, Paul, & von Heusinger, Klaus (eds), *Semantics: An International Handbook of Natural Language Meaning*, vol. 2. De Gruyter Mouton.

Leslie, Sarah-Jane. 2007. Generics and the Structure of the Mind. *Philosophical Perspectives*, 21(1), 375–403.

Leslie, Sarah-Jane. 2008. Generics: Cognition and Acquisition. *Philosophical Review*, 117(1), 1–47.

Leslie, Sarah-Jane. 2013. Generics Oversimplified. *Noûs*.

Leslie, Sarah-Jane. forthcoming. 'Hillary Clinton is the Only Man in the Obama Administration': Dual Character Concepts, Generics, and Gender. *Analytic Philosophy*.

Lewis, David. 1973. *Counterfactuals*. Oxford: Blackwell Publishers.

Lewis, David. 1979. Scorekeeping in a Language Game. *Journal of Philosophical Logic*, 8(1), 339–359.

Lewis, David. 2003. Things Qua Truthmakers. *Pages 25–38 of: Real Metaphysics: Essays in honor of D. H. Mellor*. Routledge.

Liebesman, David. 2011. Simple Generics. *Noûs*, 45(3), 409–442.

Liefke, Kristina. 2015. Codability and Robustness in Formal Natural Language Semantics. *In:* Mineshima, Koji, Bekki, Daisuke, & McCready, Eric (eds), *Selected papers of JSAI-isAI 2014 Workshops*. New York: Springer.

Link, Godehard. 1983. The Logical Analysis of Plurals and Mass Terms: A Lattice-Theoretic Approach. *Pages 127–147 of:* Portner, P., & Partee, B. H. (eds), *Formal Semantics - the Essential Readings*. Blackwell.

Lowe, E.J. 2009. *More Kinds of Being: A Further Study of Individuation, Identity, and the Logic of Sortal Terms*. Wiley.

Luhtala, Anneli. 2005. *Grammar and Philosophy in Late Antiquity: A Study of Priscian's Sources*. John Benjamins Publishing Company.

Malink, Marko. 2006. A Reconstruction of Aristotle's Modal Syllogistic. *History and Philosophy of Logic*, 27, 95–141.

Malink, Marko. 2013. *Aristotle's Modal Syllogistic*. Harvard University Press.

May, Robert. 1977. *The Grammar of Quantification*. Ph.D. thesis, Massachusetts Institute of Technology.

McCarthy, John. 1980. Circumscription: A Form of Non-Monotonic Reasoning. *Artificial Intelligence*, 13, 27–39.

McNally, Louise. *Kinds, descriptions of kinds, concepts, and distributions*. Submitted to a volume of papers from the 2014 BRIDGE-14 Workshop in Düsseldorf.

Milsark, Gary. 1974. *Existential Sentences in English*. Ph.D. thesis, Massachusetts Institute of Technology.

Montague, Richard. 1973. The Proper Treatment of Quantification in Ordinary English. *In:* Hintikka, J., Moravcsik, J. M. E., & Suppes, P. (eds), *Approaches to Natural Language*. Synthese Library, vol. 49. Dordrecht: Reidel.

Moore, G.E. 1903. *Principia Ethica*. Cambridge University Press.

Morreau, Michael. 1988. Default Formalisms for Generics: A Consumer Report. *In: Genericity in Natural Language: Proceedings of the 1988 Tübingen Conference.*

Mourelatos, Alexander. 1978. Events, Processes, and States. *Linguistics and Philosophy*, 2, 415–434.

Muskens, Reinhard. 2011. A Theory of Names and True Intensionality. *In:* Aloni *et al.* (2012).

Nickel, Bernhard. 2008. Generics and the Ways of Normality. *Linguistics and Philosophy*, 31(6), 629–648.

Nickel, Bernhard. 2010a. Generic Comparisons. *Journal of Semantics*, 27(2), 207–242.

Nickel, Bernhard. 2010b. Generically Free Choice. *Linguistics and Philosophy*, 33(6), 479–512.

Nickel, Bernhard. forthcoming. *Between logic and the world: an integrated theory of generics*. Oxford University Press, USA.

Novaes, Catarina Dutilh. 2014. *Formal Languages in Logic: A Philosophical and Cognitive Analysis*. Cambridge University Press.

Oosterhof, A. 2008. *The Semantics of Generics in Dutch and Related Languages*. Linguistics Today. John Benjamins Pub.

Parsons, Terence. 1991. *Events in the Semantics of English: A Study in Subatomic Semantics*. Cambridge, MA: MIT Press.

Partee, Barbara. 1989. Bidning Implicit Arguments in Quantified Contexts. *CLS*, 25, 342–265.

Peirce, Charles Sanders. 1940. Abduction and Induction. *In:* Buchler, Justus (ed), *Philosophical Writings of Peirce*. Dover Publications.

Pelletier, F.J. 2009. *Kinds, Things, and Stuff: Mass Terms and Generics*. New directions in cognitive science. Oxford University Press, USA.

Plato. 2002. *Five Dialogues: Euthyphro, Apology, Crito, Meno, Phaedo*. Hackett Publishing.

Pollard, Carl, & Sag, Ivan A. 1994. *Head-Driven Phrase Structure Grammar*. Chicago: The University of Chicago Press.

Pratt, V. P. 1979. Models of Program Logics. *Pages 115–122 of: Proc. 20th IEEE Conference on Foundations of Computer Science*.

Prinz, Jesse. 2006. Is The Mind Really Modular? *In:* Stainton, Robert J. (ed), *Contemporary Debates in Cognitive Science*. Malden: Blackwell Publishing.

Putnam, Hilary. 1962. It Ain't Necessarily So. *Journal of Philosophy*, 59(22), 658–671.

Putnam, Hilary. 1970. Is Semantics Possible? *Metaphilosophy*, 1(3), 187–201.

Putnam, Hilary. 1975. The Meaning of 'Meaning'. *Minnesota Studies in the Philosophy of Science*, 7, 131–193.

Quer, Josep. 2011. Quantificational Strategies across Language Modalities. *In:* Aloni *et al.* (2012).

Quine, Willard V. 1948. On What There Is. *The Review of Metaphysics*, 2(5), pp. 21–38.

Restall, Greg. 2012. A cut-free sequent system for two-dimensional modal logic, and why it matters. *Annals of Pure and Applied Logic*, 163(11), 1611 – 1623. Kurt Goedel Research Prize Fellowships 2010.

Rheins, Jason. 2012. Similarity and Species Concepts. *In: Carving Nature at its Joints: Topics in Contemporary Philosophy, Vol. 8*. MIT.

Rini, Adriane A. 1998. Is There a Modal Syllogistic? *Notre Dame Journal of Formal Logic*, 39(4), 554–572.

Rini, Adriane A. 2011. *Aristotle's Modal Proofs*. Springer Netherlands.

Roelofsen, Floris. 2011. Algebraic Foundations for Inquisitive Semantics. *Pages 233–243 of:* van Ditmarsch, Hans, Lang, Jérôme, & Ju, Shier (eds), *Logic, Rationality, and Interaction*. Lecture Notes in Computer Science, vol. 6953. Springer Berlin Heidelberg.

Rooth, Mats. 1992. A theory of focus interpretation. *Natural language semantics*, 1(1), 75–116.

Rosen, Carol. 1984. The Interface Between Semantic Roles and Their Initial Grammatical Relations. *In:* Perlmutter, David M., & Rosen, Carol (eds), *Studies in Relational Grammar*, vol. 2. The University of Chicago Press.

Rossen, M., Klima, E.S., U., Bellugi, A., Bihrle, & James, W. 1996. Interaction between language and cognition: evidence from Williams syndrome. *In:* Beitchman *et al.* (1996).

Rothstein, Susan. 2008. *Structuring Events: A Study in the Semantics of Lexical Aspect*. Wiley-Blackwell.

Russell, Bertrand. 1905. On Denoting. *Mind*, 14(56), 479–493.

Sag, Ivan A. 1976. *Deletion and Logical Form*. Ph.D. thesis, Massachusetts Institute of Technology.

Sassoon, Galit Weidman. 2009. Semantics with Clusters of Properties. *In:* Kissine, Mikhail, & De Brabanter, Philippe (eds), *Utterance Interpretation and Cognitive models*. CRISPI: Emerald Book Publishing. Ben Gurion University of the Negev.

Schubert, Lenhart K., & Pelletier, Francis Jeffry. 1987. Problems in Representing the Logical Form of Generics, Bare Plurals, and Mass Terms. *Pages 387–453 of:* Lepore, Ernest (ed), *New Directions in Semantics*. Academic Press.

Searle, J. 2009. *Making the Social World: The Structure of Human Civilization*. Oxford University Press.

Snedegar, Justin. 2013. Reason claims and contrastivism about reasons. *Philosophical Studies*, 166(2), 231–242.

Stalnaker, Robert C. 1968. A Theory of Conditionals. *Pages 98–112 of: Studies in Logical Theory*. Blackwell.

Stalnaker, Robert C. 1970. Pragmatics. *Synthese*, 22(1-2), 272–289.

Stalnaker, Robert C. 1976. Possible Worlds. *Noûs*, 10(1), 65–75.

Stanley, Jason, & Szabó, Zoltán Gendler. 2000. On Quantifier Domain Restriction. *Mind & Language*, 15(2/3), 219–262.

Stechow, Arnim von. 1984. Comparing Semantic Theories of Comparison. *Journal of Semantics*, 3(1-2), 1–77.

Sterken, Rachel. 2015a. *Generics in Context*. Forthcoming in *Philosophers' Imprint*.

Sterken, Rachel. 2015b. *Generics, Kind Predication, and Context*. Under Revision.

Stokhof, Martin. 2013. Formal Semantics and Wittgenstein. *The Monist*, 96(2), 205–231.

Stokhof, Martin, & Van Lambalgen, Michael. 2011. Abstractions and idealisations: The construction of modern linguistics. *Theoretical Linguistics*, 37(1-2), 1–26.

Strawson, P. F. 1959. *Individuals: An Essay in Descriptive Metaphysics*. Routledge.

Stump, G. 1985. *The Semantic Variability of Absolute Constructions*. Studies in Linguistics and Philosophy. Springer Netherlands.

Thompson, Michael. 2004. Apprehending Human Form. *Royal Institute of Philosophy Supplement*, 54, 47–74.

Thompson, Michael. 2009. *Life and Action: Elementary Structures of Practice and Practical Thought*. Harvard University Press.

van Benthem, J. 2010. *Modal Logic for Open Minds*. CSLI lecture notes. Center for the Study of Language and Information.

van Benthem, Johan. 1984. Correspondence Theory. *Pages 167–247 of:* Gabbay, D., & Guenthner, F. (eds), *Handbook of Philosophical Logic*. Synthese Library, vol. 165. Springer Netherlands.

Varzi, Achille C. 2002. Words and Objects. *Pages 49–75 of:* &, A. Bottani & M. Carrara, & Giaretta, D. (eds), *Individuals, Essence, and Identity: Themes of Analytic Metaphysics*. Kluwer Academic Publishers.

Veltman, Frank. 1996. Defaults in Update Semantics. *Journal of Philosophical Logic*, 25(3), 221–261.

Vendler, Zeno. 1957. Verbs and Times. *The Philosophical Review*, 66(2), 143–160.

von Fintel, Kai. 1994. *Restrictions on Quantifier Domains*. Ph.D. thesis, MIT.

Warmke, Craig. forthcoming. Modal Intensionalism. *Journal of Philosophy*.

Wiggins, David. 2002. *Sameness and Substance Renewed*. Cambridge University Press.

Williamson, Timothy. 2012. Boghossian and Casalegno on Understanding and Inference. *Dialectica*, 66(2), 237–247.

Zamparelli, Roberto. 2000. *Layers in the Determiner Phrase*. Outstanding dissertations in linguistics. Garland Publishing, Incorporated.