GENERICS IN CONTEXT: GENERALISATION, CONTEXT AND COMMUNICATION

Rachel Sterken

A Thesis Submitted for the Degree of PhD
at the
University of St Andrews

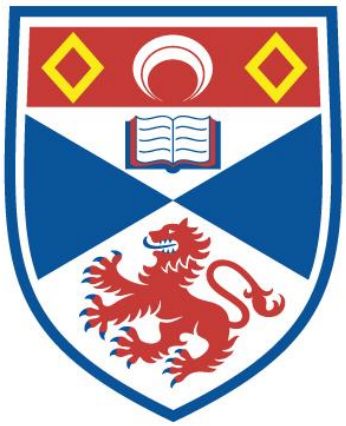

2014

Generics in Context:

Generalisation, Context and Communication

Rachel Sterken

A thesis to be submitted to
the University of Oslo and the University of St Andrews
for the degree of
Doctor of Philosophy

Department of Philosophy, Classics, History of Art and Ideas
University of Oslo

Departments of Philosophy
School of Philosophical, Anthropological and Film Studies
University of St Andrews

December 2012

I, Rachel Sterken, hereby certify that this thesis, which is approximately 60,000 words in length, has been written by me, that it is the record of work carried out by me and that it has not been submitted in any previous application for a higher degree.

I was admitted as a research student in XXX and as a candidate for the degree of Doctor of Philosophy in 2007; the higher study for which this is a record was carried out in the University of St Andrews between September 2007 to September 2010, and September 2011 to December 2012.

18 December 2012

Signature of candidate

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of Doctor of Philosophy in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree.

18 December 2012

Signature of supervisor

In submitting this thesis to the University of St Andrews we understand that we are giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. We also understand that the title and the abstract will be published, and that a copy of the work may be made and supplied to any bona fide library or research worker, that my thesis will be electronically accessible for personal or research use unless exempt by award of an embargo as requested below, and that the library has the right to migrate my thesis
into new electronic forms as required to ensure continued access to the thesis. We have obtained any third-party copyright permissions that may be required in order to allow such access and migration, or have requested the appropriate embargo below.

The following is an agreed request by candidate and supervisor regarding the electronic publication of this thesis:

Access to printed copy and electronic publication of thesis through the University of St Andrews.

18 December 2012

Signature of supervisor Signature of candidate

Abstract

This thesis consists of four chapters and an introduction.

The first chapter is concerned with cases of purported genericity which are true despite only a minority of the kind in question satisfying the predicated property and whose predicated property is somehow *striking*. I argue that such cases are pathological, and hence shouldn't be used as evidence for claims about the nature of genericity or the semantics of generics. In particular, these cases pose no difficulty for theorists who want to solve the problem of variability (the main obstacle in providing an adequate theory of generics), at least in part, by appeal to the linguistic property of quantifiers known as (contextual) domain restriction.

The second chapter is devoted to defending a quantificational analysis of generics in which the logical form of generics contains an unpronounced quantifier expression *Gen*. The chapter defends the original argument of Carlson (1989), the *multiple readings argument*, which has recently come under scrutiny. In addition, it provides a novel argument in favour of *Gen*. This chapter also responds to various objections to the quantificational analysis, in particular to the linguistic evidence offered in favour of the kind-predication analysis.

The third chapter defends the positive view that the unpronounced quantifier expression *Gen* is an indexical. The chapter argues that a given generic sentence expresses different generalisations in different contexts of utterance. A semantics for generics which treats *Gen* as an indexical A-quantifier is given. The indexical approach advocated has many virtues which are outlined in the chapter. Objections to the view are also addressed.

The fourth chapter discusses two further classes of generics — *existential generics* and *indexical singular generics*. Diagnoses of the issues arising from these classes

of data is given, and their relation to the positive view developed in the third chapter is outlined.

Acknowledgements

Very special thanks go to my supervisors Ephraim Glick at the University of St Andrews and Olav Gjelsvik at the University of Oslo, without whom this PhD project would not have succeeded. I also received invaluable feedback and support in the earlier stages from Peter Clark and Brian Weatherson.

I've been lucky enough to be a part of two very active research centres over the course of my PhD — Arché at the University of St Andrews and the Centre for the Study of Mind in Nature (CSMN) at the University of Oslo. I've met many inspiring people at both places to which thanks are due (apologies in advance for any omissions). At Arché: Derek Ball, Sebastian Becker, Mark Bowker, Jessica Brown, Herman Cappelen, Yuri Cath, Laura Celani, Stewart Cohen, Michael De, Josh Dever, Andy Egan, Noah Friedman-Biglin, Ephraim Glick, Patrick Greenough, Steven Hall, Katherine Hawley, Gudmundur Andri Hjalmarsson, Ole Hjortland, Thomas Hodgson, Torfinn Huvenes, Jonathan Ichikawa, Bruno Jacinto, Spencer Johnston, Dirk Kindermann, Ira Kiourti Julia Langkau, Martin Lipman, Ishani Maitra, Matthew McKeever, Dilip Ninan, Andrea Onofri, Walter Pedriali, Andrew Peet, Laura Cecilia Porro, Graham Priest, Simon Prosser, Francois Recanati, Marcus Rossberg, Jonathan Schaffer, Susanna Schellenberg, Anders Schoubye, Daniele Sgaravatti, Stewart Shapiro, Martin Smith, Jason Stanley, Andreas Stokke, Margot Strohminger, Joshua Thorpe, Michael Traynor, Crispin Wright, Alexander Yates and Elia Zardini. At CSMN: Nicholas Allott, Trine Antonsen, Robyn Carston, Timothy Chan, Einar Dunger Bøhn, Jan Terje Faarlund, Ingrid Lossius Falkum, Dagfinn Føllesdal, Christel Fricke, Olav Gjelsvik, Eline Busck Gundersen, Carsten Hansen, Heine Holmen, Jennifer Hornsby, Nick Hughes, Georg Kjoll, Ernest Lepore, Anders Nes, Thomas Pogge, Bjørn Ramberg, Kari Refsdal, Monica Roland, Mathias Sagdahl, Chiara Tabet and Deirdre Wilson.

Very special acknowledgement goes to the members of the Contextualism and Relativism project at Arché and the Linguistic Agency project at CSMN, and to Nicholas Allott, Derek Ball, Torfinn Huvenes, Dilip Ninan, Francois Recanati,

Jonathan Schaffer, Martin Smith, Jason Stanley and Elia Zardini who gave me comments on various parts of this dissertation.

There are also several great people working on the topic of generics who provided me with extremely valuable feedback, generous encouragement and/or endless inspiration: Nicholas Asher, Gennaro Chierchia, Ariel Cohen, Sarah Jane Leslie, David Liebesman, Michael Morreau, Bernhard Nickel and Jeff Pelletier. Very special thanks go to David Liebesman and Bernhard Nickel who provided me with detailed comments on several parts of this dissertation.

I was fortunate to have the opportunity to present chapters of this dissertation to the following audiences — these audiences include (in no particular order): the Harvard/MIT graduate conference, the Princeton/Rutgers graduate conference, the Yale/UConn graduate conference, the Society for Exact Philosophy, a workshop on the Philosophy of Language at Harvard University, a workshop on minimalism at Cambridge University, the first GENIUS generics conference, I am indebted to my respondents at these events: Mahrad Almotahari (at MIT), Philipp Koralus (at Princeton), Nick Kroll (at Yale) and Gennaro Chierchia (at Harvard), and also to audiences for their helpful comments and questions.

Many thanks also go to Lynn Hynd, Grethe Netland and Lina Tosterud for their always reliable and helpful practical support.

I also kindly acknowledge the CSMN and IFIKK at the University of Oslo and the Department of Philosophy at the University of St Andrews for their financial support.

Rachel Sterken
St Andrews
December 2012

Contents

1 Introduction

1.1 Introduction

Generics are sentences such as:

(1) Dogs bark.

(2) Candy is bad for your teeth.

(3) Lions have manes.

(4) Sharks attack bathers.

(5) Germans wear lederhosen.

(6) Sea turtles are long-lived.

These sentences communicate *something general* about dogs and candy — (1), for instance, seems to communicate something akin to *characteristically, all dogs that make noises bark* on its most favoured interpretation; while (3) seems to communicates something akin to *most male lions have manes.*

Notice, however, that in the sentences in question, there is no obvious candidate which is responsible for communicating this general content: There is no explicit quantifier and the bare plural — i.e., *dogs* in example (1) and *lions* in example (3) — combines with other predicates and doesn't express any such general content. One of the main puzzles of finding an adequate theory of generics, then, is to determine by what means they communicate their general content. The problem is even more difficult than I've let on since generics come in a variety of forms, as evidenced by the bare plural generics above and:

(7) A raven is black.

(8) This kind of animal barks.

(9) The tiger has stripes.

(10) Mary smokes.[1]

Moreover, almost all of these sentences have non-generic uses as well. What is more, no language is known to express whatever this general content happens to be with a dedicated, explicit expression.[2] How can we make sense of the logical form(s) of these sentences?

A second main puzzle for a theory of generics is to determine just what this general content(s) communicated by generics is(are) . This may seem like a reasonable task at first pass — "isn't the content just *all* or *most* suitably restricted in some way? Or, couldn't it just mean *in general*?" one might be inclined to say. But any one of these *prima facie* plausible proposals quickly runs into difficulties when considering the full range of generics. Sometimes generics sound like they communicate quite strong generalisations, as is the case in (1) and (2): Though there are a few dogs that don't bark, for instance because they are mute, (1) still seems to express that all dogs under normal circumstances make the characteristic noise, barking. Similarly, (2) seems to communicate that all forms of candy are bad for your teeth. But, sometimes generics sound like they communicate quite weak generalisations, as in (3) and (4): By no means do all lions have manes or sharks attack bathers, moreover, (3) and (4) don't seem to communicate anything nearly as strong as this. (3) seems to talk about *most* lions in some restricted domain (e.g., most male lions) and (4) could be thought of as expressing something as weak as *sharks sometimes attack bathers* — similarly, in the case of (5) and (6): very few Germans in fact wear the traditional costume, perhaps only on special occasions, and only about 1% of sea turtle hatchlings actually survive. Still further, there doesn't seem to be anything particularly unified to say about what sort of exceptions to the generalisations expressed by (1)-(2) are permissible — a restriction to normal instances is typically proposed, but this proposal has difficulty with cases like (4)-(6)[3] — for example, it is in no way abnormal for a sea turtle to die young. Further, when one ponders the intuitive content of (1)-(6) more directly, normality doesn't seem to figure prominently in all these cases, if

[1] Such sentences, are often called *Habituals*, it is still an unsettled question whether or not the class of habituals is a subset of the class of generics.

[2] Cf. Krifka et al. (1995) and Leslie (n.d.).

[3] Cf. Asher and Pelletier (2012) and Nickel (2012a) for recent defences of such normality based approaches.

it does at all.[4] What general and informative things, if anything, can we say about the intuitive content(s) of generics — i.e., their intuitive truth-conditions? Given that (1)-(6) have such different intuitive contents, can we even say that there is anything distinctive or shared about the general content communicated across this class of sentences?

The primary goal of this dissertation will be to offer solutions to these two puzzles which are informed by linguistic context-sensitivity and a less rigid conception of what these sentences have in common. The main thesis is that though generics share a logical form, there is little else they share: generics are highly context-sensitive and it is this that explains the elusive nature of their truth-conditions — i.e., that explains why providing an adequate solution to the second puzzle above is so difficult.

Put more precisely, my goal in this dissertation is to demonstrate that generics share a logical form which is quantificational, and even though there is no pronounced quantifier in their logical forms, there is nonetheless a covert quantifier which is responsible for expressing the general content, or generalisations, expressed by generics — this quantifier expression is known as *Gen*.[5] *Gen* is little more than a place-holder, however: *Gen* is a free variable of the same semantic type as a quantifier and hence, requires saturation in a context of utterance. In this way, generics do not express the same general content, or generalisations, in different contexts of utterance. In other words, the truth-conditions of generics vary with the context of utterance. Generics are context-sensitive and *Gen* is an indexical.

My approach in this dissertation, thus, is not to directly determine the content or truth-conditions of generics, but rather embrace the variability and elusiveness. I try instead to find general and informative patterns in the meaning of generic utterances across different contexts.

The dissertation is divided up as follows.

Chapter 2 discusses some purportedly problematic cases of generics — these are cases such as (4) and:

(11) Mosquitoes carry the West Nile Virus

(12) Rottweilers maul children.

[4] Cf. Chapter 4, Section 4.4.3.

[5] In particular, *Gen* is an adverbial quantifier (*A-quantifier*).

Such cases are meant to be problematic from the point of view of trying to find a distinctive semantics for generics, and are taken to motivate a wide variety of views about the linguistic and cognitive reality of generics.[6] I argue that cases like (4), (11) and (12) are not genuinely problematic as they are pathological. I provide an error theory to deal with them which is based on implicit cognitive biases. I take these examples as a case study of a broader issue linguists, philosophers of language and perhaps even language-users themselves should be wary of — the ways in which implicit biases can have an impact on interpretation and creep into the content of our utterances.

CHAPTER 3 is devoted to defending a quantificational analysis of generics in which the logical form of generics contains an unpronounced quantifier expression *Gen*. The chapter defends the original argument of Carlson (1989), the *multiple readings argument*, from some recent objections. In addition, the chapter provides a novel argument in favour of treating the logical form of generics as quantificational. The chapter also responds to various objections to a quantificational analysis, in particular, to the linguistic evidence offered in favour of a kind-predication analysis, and some recent arguments in favour of treating generics as kind-predications which are enriched by pragmatic mechanisms.

CHAPTER 4 develops and argues for the positive view that generics are context-sensitive and that *Gen* is an indexical — that is, that a given generic sentence expresses entirely different generalisations in different contexts of utterance. The view advocated has many virtues which are outlined in the chapter and objections to the view are also addressed. I hope that the indexical approach will give a radically new spin on the investigation of generics.

CHAPTER 5 provides extensive discussions of two important classes of generics — *existential generics* and *indefinite singular generics*. The chapter does not provide a settled view on either class of data. It does, however, develop further data, provide diagnoses of standard accounts, and sketches how the indexical approach in Chapter 4 relates to both classes of generics.

[6] Cf. Liebesman (2011) who uses data of the sort to argue that generics are kind-predications, and Leslie (2008, 2007) who uses data of this sort to motivate a theory of generics on which generics are sensitive to certain features of our implicit cognitive heuristics and biases, and who uses this type of data to motivate various theses about the structure of the mind.

2 Generics, Domain Restriction and Blindness

2.1 Introduction

Generics, like (1)-(3) below, communicate *something general* about dogs, tigers and candy:

(1) Dogs bark.

(2) Tigers have stripes.

(3) Candy is bad for your teeth.

Moreover, they do so without any overt or pronounced element which is responsible for communicating this content. Still, what characterises generics is meant to be their intuitive content or truth-conditions: The *something general* that they express is meant to be distinctive in that it is relatively stable across different generics and across different contexts. It is far from clear, however, just what this something general is and by what means generics express it. Further, given that theorists cannot even agree on what their intuitive content is, or how it is conveyed; this raises the concern that perhaps there isn't a distinctively generic content which is shared by this class of sentences at all.

These unclarities are apparent when one considers one of the main obstacles to providing an adequate theory of generics — namely, what I will call the *problem of variability*[1]. It is hard to pin-point just what the problem is without suggesting a

[1] This problem is articulated in many forms in the literature on generics (cf. especially Carlson (n.d.)).

particular source of the problem and hence, a path for its resolution, but the data is as follows. For many generics, like (1)-(3), the *something general* they express seems to be a quite strong generalisation — akin to *all normal dogs bark* or *generally dogs bark*, for example. Such generics tolerate exceptions, but the *something generals* they express seems to hold of nearly all or a majority of the members of the kind in question. Yet, for other generics it seems a strong interpretation doesn't suffice — consider:

(4) Books are paperbacks.

(5) Humans are Asians.

Generics like (4) and (5) are false despite nearly all or a majority of the members of the kind satisfying the predicated property. Such generics seem to require that even more instances — i.e., more than all the normal instances or a majority of the instances — satisfy the predicated property, in order for the generalisation to come out as true. Still further, for many other generics, it seems a weak interpretation is needed, weaker than a majority interpretation and often times almost as weak as an existentially quantified interpretation — consider:

(6) Birds lay eggs.

(7) Lions have manes.

(8) Mosquitoes carry the West Nile Virus.

(9) Sharks attack bathers.

(10) Rottweilers maul children.

Generics like (6)-(10) are intuitively true despite only a minority, even a very small minority of the kind satisfying the predicated property. Such generics seem to be very generous exception tolerators.

The least committal way of expressing the problem of variability is probably as follows: How can we specify the distinctive content (or a distinctive set of contents) that generics like (1)-(10) are meant to share while at the same time accounting for the variety of ways in which their intuitive truth-conditions vary? I think such a formulation of the problem of variability leaves the source of the variability open.

This chapter will focus on cases like (8)-(10): They render the problem of variability especially hard — some even argue intractable (cf. Leslie (n.d.), Liebesman (2011)). In this capacity and simply on their own, they are put to a tremendous

amount of work in the literature on generics — they are used as evidence for a wide variety of claims about the nature of genericity, the semantics of generic sentences and, interestingly, the connection between generic generalisations and social prejudice. To get a sense of all the different types of claims for which they act as primary evidence, witness the following list:

- The intractability of the problem of variability.
- The inadequacy of a *domain restriction strategy* for solving the problem of variability. (Leslie (2008))
- Generics are non-quantificational. (Leslie (2007))
- Generics are hyperintensional. (Leslie (n.d., 2007))
- Generics have a weak interpretation in addition to their strong interpretation. (Cohen (2001b))
- Generics are kind predications. (Carlson (n.d.), Liebesman (2011))
- Generics encode how kinds inherit properties from their members. (Liebesman (2011))
- The mind possesses a primitive, default, cognitive mechanism of generalisation; and generics express generalisations which track this mechanism. (Leslie (n.d., 2008, 2007))
- Using generic generalisations makes us prone to certain forms of prejudice thinking. (Leslie (2013))
- At least some generics presuppose or implicate claims about generic essences; and understanding this aspect of generics is important for social prejudice and ideology critique. (Haslanger (2011))

The main goal of this chapter is to argue that such cases of purported genericity, cases (8)-(10), are pathological, and hence shouldn't be used as evidence for the kinds of claims just outlined (or indeed any claims about the nature or semantics of generics). In particular, these cases pose no difficulty to theorists wanting to solve the problem of variability (at least in part) by appeal to the linguistic property of quantifiers known as (contextual) domain restriction (cf. Asher (2006), Asher and Morreau (1995), Cohen (n.d., 2008), Nickel (2008, 2010c), Pelletier and Asher (1997); Asher and Pelletier (2012), Schubert and Pelletier (1987)).

A *domain restriction strategy* explains the problem of variability by appeal to a variable domain of quantification. Supposing, for the sake of illustration, that generics quantify over all instances or the majority of instances, then since in different contexts, quantifiers have different domains of quantification, the domains of quantification of different generics, and different utterances of a single generic, can change depending on context. (1) for example, is true in normal contexts since the domain of quantification is suitably restricted to normal dogs that make sounds, while (6) is true in normal contexts since the domain is suitably restricted to female birds of reproductive age. Thus, it doesn't matter that only a minority of the kind in question satisfies the predicated property since nothing stops the domain of quantification from being only a very small number of instances.

Many theorists in the first camp have even proposed different mechanisms for how the suitable domain is determined relative to a given context. But cases like (8)-(10), in particular, don't seem to fit the mould. If we consider the tiny set of mosquitos that carry the West Nile Virus, the only thing that distinguishes those from the non-West Nile-carrying mosquitos is that they carry the virus. There is no systematic or significantly predictive set of features that pick out this subset. For this reason, it is very hard to see how to limit the domain of quantification in (8) with some property F so that *all F-mosquitos carry the West Nile Virus* (or *generally, F-mosquitoes carry the West Nile Virus*) come out true (of course, there's a trivial restriction if F is just those mosquitoes that carry the West Nile, but allowing such restrictions would surely cause worries of overgeneration). Similar remarks

apply to (9) and (10).[2,3]

This chapter argues that cases like (8)-(10) pose no difficulty for solutions to the problem of variability that involve domain restriction.

More importantly though, this chapter argues that cases like (8)-(10) show nothing particularly distinctive about the nature of genericity or the truth-conditions of generics — they should not be used to motivate baroque cognitively based or metaphysically based proposals about the nature of genericity, or implausible stories about the logical form of generics. Despite what many authors contend, a linguistically plausible semantics of generics can tell us things about the problem

[2]The other objection to the domain restriction strategy is that it predicts that there is a true reading of (11) when the complex predication is treated non-distributively:

(11) Birds lay eggs and are female.

Since the domain of egg-layers is a restriction of the female birds, the domain of (11) is the same as that of (6). Thus, since (6) is true, (11) must also be true. But, the objection continues, there is no reading of (11) on which it comes out as intuitively true.
The ease by which this reading is accessed and the following observation, I think are enough to abate this objection to the domain restriction strategy. First, it is easy to access the false reading of (11) where the complex predication is treated distributively: Since *birds are female* is patently false (in the relevant contexts), (11) is false. Further, consider the following:

(12) a. Tomy learned that his car had been stolen.
b. Tomy's car had been stolen.
c. Tomy's learned that his car had been stolen and Tomy's car had been stolen.

(13) a. She lactates and is female.
b. Betty gave birth and is female.
c. He pees standing up and is male.

(12c.) sounds strange in a similar manner to (11) and yet there is no domain to speak of. What likely explains the strangeness is that (12a.) presupposes that (12b.), and hence it is peculiar to proffer the content of (12b.) directly after saying something that presupposes it. Similarly, (13a.)-(13c.) sound peculiar, and there is no quantifier domain which could be responsible for the fact that it sounds peculiar. These examples demonstrate that the strangeness or unavailability of (11) is independent of anything to do with quantifier domains.

[3]With respect to the discussion in the last footnote, it is also worth noting how domain restriction works in cases where the complex predication interacts with explicit quantifiers like *all normal* or A-quantifiers. To make this clear (and bolster the point in the previous footnote), I find the following explicitly quantified sentences equally puzzling in the respect which is raised by the objection:

(14) a. All normal birds lay eggs.
b. All normal birds lay eggs and are female.
c. Generally birds lay eggs.
d. Generally birds lay eggs and are female.

Suppose, in (14a.), the domain is restricted to birds in a certain room, and suppose those birds are female. It seems strange to say (14b.), where we understand that to have a tacit restriction to birds in that room. With some work, I think you can get used to (14b.) and hear it as fine and true. In the same way, I think speakers have to do some work to get themselves to hear (11) as fine and true as well. A-quantified sentences, like (14c.) and (14d.) work even better to make this point. At minimum, I think it is prudent to hold off on giving a verdict on these examples.

of variability (i.e., about the ways in which generics tolerate exceptions and their truth-conditions vary). The manners in which variability arises needn't be a result of something non-semantic like our cognition (Leslie (n.d., 2008, 2007)), the metaphysics of kinds (Liebesman (2011)) or simply because generics are predications and it is not the job of semantics to tell us when an item satisfies a predicate (Carlson (n.d.)). Authors who defend such claims see cases like (8)-(10), as resounding evidence for these types of claims, but this is a mistake: (8)-(10) are pathological cases of genericity. I reiterate: (8)-(10) should not be used to baroque cognitively based or metaphysically based proposals about the nature of genericity, or implausible stories about the logical form of generics.

Finally, the purported connections between genericity and social prejudice which are supported by appeal to (8)-(10) are not features of genericity *per se*, rather the pathology of cases like (8)-(10) is a result of the same sorts of biases and errors which cause social prejudices in the first place. The fact that (8)-(10) are prejudice utterances can tell us something about the mechanisms of social prejudice, but (8)-(10) do not tell us about the nature or truth-conditions of generics themselves, or about how generics manifest and reinforce prejudice.

Given the variety of claims to which (8)-(10) are taken as primary evidence, the importance of these cases cannot be underestimated. I will argue that cases like (8)-(10) are pathological, but I won't deny that they are interesting. They are interesting because identifying the source of their pathology allows us to understand the types of circumstances and cognitive mechanisms which cause us to form mistaken generalisations, some of which will be unfair and will count as prejudice. Not only that, but they can also act as an important case study for linguists and philosophers of language about how cognitive mechanisms can affect our intuitive judgements and impact the communicated content of utterances.

The fact that we are inclined to judge generics like (8)-(10) as true does not entail that we are judging their semantic contents as true — it could be that we are judging some sort of pragmatically conveyed contents as true. Alternatively, especially given the nature of the examples, it could be that our judgments that cases like (8)-(10) express truths are merely mistaken. I will defend the latter view, which I call the *Blindness approach* to generics like (8)-(10). In what follows, I will provide evidence in favour of the Blindness approach, and arguments against the existing approaches to cases like (8)-(10) — including those of Cohen (n.d.), Asher (2006), Nickel (2010c) and Leslie (n.d., 2007, 2008).

The Blindness approach is an error theory for cases like (8)-(10). Throughout this chapter, what I will mean by generics like (8)-(10) are generics that express or convey contents, relative to contexts, that are:

1. Intuitively true even though the (vast) majority of the kind in question are exceptions to the generalisation, and;

2. the generalisation in question is Φ where Φ is the property the generalisations expressed must have in order for these minority readings to come out as intuitively true — whatever happens to be common amongst the generalisations in (8)-(10) which makes them intuitively true.

According to Leslie, Φ is something like the *strikingness* of the predicated properties (as attributed to the kind in question), where she characterises strikingness as follows:

> The examples… have something in common: In all of them, the sentence attributes harmful, dangerous, or appalling properties to the kind. More generally, if the property in question is the sort of property of which one would be well served to be forewarned, even if there were only a small chance of encountering it, then generic attributions of the property are intuitively true. (2008, p.15)

The notion of strikingness that Leslie appeals to has a basis in empirical psychology (cf. Rothbart, Fulero, Jenson, Howard, and Birrell (1978)), in particular, as it is understood as a cognitive heuristic or bias — i.e., a mental shortcut that allows our minds to make efficient judgements, and a hard-wired or systematic mental tendency which leads to errors in our judgment. Here the heuristic or bias would be understood as something like the tendency to generalise properties more easily when the property in question is a striking property for the kind. Though, I find Leslie's characterisation helpful for (8)-(10) — strikingness does seem to track the cases at issue — I think a more general proposal of what Φ amounts to is more plausible. A wide variety of cognitive heuristics might impinge on our judgements about whether a given generalisation holds, and affect how speakers produce and interpret generics. Further, many generic generalisations require quite complex justification to support their truth. As such, though the discussion below will focus only on cases like (8)-(10), which seem to follow a similar pattern of error, (8)-(10) are examples of a more general class of potentially mistaken generalisations and we have no reason to believe that potential mistakes are a result of a unique source — i.e., as a result of a single cognitive heuristic or bias. Thus, in identifying cases of potentially mistaken generalisations and the sources of error, I propose that we not limit ourselves to the strikingness heuristic, but look at cognitive heuristics more generally

With these preliminaries in place, I now go on to defend the Blindness approach.

2.2 The Evidence for Blindness

Blindness claims that intuitions about the truth and falsity of cases like (8)-(10) should not be taken at face value. Their interpretation is defective: When we intuit (8)-(10) as true, we are making a mistake. It should not be a condition on a theory of genericity that it makes generics like (8)-(10) come out as true. We might have intuitions to the effect that they are, but these intuitions are mistaken.

Having said that though, we cannot simply deny the intuitions that such generics are true, we do have them after all. What accounts for the fact that we are making errors? As mentioned above, there are obviously many ways of making mistakes. I will focus on how cognitive factors, cognitive heuristics and biases, like those Leslie appeals to in the above quote, can cause us to make errors. It could be that such cognitive factors interfere with how we form beliefs about the world, causing us to have incorrect beliefs — in this case, speakers exhibit a certain *blindness or ignorance about the world.* Alternatively, it could be that such factors interfere with how speakers interpret generics, causing us to get the content of the generic utterance wrong — in this case, speakers exhibit a form of *semantic blindness.* For simplicity, my focus will primarily be on the former.

The Blindness solution has the consequence that it serves as a vindication of the cognitive features Leslie (n.d., 2008, 2007, 2013) appeals to in her account of genericity (if one so wishes), while at the same time not forcing any drastic claims about the truth-conditions of generics, the nature of genericity, or for that matter, the metasemantics of domain restriction or the magic of pragmatics. Not only that, but it allows these troublesome cases to be merrily set aside.

Further, it takes seriously one key property which sets (8)-(10) apart. When one comes across (8)-(10), one has the distinctive initial reaction that they are somehow *unfair* or *unreasonable* generalisations. This reaction should not be ignored. Blindness says that although we are inclined to accept (8)-(10), they are in fact false; our reaction that they are unfair or unreasonable is a reaction which is caused by having knowledge which supports a judgement of falsity, but being unable to provide that verdict due to a conflicting inclination causing the opposite judgment.

Before going on to present the two very compelling pieces of linguistic evidence in favour of the Blindness view, I will attempt to make this last point into a key motivation in favour of the approach.

2.2.1 Motivating Idea: Unfair or Unreasonable Generalisations

The basic idea here is that when objects have salient features of certain kinds, we are inclined to overgeneralise, sometimes in unfair or unreasonable ways to the members of the relevant kind. The most apparent examples of this are prejudice or (negative) stereotypical generalisations — for example:

(15) Muslims are terrorists.

(16) Blonds are dumb.

(17) Homosexuals carry the HIV virus.

These are prejudicial, unfair and irrational generalisations. They are not propositions anyone should endorse. Nonetheless, people often do. They form unfair general beliefs of this kind.

The Blindness proposal sees minority generics like (8)-(10) as just like (15)-(17). The fact that we take generics like (8)-(10) as intuitively true, when they are in fact false, is a result of unfair and poor generalising. There is no objective difference between (8) and (17) for example, they are a result of the same unfair reasoning and biases. The Blindness approach says that they are all false since the generic requires that in general the kind possess the predicated property; the fact that we think they are true is based on irrational mechanisms that cause us to overgeneralise and form a mistaken belief about the generality over which mosquitoes carry the West Nile virus.

Leslie (2013) discusses generalisations like (15)-(17) as well. The main difference between our accounts is that we draw the line between good and poor generalisations differently. I think (15)-(17) and (8)-(10) are poor generalisations, whereas Leslie thinks that some of these generalisations as fine in virtue of a minority of the kind satisfying the predicated property — these generalisations are good since the members of the kind that don't satisfy the predicated property are at least disposed to have it. Since we judge that the non-West-Nile-carrying mosquitoes are disposed to carry the West Nile and hence, (8) is true, whereas (17) is false since the non-HIV-carrying homosexuals are not disposed to have HIV.[4] The Blindness view, thus, undermines an important part of her theory.

[4]Leslie needs some story about what accounts for the difference in disposition attribution of the diseases in each case. But no plausible story appears to be forthcoming: One would expect that disposition attributions of carrying viruses rely on relevantly similar circumstances to support their truth. Nothing about the difference in virus or the difference in species appears to explain any difference in disposition attribution in the case of (8) and (17), but Leslie must claim that they are different in some significant way in order to account for the difference in truth-value of (8)

It is worth noting at this point that the Blindness account can also be extended to account for *overly-fair* generalisations or positive stereotypes as well — consider the following:

(18) Black people are good at basketball.

(19) Asians are good at math.

(20) Liberals are intellectuals.

Such generics are intuitively true despite only a minority of the kind in question satisfying the predicated property; just as is the case for (8)-(10). Moreover, there doesn't seem to be a sufficiently predicative subset of the kind in question which could act as an appropriate restrictor on the domain restriction strategy; just as (8)-(10). What's interesting about (18)-(20) is that we get the sense that they are overly-fair generalisations, which gives us motivation to question whether they are strictly speaking true despite their compellingness as apt generalisations.[5]

Next, I will go on to present two pieces of linguistic evidence for Blindness. Later in this section, I will offer some additional considerations in favour of Blindness and reply to some *prima facie* objections.

2.2.2 Evidence 1: Contradictory Conjunctions

Consider the following paradoxical sounding sentences:

(21) a. Mosquitoes carry the West Nile Virus, but typically they don't.
b. Typically mosquitoes don't carry the West Nile Virus, but mosquitoes carry the West Nile Virus.

(22) a. Sharks attack bathers, but generally they don't.
b. Generally sharks don't attack bathers, but sharks attack bathers.

(23) a. Rottweilers maul children, but normally they don't.
b. Normally rottweilers don't maul children, but rottweilers maul children.

and (17). The Blindness view has no such problem since (8) and (17) share a truth-value — they are both false despite the irrational mechanisms which might compel us to judge them as true.

[5]Unfairness and overly fairness is being used in this chapter as a guide for finding cases: I am not saying unfair is synonymous with false and that fair is synonymous with true!

(21)-(23) sound contradictory or at least bad, but according to those advocating the truth of (8)-(10), (21)-(23) could turn out as true since these theorists are committed to both conjuncts in each case being true.[6]

How can Blindness explain their falsity? If we accept Blindness, then we can simply say that the first conjunct of (21a.) is close in meaning to that of the negation of the second conjunct, and so their conjunction is false in the relevant contexts. Our initial reaction is that the first conjunct of (21a.) is true, but this initial reaction is a mistake. Cases like (21)-(23) make their falsity clear to us.

The point above can be made more precise in several ways — and I won't choose between the possible elaborations here. I will simply note, first, that what I say is compatible with treating the first conjunct of (21)-(23) as close in meaning to *It's not the case that generally sharks don't attack bathers*, and also close in meaning to *Generally sharks attack bathers*. Either way would predict the paradoxical feeling (though, note that the former might be consistent with *It's not the case that generally sharks attack bathers*).[7] Second, I should note that I have used the vague expression "close in meaning" but there are alternatives. A more radical option,[8] is to say that the first conjunct entails the negation of the second. The first conjunct could be stronger, and so not close in meaning to the first conjunct. The paradoxical feeling is still explained.

One might worry that the conjunction test over-predicts in various ways — consider for instance:

(24) a. Ravens are black, but some aren't.
b. Some ravens are white, but ravens are black.
c. Ravens are black, but sometimes they're not.

(25) a. Mosquitoes carry the West Nile virus. Rarely, but they do.
b. Mosquitoes carry the West Nile virus, even though most don't.

(24) and (25) bring up different issues, so I will discuss each in turn. Some of my informants claim that (24a.)-(24c.) sound bad, maybe in a way analogous to how (21)-(23) do. But I do not think that they sound equally bad or even bad in a similar way (i.e., a contradictory way). It is, for example, fine to say:

(26) Ravens are black, but there are a few exceptions.

If we are fine with (26), then we should be fine with (24a.)-(24c.) as well. Further, recall that we noted above that it could be that *Sharks attack bathers* is stronger

[6] These examples were generated out of fruitful discussions with Jonathan Schaffer.
[7] Thanks to Ephraim Glick for pointing out these options to me.
[8] Thanks again to Ephraim Glick for the suggestion

than *It's not the case that generally sharks don't attack bathers.* If one makes this move, then the corresponding argument in the case of (24) would conclude that *Ravens are black* entails *It's not the case that some ravens aren't black* — i.e., *All ravens are black.* Paired with a domain-restriction on *all* (and *some*), this is better than concluding that *Ravens are black* is as weak as an existential, as would be the case on the wide-scope reading of the second conjuncts in (24). Still further, if one insists on hearing some infelicity in (24a.)-(24c.), then I claim that this is due to a tension in the relevance of the *some*-claim or the domains being talked about. Thus, I contend, the conjunctions in (24) do not entail that generics are close in meaning to existentially-quantified sentences.

Other troubles for the conjunction test arise with (25a.) and (25b.): (25a.) and (25b.) sound relatively fine, but this makes it seem like the existentially-quantified readings of (8)-(10) are legitimate (and perhaps also, that a wide-scope reading of the negations of the second conjuncts in (24) and the bad readings of (24) are indeed possible). Though, (25a.) and (25b.) are compatible with an existentially-quantified reading of (8), they are also compatible with a strong reading as well: (25a.), I contend, is what Cohen (2003) calls a *quasi-existential* reading which he argues is consistent with a strong interpretation (cf. Chapter 5 of this dissertation). The only way I hear (25b.) as acceptable is if the generic is given a capacity reading which is consistent with the strong interpretations, though due to the presence of the expression *even though*, it is plausibly as a quasi-existential reading (again, cf, Chapter 5).

2.2.3 Evidence 2: Disagreement and Retraction

Some further evidence that we should not trust our initial reaction to generics like (8)-(10) comes from examples involving disagreement and retraction. Consider:

(27) A: Let's stay inside. Mosquitoes are out there and they carry the West Nile.
B: That's ridiculous. Almost none of them do.

(28) A: Sharks attack bathers.
B: No, they almost never do.

(29) A: Rottweilers maul children.
B: Don't be silly. There have only been a few isolated incidences.

The dialogues in (27)-(29) sound like genuine disagreements. In (27), for example, speaker B believes that the fact that almost no mosquitoes carry the West Nile is good reason for her to disagree with A's assertion of (8). If (27)-(29) are

indeed genuine disagreements, then we have evidence that the generics (8)-(10) cannot be true when speaker B's assertions — that almost no mosquitoes carry the West Nile, that sharks almost never attack bathers and that there have only been a few isolated incidents of rottweilers mauling children — are true. Since B's assertions are true, we have evidence that (8)-(10) are false and hence, for Blindness.

These observations are further strengthened by the way one would expect speaker A to retract. In examples like (27)-(29), speaker A would naturally retreat to something like: *Well, at least some of them do.* A natural explanation of this retraction is that A's initial statement — i.e., the generic claim — is false. These kinds of dialogues I suggest should be taken as evidence that (8)-(10) are not true in general — when we think they are true we are suffering from a form of blindness.

Similar concerns to those raised in the previous section with (24) arise for the disagreement diagnostic proposed here — consider for instance:

(30) A: Sharks attack bathers.
B: No, some are perfectly gentle.

My reaction to cases such as these is the same. Consider for example:

(31) A: Sharks normally attack bathers.
B: No, some are perfectly gentle.

(32) A: Sharks always attack bathers.
B: No, some are perfectly gentle.

If the same sense of disagreement can be generated with explicitly quantified sentences as in (31)-(32), then this obviously does not entail that the explicitly quantified sentences in question are close in meaning to their existentially-quantified counterparts. Thus, I contend, disagreement dialogues like the one in (30) do not entail that generics are close in meaning to existentially-quantified sentences.

Next, I will consider various considerations in favour of Blindness. Later in this section, I will go on to defend the view from a crucial objection.

2.2.4 Mistakes and Danger

I think a somewhat speculative, but *prima facie* plausible, generalisation about human psychology supports the Blindness view (and gives it an edge over, e.g., Leslie's proposal): It is natural for people in general (and also for theorists)

to make mistakes in attributing and generalising properties when they are confronted with potentially scary, appalling, worrying, harmful or dangerous information. Even just thinking about striking information of this sort, can put thinkers in an uneasy, worried or alarmed state of mind, and being in such states of mind makes it more likely that thinkers will make errors — at the very least, it makes the task of generalising a cognitively higher-demand task than it might have been otherwise. According to Blindness, our attributions and generalisations are not true in virtue of their strikingness, but rather it is their strikingness that is causing us to make mistakes in attribution and generalisation.

To see this, consider the following examples:

(33) a. Airplanes crash.
b. Airplanes are dangerous.

When a speaker is about to embark on an voyage by air, (33a.) and (33b.) are rather scary propositions to contend with. They are not factually true: there are very few airplanes that crash or are dangerous — one is much more likely to die in a car crash than in an airplane crash; and yet we judge them as true and do not judge *cars crash* and *cars are dangerous* as true.. Even so, this is a very natural mistake for someone to make in this kind of setting — the speaker is in a particular state of mind that causes her to believe that the generalisation is true. Similarly, when a speaker utters a mere attribution, *this plane is going to crash*; her state of mind can be what causes her to treat the generalisation as true, even though it may be false (in the relevant context).

I take it as a virtue of the Blindness approach that it leaves open the possibility that the states of mind in question can contribute to our errors — they are not irrelevant to the misattributions and overgeneralisations at issue. This contrasts with the account given by Leslie. According to her account, in cases of misattribution or overgeneralisation it is our lack of knowledge of the dispositions associated with the having of these properties that can cause us to misattribute or overgeneralise striking properties. Another way to put this is as follows: The "danger factor" or "scariness impact" of striking properties on the Blindness view is responsible for causing errors, while on Leslie's account the "danger factor" merely warrants a different mode of attribution and generalisation, the "the danger factor" doesn't have anything to do with what causes a misattribution or overgeneralisation. Specifying the correct role of the "danger factor" or "scariness impact", I contend, is a consideration in favour of the Blindness approach.

2.2.5 Primitive Generalisations vs. Mistaken Generalisations

Is the empirical evidence which supports treating (8)-(10) as primitive generalisations in conflict with the Blindness approach? Are the psychology studies that Leslie (n.d., 2008, 2007, 2013) argues support her theory, that generics are a special type of primitive generalisation which is sensitive to certain cognitive heuristics or biases — e.g., strikingness — compatible with the Blindness approach?

In addition to the Rothbart et al. (1978) study, Cimpian, Brandone, and Gelman (2010) found that generic generalisations involving properties that were described as distinctive or dangerous were accepted more often than generic generalisations of other similar properties. These studies are of course perfectly consistent with the Blindness approach: The studies do not say what the truth-conditions of generic sentences are, rather they demonstrate a heuristic or bias in how the mind handles striking information — this demonstrated heuristic or bias can be treated as impacting the truth-conditions of generic sentences (Leslie's view) or as merely impacting our judgments about the truth-values of such sentences.

In one study, conducted by Abelson and Kanouse (1966), some generics were found to require very little evidence for acceptance (like (8)-(10)), and yet "once accepted psychologically they appear[ed] to be commonly taken in a rather strong sense, as though the quantifier *always* had implicitly crept into their interpretation". This finding is consistent with the Blindness approach: If a speaker thinks that, for example, the generic (8) is true and goes about her thinking as if it has a strong interpretation, then when thinking about mosquitoes, her tendency to overgeneralise will extend to sentences with explicit quantifiers. That is, she will also often be inclined to accept *many mosquitoes carry the West Nile* or *lots of mosquitoes carry the West Nile* as true, and *mosquitoes don't carry the West Nile* as false. This is what one would expect on the Blindness view.

The studies which could be seen as having a bearing on how we treat the truth-conditions of (8)-(10) are neutral between Leslie's approach and the Blindness approach. I will now turn to developing a test which can help decide the matter.

2.2.6 Kripke's Test and Weak Interpretations of Generics

In 'Speaker's reference and semantic reference', Kripke recommends a test to determine whether alleged counterexamples to a semantic proposal should count as genuine counterexamples. According to Kripke's test:

> If someone alleges that a certain linguistic phenomenon in English is a counterexample to a given analysis, consider a hypothetical lan-

> guage which (as much as possible) is like English, except that the analysis is stipulated to be correct. Imagine such a hypothetical language introduced into a community and spoken by it. If the phenomenon in question would still arise in a community which spoke such a hypothetical language (which may not be English), then the fact that it arises in English cannot disprove the hypothesis that the analysis is correct for English. (1977, 16)

In this context, something like Kripke's test can be a helpful heuristic or diagnostic for thinking about the examples at issue.[9] Kripke's original test was perhaps not intended for these purposes, but a variant can be used. The thought is that a semantic theory that treats our (initial) positive judgements about (some) minority generics as mistaken is supported by Kripke's test: In the language where we stipulate that such generics are false and speakers' psychological propensities remain as now, the initial (now clearly mistaken) intuition is still likely to be present.

Thus, the idea is that this variant of Kripke's test can be used in favour of the conclusion that cases like (8)-(10) do not count as counterexamples against the hypothesis that they receive a strong interpretation, i.e., the same interpretations as paradigmatic generics like (1)-(3). In other words,(8)-(10) needn't be given a weak interpretation whereby they are only true in virtue of a minority of the instances in question satisfying the predicated property.

Applying a variant of Kripke's test, suppose that L is a language in which generics like (8)-(10) are given an analysis consistent with the strong interpretation — e.g., for (8), something along the lines of *all normal mosquitoes carry the West Nile Virus* or *generally, mosquitoes carry the West Nile Virus.* Further stipulate that L is a language just like English in every other respect and that the psychology of speakers of L is as close as possible to the psychology of speakers of English. Now we ask the question: Would speakers of L make the mistake of thinking that (8)-(10) are true? Would a minority of the kind in question having some salient striking property lead speakers to make a mistaken generalisation? If your inclination is to answer yes to these questions, which is the inclination of my informants, then this is no argument against English being just like L — i.e., English having

[9]This isn't the place for an evaluation of Kripke's test. It is generally taken to be at least a helpful heuristic for determining whether a given linguistic phenomenon is semantic content. That said, it is worth noting that there is disagreement about just what the test is a diagnostic of. Kripke takes it to provide support for semantic analyses (facing a set of counterexamples), but it takes argument to show that a given analysis is a correct account of the semantics, rather than some other level of content. I will not try to settle that issue here.

a semantics for cases like(8)-(10) on which they are given an analysis consistent with the strong interpretation.

With sufficient positive motivation and evidence in favour of Blindness, I now move on to the main objection to the blindness approach — namely, that the evidence can cut in some questionable ways.

2.2.7 Objection and Reply: Overgeneration with Other Weak Interpretations of Generics

A genuine worry with both the contradictory conjunction and the disagreement and retraction data is that it overgenerates — that is, that it applies equally well to other seemingly or indisputably weak interpretations of generics; indicating that the advocate of Blindness needs to claim blindness for these legitimately weak interpretations as well. Consider again example (6), plus the following:

(34) Bees reproduce.

(35) Dobermans have floppy ears.

(36) Frenchmen eat horse meat.

(37) Scots wear kilts.

Each of these generics is intuitively true despite only a minority of the kind in question satisfying the predicated property: Only the adult female birds lay eggs; there is normally only one adult, mated queen bee that reproduces for each bee hive; Dobermans are born with floppy ears but breeders chop them off so that most have short pointy ears; very few Frenchmen actually eat horse meat and finally, generally Scottish men nowadays only wear kilts for special occasions.

The advocate of Blindness does not, or at least may not, want to be committed to the claim that generics like (6) and (34)-(37) are false, but it looks like the arguments she used in favour of Blindness work in the case of (6) and (34)-(37) as well. Consider:

(38) a. Birds lay eggs, but typically they don't.
b. Typically birds don't lay eggs, but birds lay eggs.

(39) a. Bees reproduce, but generally they don't.
b. Generally bees don't reproduce, but bees reproduce.

(40) A: Let's get a Doberman. Dobermans have floppy ears.
B: That's ridiculous. Almost none of them do.

(41) a. Frenchmen eat horse meat, but typically they don't.
b. Typically Frenchmen don't eat horse meat, but Frenchmen eat horse meat.

(42) A: Scots wear kilts.
B: Actually, hardly any of them do.

(38a.) and (38b.), (39a.) and (39b.), and (41a.) and (41b.) sound contradictory, and the disagreements in (40) and (42) sound like those in (27)-(29).

The advocate of Blindness could of course go all out and claim that (6) and (34)-(37) are false as well. She could for instance appeal to the idea that (36) and (37) are clichéd or silly stereotypes and that such generalisations should not count as true generics. However, the advocate of Blindness need not go this route.

If the advocate of Blindness does not want to claim that generics such as (6) and (34)-(37) are false, then she needs to say what separates the interpretation of (8)-(10) from that of (6) and (34)-(37)? The easy answer is that our best theories of domain restriction apply straightforwardly to make (6) and (34)-(37) true (on the strong interpretation), while they fail for generics like (8)-(10) true (on the strong interpretation).

2.2.7.1 Seemingly Weak Truth-Conditions

If domain restriction applies straightforwardly to (6) and (34)-(37), then what appears to be a weak interpretation is in fact a strong interpretation suitably restricted. Moreover, provided domain restriction works for (6) and (34)-(37), the advocate of blindness has an explanation of the contradictory conjunctions and disagreements in (38)-(42).

For example, (38a.) and (38b.) sound contradictory, because they *are* contradictory. According to a fairly standard theory of domain restriction for quantifiers (in particular, A-quantifiers) (cf. Cohen (2008), Beaver and Clark (2008)), the domain of quantification is restricted by the maximal set defined by the presuppositions of the predicate. On such an account, the first conjunct of (38a.) says something akin to (43a.) or (43b.) and the second conjunct of (38a.) says something akin to (44):

(43) a. All normal birds that extrude offspring, lay eggs.
b. Generally, birds that reproduce in some way, lay eggs.

(44) Typically, birds that reproduce in some way, don't lay eggs.

On this construal, the first conjunct of (38a.) is true, while the second conjunct is false, which explains why their conjunction sounds bad – it sounds bad, because it is false. Thus, the advocate of Blindness has a plausible explanation of the contradictory conjunction data, in terms of domain restriction, for generics that exhibit seemingly weak truth-conditions.

The disagreement data in (40), for example, can be dealt with using the properties of domain restriction as well. In (40), the sense of disagreement results from A making a mistaken claim which B is correcting by pointing out that what A has said is false and why it is false. According to our best theories of domain restriction, topic constrains the selection of the domain of quantification. Thus, in the case of (40) for instance, if the topic of the dialogue is pet Dobermans, then the domain picked out by A's statement is not the floppy eared Dobermans, but the pointy eared ones, so what A says is false. B points this out. The advocate of blindness has an explanation of disagreements for generics that exhibit seemingly weak truth-conditions.

The difference between cases like (8)-(10) and those that exhibit seemingly weak truth-conditions is that standard theories of domain restriction explain their seeming weakness, and the contradictory conjunction and disagreement data. But are there examples which are not seemingly weak, but legitimately weak interpretations of generics?

2.2.7.2 Indisputably Weak Truth-Conditions: Distinguishing Property Generics

In the same way, that domain restriction doesn't work for (8)-(10), it seems to not work for generics like (36) and (37). If we consider, for example, the minuscule group of Scots that actually wear kilts, the only thing that distinguishes them from the non-kilt-wearers is the kilt-wearing itself. There is no systematic or significantly predictive set of features that pick out this subset. For that reason it is very hard to see how to limit the domain of quantification in (37) with some property F so that *all F-normal Scots wear kilts* or *generally, F-Scots wear kilts* come out true.

Such sentences receive what theorists call a *distinguishing property interpretation.* Other possible examples include:

(45) Germans wear lederhosen.

(46) Dutchmen are good sailors.

(47) Oysters have pearls.

(48) Inuits live in igloos.

Krifka *et al.* characterise the distinguishing property interpretation as follows:

> The somewhat weaker claim of [(36)] can be paraphrased as: The [French] are known to [eat horse meat]. Or: The [French] distinguish themselves from other comparable nations by [eating horse meat]. It is this somewhat weaker interpretation we call the distinguishing property interpretation, as it seems to express a property that distinguishes the subject referent from other entities that might belong to the same category. (1995, 81-3)

It is much more difficult to say what the difference is between cases like (8)-(10), and (36)-(37) and (45)-(48). Indeed, distinguishing properties are often striking, and if we take strikingness as that which identifies cases like (8)-(10), then what is the difference between (8)-(10) and (36)-(37) and ((45)-(48)?

I take it that there are two options for how to deal with cases like (36)-(37) and (45)-(48) for the advocate of Blindness: The first option treats them like (8)-(10) and claims that they are false generalisations which we mistakenly take to be true; on this view the distinguishing property interpretation is not a legitimate generic interpretation. The second option treats the distinguishing property interpretation as a distinctive weak interpretation of generics — perhaps treating them as Krifka *et al.* suggest in the above quote or perhaps in some other way (cf. Cohen (2001b) and Nickel (2010b) for a treatment of (46)).

The trouble with cases like (36)-(37) and (45)-(48) is that no matter which option we choose for how to treat them, there are reasons against it.

On the one hand, if we treat (36)-(37) and (45)-(48) along the lines of the first option and say that they (semantically) express clichés or stereotypes, then there doesn't seem to be any substantive sense in which they are unfair or overly-fair generalisations. Without the unfairness or overly-fair motivation, it is more difficult to see why we would want to treat them as false. Further, the retraction patterns for cases like (36)-(37) and (45)-(48) seem to be different from regular generics and cases like (8)-(10) — consider again:

(49) A: Let's stay inside. Mosquitoes are out there and they carry the West Nile Virus.
B: That's ridiculous. No they don't, almost none of them do.
A: Well, at least some of them do.

(50) A: Let's get a Doberman. Dobermans have floppy ears.

B: That's ridiculous. Almost none of them do.
A: Well, some of them do, but I guess those aren't available as pets.

(51) A: Scots wear kilts.
B: Actually, hardly any of them do.
A: Yes, but it is very distinctive of Scots that they wear kilts.

In the case of (49) and (50), A seems to be genuinely retracting her initial generic claim, but in the case of (50), it seems possible for A to, rather, explain why she thinks it suffices that only a minority of Scots wear kilts in order for her initial generalisation to be true. So long as dialogues like (50) are appropriate — i.e., that A is not correcting herself, but rather pointing out that she is appealing to distinctive properties — we have evidence in favour of the second option and against the first option.

On the other hand, if we do take the differences in retraction patterns seriously, and advocate the second option, then we must explain why conjunctions of the following sort seem fine:

(52) Germans wear lederhosen, and many other people do too.

(53) Germans wear lederhosen, but it's not very distinctive.

If we are to take the distinctive part of the distinctive property interpretation seriously, then (52) and (53) should sound contradictory or at least bad in some way; but they don't and hence, we have evidence in favour of the first option and against the second.

Thus, no matter which account of cases like (36)-(37) and (45)-(48) the advocate of Blindness opts for, there seems to be evidence against it. The advocate of Blindness needn't be overly concerned, however, since this is a problem no matter which account of cases like (8)-(10) one chooses — they all need to explain the distinctive property interpretation, the difference in retraction patterns and why conjunctions like (52) and (53) sound fine.

As should be clear from the last section, the Blindness solution faces some significant objections and it is not entirely straightforward to defend it. It is worthwhile, therefore, to compare it to other proposals. I'll conclude these other proposals do even worse and this makes the Blindness solution all the more plausible.

2.3 Blindness does Better: Other Accounts

There are three existing proposals for generics like (8)-(10): Cohen (n.d., 2001b) treats them along side other weak-sounding generics and claims that generics,

more generally, are systematically ambiguous between a strong (what he calls *absolute*) interpretation and a weak (what he calls *relative*) interpretation; Asher (2006) and (Nickel, 2010c) argue that (8)-(10) are not weak, but rather strong generic capacity readings; and Leslie (n.d., 2008, 2007) gives them great prominence and argues that they are evidence for a cognitively based approach to their semantics. I will maintain that the Blindness approach should be preferred over all of these proposals on the grounds that there are objections to each of them; I consider each proposal in turn. At the end I'll consider a *prima facie* plausible pragmatic proposal and argue that upon closer inspection it is not plausible after all.

2.3.1 Cohen on Relative Readings of Generics

Cohen (n.d., 2001b) argues that generics are systematically ambiguous between a strong (what he calls *absolute*) interpretation and a weak (what he calls *relative*) interpretation. Generics like (8)-(10) are a species of a more widely available weak interpretation he calls the *relative reading* of generics. The relative reading asserts that the members of the kind at issue are more likely than an arbitrary member of other kinds (in the same category), to satisfy the predicated property:

> Relative readings of generics are more common than may seem at first sight. Indeed, many naturally occurring generics are only true if given a relative interpretation. Consider [(54)], for example:
>
> (54) [Tigers]$_B$ eat people.
>
> This sentence is false under the absolute reading: it is not the case that, in general, tigers eat people — very few do, in fact. Nor is it true that, in general, people who are eaten by some animal, are eaten by tigers. However, [(54)] is true under the relative interpretation, since a tiger is more likely to eat people than an arbitrary animal is. (2001b, p.10)

Leslie (2008) offers the following criticism of Cohen — namely, that Cohen over-predicts true readings of generics:

> The notion of a relative generic is also troublesome. The category of relative generics was introduced to deal with sentences like [(11)]. The incidence of West Nile virus-carrying is higher among mosquitoes than among arbitrary insects, so it is a true relative generic. We should worry, though, that the possibility of relative generics opens

> the door too wide… [A]ny time, say, a handicap or birth defect D is more frequent among one kind than all others, then *Ks have D* will be a true relative generic. Thus, for example, *humans are one legged*, or autistic, or blind, would be true relative generics by Cohen's lights. (2008, p.12)

Leslie's criticism fails, however, since it leaves out a crucial feature of Cohen's account of relative generics and inappropriately assumes that there are no contexts in which her counterexamples are in fact true.[10]

Cohen specifies that speakers select the relative interpretation in contexts where the generic is *B-marked* — that is, when the intonation contour (in English) of the generic is *fall-rise*, or when there is the appropriate topic/focus markers present in the logical form of the generic (cf. Buring (2011) and Krifka (2006)). Leslie's counterexamples, thus, only work if they are in fact false when they are B-marked; otherwise, they are merely false absolute readings according to Cohen's theory.

Briefly, given some bare basics of standard theories of information structure, B-marking induces in the context of utterance a particular question under discussion, which the utterance addresses. In the case of Leslie's examples, some candidate corresponding questions under discussion induced by the B-marked utterance are as follows:

(55) Candidate questions: What things are autistic? / What animals are autistic? / What animals have diseases?
[Humans]$_B$ are autistic.

(56) Candidate questions: What things are blind? / What animals are blind? / What animals have eye-troubles?
[Humans]$_B$ are blind.

As answers to the corresponding questions under discussion, Leslie's counterexamples sound intuitively true or at least it is not straightforward that they are plainly false. Thus, Cohen's account of generics like(8)-(10) is unscathed by Leslie's objection.

Though Cohen has escaped Leslie's objection, there are still problems with his proposal for cases like (8)-(10).[11] The first is that there are intuitively true

[10] I am also leaving out crucial features of Cohen's account — namely, his condition of *homogeneity* and a detailed discussion of his interpretation of probability. Such details are not directly relevant here and so I leave them out for simplicity's sake.

[11] It is important to note that I am not claiming that the relative interpretation itself does not exist. I am merely claiming that it is not the interpretation of cases like (8)-(10).

striking property generics which are not B-marked and thus, cannot receive a relative interpretation. The second, very potent worry is that his account does make too much true just as Leslie had said.

2.3.1.1 Mosquitoes, Sharks and B-marking

One should, rightfully I think, refrain from doing armchair philosophy/linguistics when subtle questions of intonation arise; however, since a full-scale empirical study is not in the cards, I will do so in any case. One thing to note from the get-go is that there is no reason to think that contexts with particular formal properties like B-marking should go hand-in-hand with contexts in which the predicated properties are somehow striking or appalling (which Leslie takes as the hallmark of examples like (8)-(10), and which is an observation which generally tracks the cases at issue). To this end, here are some examples where they come apart — consider:

(57) Doctors don't wash their hands.

(58) Drunk drivers are murderers.

In a similar way to (8)-(10), (57) and (58) are, with the appropriate contextual priming, intuitively true, despite only a minority of the kind satisfying the predicated property.[12] In addition, like (8)-(10), they needn't be B-marked in order for this to be the case: Imagine uttering (57) or (58) with different intonation (instead of a fall-rise intonation) and they still come out as intuitively true. They are also intuitively true in contexts where an explicit question isn't compatible with B-marking — consider:

(59) A: What alarming things do people fail to do?
B: Well, doctors don't wash their hands.

(60) A: What about drunk drivers?
B: Drunk drivers are murderers.

Note that (57) and (58) as they appear in (59) and (60) are not B-marked since the predicate is focused in each case.

Further, it is not the case that doctors are more likely than an arbitrary person, or even an arbitrary health care worker, to not wash their hands; and it is not the

[12] One might get the sense that (57) and (58) feel different in some important respect from (8)-(10), however, it is important to ignore any feeling that arises from the fact that it is simply easier to be in the right context to hear (8)-(10) as intuitively true, whereas (57) and (58) require some contextual priming.

case that drunk drivers are more likely to be murderers than the arbitrary person is. Thus, both (57) and (58) are false according to the relative interpretation, and yet they are intuitively true examples of cases like (8)-(10). Cohen's relative interpretation will not work for the cases at issue.

2.3.1.2 Relative Readings and Prejudice Generalisations

Another worry with Cohen's account is that although it does make cases like (8)-(10) come out as true. It makes a lot more come out this way too — consider:

(61) Blacks are criminals.

(62) Women wear lipstick.

Examples (61)-(62) (when B-marked) are true relative generics according to Cohen since blacks are more likely than an average American of other races is to be a (convicted) criminal and women are more likely than men to wear lipstick. But they are obviously intuitively false and are extraordinarily unfair generalisations. Legitimising unfair generalisations about mosquitoes and sharks is one thing, but legitimising unfair generalisations about ethic or racial groups is not something a semantics for generics should do.[13]

Thus, given these two objections, I conclude that Cohen's account of generics like (8)-(10) is not adequate. Moreover, it seems implausible that any formal account along the lines of Cohen's is forthcoming.

2.3.2 Nickel and Asher on Generic Capacity Readings

Both Asher (2006) and Nickel (2010c) claim that (8)-(10) are unproblematic since the sentences in question receive a capacity reading which allows them to be interpreted along quantificational/majority based lines without any difficulty.

> ... let's look at the first class of counterexamples, according to which, on Leslie's account, the satisfiers of the restrictor of a generic claim exhibit some striking property. Here it seems that modal approaches which take into account the differences between eventive and stative sentences and which take the *Gen* operator to bind event variables as well as individual variables already do the trick. Consider for example [(8)].

[13] Of course, there are interesting moral considerations here, but the *should* is only meant to be read in the normal metaphysical sense.

> With the accommodation of an underspecified restrictor, [(8)] says something like *all mosquitoes in appropriate circumstances carry the West Nile virus*, which is close to something like *all mosquitoes normally can carry the West Nile virus*. Striking eventive properties may just be modalized in an appropriate fashion. Such interpretations are of course perfectly consistent with the observation that in fact very few mosquitoes actually carry West Nile virus. (Asher, p.4)

Nickel has a similar line:

> Of course, [(8)] and [(9)] are troublesome only if they predicate a property that only a minority of members of the kind at issue have. Importantly, they also have a weaker reading on which [(8)] and [(9)] are completely unproblematic. Put intuitively, the difference in readings is a difference between an ascription of a habit and of a capacity. If the latter was at issue, a normality-analysis would predict [(8)] to roughly mean *all normal mosquitos (in the relevant domain) have the capacity to carry [the West Nile virus]*, and this is indeed true. (2010c, p.3)

The problem with both Asher's and Nickel's proposals is that there are examples like (8)-(10) with non-eventive predicates — consider again (58) and the following:

(63) Snakes are poisonous.

(64) Mosquitoes are carriers of the West Nile virus.

Both (58) and (63) are intuitively true striking property generics and yet they cannot be construed as a generic capacity interpretation since the capacity interpretation is only available with eventive predicates. With the absence of an event variable in the predicates in (58) and (63), it is not possible for Asher and Nickel to get the *can* into their paraphrases of (58) and (63) — i.e., (63) cannot mean that all snakes normally can be poisonous or that all normal snakes have the capacity to be poisonous, and (58) cannot mean that all drunk drivers normally can be murderers or that all normal drunk drivers have the capacity to be murderers. The same point can be made even more vividly with (64): The predicate in (64) is obviously non-eventive and thus, the capacity interpretation is not available.[14]

[14]Of course, one could argue for a different understanding of the availability of the capacity reading — e.g., an unpronounced modal. I will not discuss such options: The subsequent discussion on diagnostics for the capacity reading omits this option as possibility as well.

If you find yourself unconvinced by examples (58) and (63), then I offer the following two diagnostics as a means to distinguish the capacity reading from the generic reading. If (8)-(10) are generic ascriptions of capacities, then they should pass my proposed tests. Since they do not, I conclude that (8)-(10) are regular generics after all. In other words, the diagnostics reveal that the cases in question, (8)-(10), are not capacity readings, and hence, that they are generic.[15]

2.3.2.1 Diagnostic 1: Actuality Entailments and Capacities

One feature of capacity and ability attributions, like (65a.) and (66a.) below, is that they do not generally entail that the event related to the embedded predicate took place, will take place or is taking place (cf. Bhatt (1999)). Consider the following:

(65) a. Yesterday, it could have rained here. / Yesterday, it was able to rain here.
b. Does not entail: Yesterday, it rained here.
c. It could have rained yesterday, but it didn't.

(66) a. The mailman can deliver the mail. / The mailman is able to deliver the mail.
b. Does not entail: The mailman delivered/will deliver/is delivering the mail.
c. The mailman can deliver the mail, but he didn't/he won't/he isn't.

For example, the ability attribution in (65a.) does not entail (65b.), nor does the conjunction in (65c.) sound contradictory or bad. Similarly, in the case of (66). Given this property of capacity and ability attributions, the following acts as a diagnostic: If a conjunction with (8)-(10) as conjuncts analogous to those in (65c.) and (66c.) sound bad, then this is evidence that (8)-(10) are not capacity readings. Applying the diagnostic, consider:

(67) a. Mosquitoes carry the West Nile virus. # But they don't because they've been vaccinated.
b. Mosquitoes can carry the West Nile virus. But they don't because they've been vaccinated.
c. Mosquitoes generally carry the West Nile virus. # But they generally don't because they've been vaccinated.

[15] Again, my aim here is not to argue that generic capacity readings aren't ever available, rather my aim is to argue that cases like (8)-(10) are not generic capacity readings.

(68) a. Sharks attack bathers. # But they don't because there are shark nets on all beaches.
b. Sharks can attack bathers. But they don't since there are shark nets on all beaches.
c. Generally, sharks attack bathers. # But generally they don't since there are shark nets on all beaches.

The conjunctions in (67b.) and (68b.), where the generics contain an explicit capacity modal, sound fine. By contrast, the conjunctions in (67a.) and (68a.) sound contradictory or at least bad. Thus, given our diagnostic, this is evidence that cases like (8)-(10) are not capacity readings. Further support for the claim that (8)-(10) are interpreted as plain generics is lent by the fact that the contradictoriness or badness exhibited by (67a.) and (68a.) is analogous to the conjunctions in (67c.) and (68c.) where there is an explicit A-quantifier.

2.3.2.2 Diagnostic 2: Conditional Restriction

In recent work, Kratzer (2012) claims that *if*-clauses cannot restrict ability modals. She provides the following illustrative examples (cf. p.107):

(69) If he has a kitchen, he can cook.

(70) If I was taller, I could reach the ceiling.

(69) is not naturally read as saying of someone that he has the ability to cook on the condition that he has a kitchen. Similarly, (70) is not naturally read as saying that I have the ability to reach the ceiling conditional on being taller. These would be quite bizarre capacities. Given examples like (69) and (70), Kratzer concludes that *if*-clauses do not restrict ability modals. But, of course, *if*-clauses can restrict plain generics — consider:

(71) If shops are open on Sundays, then they get customers on Sundays.

(72) If cats jump, then they land on their feet.

(71) is naturally read as saying that shops get customers on Sundays given that they are open on Sundays, and (72) as saying that cats land on their feet conditional on them jumping.

Thus, if Kratzer is correct,[16] then the diagnostic is this: If *if*-clauses can restrict (8)-(10), then (8)-(10) are not capacity readings, rather (8)-(10) are plain generics. Consider the following:

(73) If mosquitoes are found in Africa, they carry the West Nile virus.

(74) If sharks aren't given adequate space, they attack bathers.

(73) is naturally read as saying that mosquitoes carry the West Nile virus conditional on being found in Africa, and (74) as saying that sharks attack bathers conditional on them not having enough space. Hence, *if*-clauses can restrict cases like (8)-(10) and our diagnostic reveals that generics like (8)-(10) are not capacity readings.

2.3.3 Leslie on "Type B Troublesome Generics"

According to Leslie (2008, 2007, n.d., 2013), the truth-conditions for generics like (8)-(10) are as follows:

> *Ks are F* is true if:
>
> > (i) the counter-instances (if any) are negative and;
> >
> > (ii) if F is striking, then some Ks are F and the others are (typically) disposed to be F.

The most pressing problem for Leslie is that there are several counterexamples to her proposed sufficiency conditions — that is, there are cases in which all the conditions specified are satisfied and yet the generic comes out as false. Consider, the following cases which satisfy the conditions outlined above and yet we judge them as false in normal contexts:

(75) Humans kill themselves.

(76) New Yorkers start fires at night.

I take it when (75) and (76) are uttered on their own, they are not intuitively true.[17] Yet killing themselves and starting fires at night are quite striking properties for human beings and New Yorkers to possess, respectively. Moreover, the

[16] Schaffer and Szabo (n.d.) claim in a footnote that Kratzer's conclusion is too strong. They claim that ability modals can be restricted by *if*-clauses and that her data warrants a pragmatic explanation.

[17] Given the right context, they can come out as intuitively true — for example, consider:

counter-instances — i.e., the humans and New Yorkers who don't satisfy the predicated property — are appropriately negative. That is, there is no perceptually salient or striking alternative property which the non-suicidal humans and non-pyro New Yorkers satisfy. Thus, it seems, at least at first blush, that Leslie's conditions are satisfied making (75) and (76) true even though they are intuitively false.

Perhaps, though, the falsity of (75) and (76) stems from the disposition clause of her conditions. Leslie suggests that our primitive cognitive mechanism of generalisation:

> ... looks for a good predictor of the property in question; it avoids generalising to overly broad kinds, or to irrelevant kinds. In particular, for a kind to be the locus of a striking property generalisation, it seems that the members of the kind that lack the property must at least be disposed to have it. It is important, for example, that the virus-free mosquitoes be capable of carrying the virus. If there is no such shared disposition, the generalisation is not made. (2007, p.385)

Are humans and New Yorkers good predictors of the property? Are the relevant disposition attributions true? Amongst living creatures, humans are the best predictor of the property killing themselves — in fact, other animals rarely, if ever, do so:

> ... such phenomenon occurs very rarely. Most animals are prone to the instinct of living, however, in some animals, overwhelmingly the human, this instinct may die, or the reason/s of living may not seem very appealing, especially when a dream that one wished to be true vanishes. Humans, having very developed brains, may bring their lives to an end when they see no point in living any longer. Other non-human animals have less-developed brains and they are happy with simply a meal, while humans need a lot more to be content, since our mental abilities require a lot of distressful stuff that entertains us.

(77) A: What surprising things do humans do?
B: Humans kill themselves.

However, these are not the right sort of contexts for Leslie. Indeed her disquotational account of *Gen* does not permit any context-sensitivity inherent in *Gen*, thus any context-sensitivity arising must be accounted for by Leslie by appeal to other expressions in the sentence or via the contextual aspects of the psychological features she appeals to — not the contextual features which distinguish a null utterance of (75) and (75) as it appears in the context specified.

> ... some grey wolves can indeed intentionally starve themselves to death when their partners have died, since this species of animal mates for life. The domestic dog, which is the descendant of the grey wolf, also has shown suicidal behaviours, but different from those of the wolf — dogs usually starve when their owner passes away. (Yahoo Answers, 2008)

Perhaps amongst humans, there is a subclass which serves as a better predictor — for example, depressed people or deeply troubled people. However, the corresponding generics are not intuitively true in normal contexts either:

(78) Depressed people kill themselves.

(79) Deeply troubled people kill themselves.

Similar arguments hold for (76). Thus, I take it that these examples satisfy the requirement of being a good predictor. What about the disposition requirement?

In what sense are mosquitoes disposed to carry the West Nile virus or sharks disposed to attack bathers? Leslie must construe the disposition condition in an extremely weak way: In no substantive way are all mosquitoes disposed to carry the West Nile virus or all sharks disposed to attack bathers. There are conditions under which they will (e.g., when they are hungry, scared, confused and mistake a swimmer for some- thing fish-like), but there are also such weak conditions under which humans will kill themselves and New Yorkers will start fires at night. What is more, if there weren't such conditions, then it would be impossible for cases like (75) and (76) to be striking or dangerous properties for the kind to possess. Thus, I take it that the disposition clause is not doing the work Leslie wants it to be doing. Examples (75) and (76) are counterexamples to her proposed sufficiency conditions.

Further, as mentioned above, it is hard to see how Leslie's disposition clause will draw the line between (8)-(10) and outwardly false prejudicial generalisations like (15)-(17) — contrasting (8) and (17) is particularly telling: If virus-carrying is dispositional for mosquitoes, then why not for humans.

As Leslie (2013) points out, understanding dispositional mechanisms is very hard, and requires a great deal of empirical knowledge that is not obvious or immediately available. Thus, even if mosquitoes do have such a disposition for carrying the West Nile virus, very few ordinary speakers would possess the relevant knowledge (not even by testimony) and hence, would not be justified in making the corresponding generic claims. Such considerations bolster the case against Leslie's account of (8)-(10).

2.3.4 Pragmatic Proposals

So far, I have eliminated the accounts that purport to show that generics like (8)-(10) are true, but what about pragmatic accounts of (8)-(10)? So far, only one broadly relevant account has been suggested. In this section, I will address the account of Haslanger (2011) and also another *prima facie* plausible pragmatic proposal that I think warrants discussion.

2.3.4.1 Haslanger on Generics and Ideology

Haslanger (2011) claims that the generic generalisations at issue presuppose or implicate an objectionable claim about what the corresponding kind is by nature or intrinsically. It is the presupposed or implicated content that is objectionable, and needn't be the content of the generic generalisation itself which is false or objectionable:

> The general strategy of argument is to show that there is a set of problematic generics that introduce implicitly into the common ground a proposition about a generic essence, about how beef or women or blacks are by nature or intrinsically. These cases are problematic because the introduced proposition is false, so we have a reason to block it. (2011, p.193)

These presuppositions/implicatures are hard to block since they form part of an ideology and/or are hegemonic, according to Haslanger. They enter into the common ground with relative ease: To challenge them would be to challenge something like the ideology of the speaker, and not merely the generalisation expressed. Leslie (2012) offers the following useful illustrative example of Haslanger's view:

> Suppose, for example, that society punishes assertiveness in women to such an extent that they rarely, if ever, are other than perfectly submissive. It would be hard to argue that the generic *women are submissive* is then false, but there nonetheless would seem to be something damaging about asserting it. Haslanger proposes that the damage comes, at least in part, from our tendency to suppose that generic generalisations obtain because of inherent features of the kind in question. That is, even if *women are submissive* is made true by purely external sociological factors, it still may by default communicate that there is something in the nature of women that makes them submissive. That is, it may reinforce essentialist beliefs about women, and

> further communicate that this shared essence causes women to be submissive. Leslie (2012)

Though I've added Haslangers' account here as an alternative proposal, I do not see it as a genuine alternative: It is not at odds with what I am defending in this chapter; her observations could be adopted by various views. The objectionable contents about the essence of kinds are not communicated by generic utterances in general, or even cases like (8)-(10) in general. Further, I see no real reason to see the cases (8)-(10) as ideological or social claims in Haslanger's intended sense. Thus, I take it that the considerations she discusses will not bring to bear on the issues of interpretation and the correctness of our judgments, discussed in this chapter.

2.3.4.2 Overstatement

In communication speakers often overstate, exaggerate or understate the things they say. For example, when a speaker utters *That was the greatest ride in the world!* after getting off an exhilarating ride on a roller-coaster, she does not literally mean that there are no other rides in the entire world that are better than the one she's just been on. She has merely overstated how great the ride was to communicate that she found it very fun and exhilarating. Generics like (8)-(10) could plausibly be treated along the same lines. When a speaker utters a generic like (8), she is overstating how pervasive the West Nile virus is amongst mosquitoes in order to communicate how scary or alarming she finds this property of mosquitoes.

This proposal has at least four issues: First, it has a similar problem to that of Cohen's account. It seems to require that (8)-(10) are uttered with certain intonation or accompanying gestures in order for the intended content to be communicated.

Secondly, this account fails to explain why we find generics like (8), for example, intuitively true *qua* generalisation about mosquitoes. In overstatement we reason from the premise that what is said is false to the true communicated content (i.e., the communicated content that the speakers finds this property striking for mosquitoes).

Thirdly, if (8)-(10) are cases of overstatement, then speakers should be able to achieve the same effect no matter what the subject matter, as long as they exaggerate the claim. Take, for instance:

(80) Dogs live in Bergen.

(81) Norwegians are terrorists.

(80) and (81) should, if the overstatement account is a good one, have a over-generalisation reading that makes them sound like (8)-(10), but (80) and (81) do not have such a reading.

Finally, overstatement as a species of hyperbole would require speakers to have the Theory of Mind or meta-representational capacities necessary to understand this type of hyperbole. This is at odds with current empirical research on the early acquisition and comprehension of generics. If children are to acquire and comprehend generics like (8)-(10) as overstatements, then they would need the relevant Theory of Mind or meta-representational capacities in order to do so. But current research places the acquisition of such meta-representational capacities later than the acquisition of generics (cf. Hollander, Gelman, and Star (2002)).[18]

Given the multitude of objections, I conclude that despite its initially plausibility, the overstatement proposal is not a viable proposal to account for cases like (8)-(10).

2.4 Consequences and Conclusions

This first chapter concerned cases of purported genericity which are true despite only a minority of the kind in question satisfying the predicated property and whose predicated property is somehow *striking*. I argued that such cases are pathological, and hence shouldn't be used as evidence for claims about the nature of genericity or the semantics of generics. In this concluding section, I will outline two important consequences of the arguments given in this chapter: The first is in how we understand the metasemantics and pragmatics of natural language, and the second is in how generics interact with our mind to yield prejudice.

2.4.1 Metasemantics, Pragmatics and Blindness

Cases like (8)-(10) bring up an interesting more general issue about where to locate the effects that cognitive heuristics and biases have on interpretation. For instance, when it appears as though the truth-conditions of sentences depend somehow on cognitive heuristics or biases, such as the strikingness of the predicated property, there are several options for how to incorporate their impact: As part of the content of some expression in the utterance (as in for example, response-dependent terms), as part of the pragmatically derived content, as part of the metasemantics of context-sensitive terms (as in for example, the selection

[18] Note that specific empirical research on cases like (8)-(10) is unavailable; the research I cite here deals with generics more generally.

of a salient quantifier domain or salient referent of *he*), or finally, as no part of the communicated content at all. The latter option is where Blindness stands.

Cases like (8)-(10) are certainly not alone in exhibiting the effects of cognitive features, like strikingness, on interpretation — the question is just how far we let such features determine the content of utterances. We certainly do not want them to have no impact at all on the various levels of content, but we also need to allow for some objective content and indeed, the possibility of error. This question, I think, must be answered on a case by case basis and it would be useful, as theorists, to have some diagnostics and some knowledge of the types of cognitive features to look out for, so to speak.

It is worth pointing out that there are interesting non-generic analogues to cases like (8)-(10), which bring up the same issues quite vividly. It is well-known that probability judgements are subject to certain heuristics and biases — for example, the *representativeness heuristic*, which is the tendency to judge an event or hypothesis as more probable by considering how much the event or hypothesis resembles available data as opposed to using the probability calculus as a basis for the judgement — it would be interesting to see the impact of such features on our accounts of the truth-conditions of sentences containing probability operators. Suppose we have semantic data analogous to the famous Linda-the-bank-the-teller cases of Tversky and Kahneman (1983), and hence, that speakers intuitions about the truth-conditions of sentences containing probability operators inherit somehow the pitfalls of the conjunction fallacy. Should we then automatically trust our intuitions and assume a semantics of probability operators that invalidates such inference patterns as the distribution of probabilities over conjunction, or disjunction introduction of probabilities? The answer, I think, is not obvious, and that is why as theorists we need tests and arguments to answer these types of questions on a case by case basis. In the case of examples like (8)-(10), I have provided the tests, and they indicate Blindness. Another interesting case-study is in the "stakes-sensitivity" of knowledge ascriptions. Nagel (2008), for instance, argues that knowledge ascriptions are invariant despite intuitions to the contrary, and what accounts for these intuitions is a psychological factor called *need for closure*.

As should be clear, this issue is not an issue of generic versus non-generic, it is not an issue of semantics versus pragmatics, and it is not an issue of whether the constituents of the sentence in question are explicit versus implicit.

2.4.2 Generics, Mistaken Generalisation and Prejudice

In Leslie (2013), Leslie sets out to:

> ... identify a particular and pervasive aspect of our thinking which itself leads us down the dark road to prejudice. (2013)

She provides the following account:

> ... our most primitive generalisations, voiced in language as generics, are especially sensitive to information that is particularly striking, horrific, or appalling. When we encounter individuals engaging in such an act, we are naturally inclined to seek to generalise this action to a kind to which the individuals belongs. The correctness conditions of these generalisations require that some members of the kind must indeed have the relevant property, and also that the other members must be typically disposed to have the property. We do not, however, normally have good information about unobservable dispositions available to us, so as a proxy, we generalise the property to a kind that we perceive to have a highly predictive and distinctive essence. (2013)

The examples discussed in this chapter are some of the main evidence used in favour of her account. Leslie (2012) also offers accounts of related social phenomena. Again, some of the main evidence offered are the generics discussed here.

This chapter gives reasons to doubt the plausibility of these accounts. It offers systematic counterexamples to Leslie's attempt to characterise generics like (8)-(10). These undermine her thesis about the mind — the thesis that our minds are furnished with a "primitive cognitive mechanism of generalisation" — and her view about the associated manifestation of prejudice and related social phenomena.

The view proposed in this chapter offers an alternative. It says that there are cognitive biases which can cause us to form mistaken beliefs or to make mistakes about what the contents of our utterances are, and it might be that these mistakes sometimes cause us to act unfairly or in prejudicial ways. Importantly, these cognitive biases are not directly connected to the content of generics, nor do generics enshrine a distinctive form of generalising which is systematically related to any given bias.

I think this alternative picture opens the door to a more general study of the manifestation of error and prejudice in the relevant sense. On the broader

picture I have in mind, Leslie has identified one way in which error and prejudice occurs when generalising with generics, but there are surely others.[19]

[19]In Chapter 4 of this dissertation, I discuss a general metasemantic constraint on the interpretation of *Gen*, which stipulates that there must be some contextually relevant grounds or evidence which helps fix the generalisation expressed by a given generic (cf. Constraint 4 in Section 4.5) — that is, the content of *Gen* is *evidence-sensitive*. The evidence, in a particular context, which is relevant to determining the content and to the evaluation of a generic utterance could be a shared essence or inherent property, a shared property which causes the instances to have the property F, a shared disposition to have the property F, a shared capacity to have the property F, a shared function or purpose which is to F, or whatever sets of propositions adequately and appropriately satisfies *because* clauses.

If generics are evidence-sensitive in this way and our minds often rely on heuristics and biases to form generalisations in the absence of evidence, then there are surely some further intriguing connections between generics, mistaken generalisations and prejudice, and cognitive heuristics and biases. One such heuristic I find particularly intriguing is known as the *explanation effect*. The explanation effect is a tendency of our minds to treat explanation as sufficient for evidence for a particular claim.

Reliance on such a heuristic to use and interpret generics also explains the observation that generics often do not feel as though they are relying on statistical correlations as the basis of the generalisation. Several psychologists study this phenomena — cf. for example, (Lombrozo, 2009), Lombrozo (2006), Brem and Rips (2000), Ramney, Schank, Hoadley, and Neff (1994), Ross, Lepper, and Hubbard (1975) and Anderson, Lepper, and Ross (1995). I find this work intriguing and what I hypothesise is that studying generics from the stand-point of these results in psychology will be fruitful in understanding what role generics play in mistaken generalisation, prejudice and related social phenomena.

3 Pro Gen

3.1 Introduction and Landscape

When speakers intend to communicate *something general* about the world they often use sentences like the following:

(1) Dogs bark.

(2) The tiger has stripes.

(3) Candy is bad for your teeth.

These are examples of generic sentences (or simply *generics*). Generics come in many forms, all of which have non-generic uses as well. What characterises (1)-(3), and other generics, is their intuitive content — they seem to express something general about dogs, tigers and candy, and this something general is meant to be distinctive in that it is relatively stable across different generics and across different contexts. It is far from clear, however, just what this something general is, and by what means generics express it. This paper will be concerned with the latter issue. There are two broad approaches theorists have taken towards this latter issue: The *Quantificational Analysis* and the *Kind-Predication Analysis*.

The Quantificational Analysis: According to the Quantificational Analysis, generics have a logical form akin to sentences containing explicit quantifiers. On this analysis, generics have a tripartite structure consisting of a *quantifier position*, a *restrictor* and a *scope*. Since the quantifier in the case of generics is not explicit, theorists posit a covert, unpronounced quantifier expression they call *Gen*. *Gen*

expresses generality, though there is little agreement amongst theorists about what its meaning is. As a useful heuristic, *Gen* can be taken to mean something akin to *in general* or *all normal*, so that (3) says something like (4) or (5):

(4) In general candy is bad for your teeth.

(5) All normal candy is bad for your teeth.

Though this is only a helpful paraphrase, not a semantic proposal.

The word "quantifier" is ambiguous and is most often used to talk about quantificational determiners (i.e., *D-quantifiers*), an important thing to note is that according to the Quantificational Analysis, at least as I advocate it in this chapter, *Gen* is not a D-quantifier, but rather a quantificational adverb (i.e., an *A-quantifier*).[1] Both syntactically and semantically generics share many more features with A-quantified sentences than with D-quantified sentences.[2] Unfortunately, this point is often overlooked in discussing important syntactic and semantic issues regarding generics. The goal of this paper, in part, is to recast some of these discussions in light of an A-quantificational analysis of *Gen*.

The main advantage of the quantificational analysis is that generics can be said to express genuine *generalisations*. Generics make generalisations over instances of candy, for example, and they are, correspondingly, *about* such instances.

Other advantages include an account of the focus sensitivity of generics — that is, depending on the intonation with which the generic is pronounced or structure of the text or discourse in which it is embedded, the logical form of that generic will be different. I won't go into detail, but the basic idea is that quantificational structure allows the appropriate variety of logical forms which can be the output of the different intonation or information structural properties of the given generic utterance. Moreover, if *Gen* is treated as an A-quantifier, then we have an account of how, like explicit A-quantifiers, the truth-conditions of generics can vary depending on focus — that is, how *Gen* can *associate with focus*.

The quantificational analysis has its fair share of difficulties too. For instance, it has the challenge of specifying a meaning for *Gen* in the face such varied generics as:

(6) Tigers are fierce.

[1] I mean this to include A-quantificational operators that have a modal component to them, in addition to a quantificational component.

[2] For instance, *Gen* and explicit A-quantifiers have similar distributions (i.e., they compose with the same sorts of expressions in the same sorts of ways). We will see this property in action in Section 3.3. In addition, generics are focus sensitive, and appear to *associate with focus* in similar ways to explicit A-quantifiers (cf. Rooth (n.d.) and Beaver and Clark (2008)). For other shared properties, see Cohen (2001b) and the further sections of this chapter.

(7) Philosophers give the Locke lectures.

(8) Dogs have four legs.

It has proved difficult to find a unique characterisation of *Gen* which captures the intuitive meaning across the full range of generics, like (6)-(8). If (6)-(8) express generalisations, then they express very different generalisations. To get a sense of this, consider the paraphrases in (9)-(11) which reflect common judgments about what (6)-(8) communicate on their most salient readings:

(9) In general, tigers are disposed to be fierce.

(10) Traditionally, only philosophers give the Locke lectures.

(11) All dogs without defects have four legs.

This makes it hard to see what sort of meaning to assign *Gen* which is invariant across all these seemingly different generics. As we will see, the kind-predication analysis, on the other hand, can simply locate these difficulties elsewhere (e.g., the metaphysics of kinds or features of our semantic cognition).

This contrasts with the other analysis theorists have offered — the Kind-Predication Analysis.

The Kind-Predication Analysis: According to the Kind-Predication Analysis, generics have a logical form akin to direct kind-predications[3] — that is, a simple *subject-predicate*, or *noun phrase-verb phrase*, form.[4] On this view, generics like (1)-(8) are taken to express genuine kind-predications like those in (12) and (13):

(12) Dinosaurs are extinct.

(13) Mosquitoes are widespread

[3]This view was originally defended by Carlson (n.d.). More recently, Liebesman (2011) and Cohen (2012) have offered new kind-predication proposals.

[4]A more general approach to the kind-predication analysis would just be a *Predication Analysis* where the predications have certain properties which make them generic, like non-episodic group reference. As far as I am aware this has not been proposed, perhaps because it is not motivated, as the kind-predication analysis is, by independent evidence like the use of bare plurals as referring to kinds more generally. Though I am in no way advocating this strategy, it would perhaps deal with some of the original worries of Carlson (1989) about the existence of generic readings without a distinctively kind referring component of the logical form.

On this analysis, there is no expression or component of logical form which is responsible for expressing a generalisation,[5] rather the something general is expressed by virtue of the kind reference itself or some other communicative mechanism.

An obvious advantage of the Kind-Predication Analysis is that it has a readily available explanation of why *Gen* is never pronounced — it is never pronounced since it doesn't exist! But an equally obvious disadvantage of this analysis is that one has no readily available explanation of why generics seem to be *about* instances of the kind and not the kind itself: Intuitively, (3) is about individual instances of candy, not about the kind candy. One way to make this disadvantage apparent is to observe that dialogues of the following sort seem unacceptable:

(14) A: Dogs have four legs.
B: ?? That's absurd; the kind dog doesn't have four legs.

(15) A: Dogs bark.
B: ?? That's absurd; the kind dog doesn't bark.

The Kind-Predication Analysis takes the literal (or semantic) content of generics like (14A)-(15A) to be kind-predications, but as these types of dialogues reveal, this content seems entirely unavailable to ordinary speakers: The kind-predication analysis predicts that (14B)-(15B) should be fine since the literal content of the generic uttered by A should obviously, by the lights of the kind-predication analysis, be an interpretation which is available to B.

An additional worry with the kind-predication analysis is that many generics seem to have subjects which are not kind referring. The indefinite singular, for example, often yields generic readings as in (16a.), but cannot be used in direct kind-predications as in (16b.):

(16) a. A dodo has wings.
b. * A dodo is extinct.

Despite these difficulties, however, there is plenty of linguistic evidence which supports treating generics as kind-predications (cf. Carlson (n.d.), Liebesman (2011)). These will be considered in detail in Section 3.3.

In this chapter, I will argue for the Quantificational Analysis, in particular, by defending the original argument of Carlson (1989) for the existence of *Gen*,

[5]Carlson (n.d.) argued that there is an unpronounced verb-phrase operator, *Gn*, which was part of the logical form of generics. Cohen (2012) thinks *Gen* is responsible for communicating the general content, but that it is pragmatically derived. Liebesman (2011) is not entirely explicit about the matter, but seems to think that the kind-term itself communicates the general content, perhaps together with some pragmatic factors.

which has recently come under scrutiny — the *multiple readings argument*. In addition, I will provide a novel argument in favour of *Gen* (Section 3.2). In Section 3.3, I will respond to various objections to the Quantificational Analysis, in particular to some of the linguistic evidence offered in favour of the Kind-Predication Analysis.

This paper comes at an important juncture since there has been a recent upswing in the endorsement of the Kind-Predication Analysis with novel arguments offered in its favour. My goal is to argue that these novel arguments do not warrant abandoning the Quantificational Analysis which has been, more or less, the standard view, especially amongst linguists, since Carlson's paper which presents the multiple readings argument.

It is worth emphasising the importance of settling the question of which analysis to give. The analysis we endorse, not only has consequences in linguistics and philosophy of language, but also for a wide range of philosophical and worldly endeavours:

- Ceteris paribus laws of the special sciences are arguably expressed by generics (Nickel, 2010a);
- as are moral generalisations (Varynen, 2004);
- as are disposition attributions (Fara, 2005);
- as are premises in thought experiments (Ichikawa & Jarvis, 2009);
- as are prejudice generalisations (Leslie, 2013);
- as are misleading summary statements of statistical facts in scientific discourse and media (Liberman, 2009);
- as are laws in our legal systems (Sterken (ms.) pace Holton (2011)).

3.2 Pro Gen

3.2.1 Multiple Readings: The Classic Argument in favour of Gen

The observation which swung theorists towards a consensus post Carlson (1989), to the effect that generics have a tripartite structure including *Gen*, is that some ambiguous sentences appear to have more than one salient generic interpretation. The most discussed example is (17a.) which is thought to be ambiguous between the two readings, paraphrased in (17b.) and (17c.) (cf. also Milsark (1974)):

(17) a. Typhoons arise in this part of the Pacific.
b. Typhoons are such that they (generally) arise in this part of the Pacific.
c. This part of the Pacific is (generally) such that typhoons arise in it.

(17a.) is ambiguous since there are contexts in which the two readings (17b.) and (17c.) are truth-conditionally distinct. The (17b.) reading says of typhoons that they generally have a common origin, namely the part of the Pacific under discussion. The truth of this reading requires that at least a majority of, or possibly even all, typhoons originate in the geographic area in question. (17c.), on the other hand, says something like characteristically or generally some typhoon or other arises in the part of the Pacific under discussion: This reading asserts a general property of the geographical area; the subject term *typhoons* has more of an existential force. Unlike (17b.), the truth of (17c.) is compatible with there being several geographical areas where typhoons generally originate.

The classic argument in favour of *Gen* is that the quantification theorist has an easy time explaining the ambiguity present in sentences like (17a.), whereas the kind-predication theorist does not (if he can at all).

Since *Gen* is a quantifier expression according to the quantification theorist, she can explain the two readings, (17b.) and (17c.) by appeal to independently motivated properties of sentences involving determiners or quantificational adverbs more generally. According to the quantification theorist, the ambiguity exhibited by (17a.) is a familiar sort of structural ambiguity exhibited by sentences involving quantifier expressions. Such sentences are often times ambiguous between different readings which result from different material being mapped into the restrictor and scope of the tripartite structure. Consider, for example:

(18) a. Typhoons typically arise in this part of the Pacific.
b. Typhoons are such that they typically arise in this part of the Pacific.
c. This part of the Pacific is typically such that typhoons arise in it.

(18a.), a sentence involving the quantificational adverb *typically*, is thought to be structurally ambiguous between the two readings in (18b.) and (18c.) as a result of, in the former case, *typhoons* being mapped to the restrictor of the tripartite structure, and in the latter case, *typhoons* being mapped to the scope of the tripartite structure. The different mappings result in the following distinct logical forms being assigned to (18b.) and (18c.) respectively:

(19) Typically x [Typhoon(x)] [Arise-in-this-Part-of-the-Pacific(x)]

(20) Typically x [This-Part-of-the-Pacific(x)] $\exists y$[(Typhoons(y) $\wedge$ Arise-in(y, x))]

With this observation in tow, the quantification theorist goes on to observe that (18a.) is ambiguous between the two readings (18b.) and (18c.) in an analogous way as (17a.) is ambiguous between the two readings in (17b.) and (17c.). In (18b.) typhoons is given the force of *typically* and in (18c.) *typhoons* is given what seems to be existential force: (18c.), but not (18b.), is compatible with there being several geographic areas where typhoons typically arise.

If this analogy withholds scrutiny, then there is good reason to think the ambiguity exhibited by (17a.) is the same sort of structural ambiguity; and hence, the quantification theorist has a readily available explanation of the two interpretations. While the (17b.) reading is a result of *typhoons* being mapped to the restrictor and the remaining material being mapped to the scope; the (17c.) reading is a result of *this part of the pacific* being mapped to the restrictor and the remaining material being mapped to the scope. Further, since tripartite structures have quantifier positions, this is evidence that absent an overt quantifier expression, the unpronounced quantifier expression, *Gen*, is what fills this position. The quantification theorist, thus, can claim that the readings in (17b.) and (17c.) are assigned the following logical forms respectively:

(21) Gen x [Typhoon(x)] [Arise-in-this-Part-of-the-Pacific(x)]

(22) Gen x [This-Part-of-the-Pacific(x)] [$\exists y$(Typhoons(y) $\wedge$ Arise-in(y, x))]

The Kind-Predication Analysis, on the other hand, doesn't seem to fare as well. On the face of it, treating (17a.) as a kind-predication cannot do the same explanatory work. After all, at least initially it seems that the proponent of kind-predication can only provide (17a.) with a single logical form (which results in a generic interpretation), the kind-predication glossed in (23), which corresponds to the (17b.) interpretation:

(23) λx. Arise-in-this-Part-of-the-Pacific(x)(typhoons)

The challenge for the proponent of kind-predication, then, is to account for the second reading of (17a.) — i.e., the (17c.) reading.

At this point the kind-predication theorist has two options: He can deny that (17c.) involves genericity at all, or he can maintain that (17c.) is generic and explain the ambiguity on other grounds. Liebesman (2011) has attempted the former, and Koslilcki (1999) has proposed the latter as a potential solution. I will respond to both proposals in turn.

3.2.1.1 Liebesman on Multiple Readings

Liebesman (2011) has responded to the multiple readings challenge by claiming that (17c.) is not a generic — he claims it is an existentially quantified sentence. This response, he says, is well-motivated since most theorists agree that bare plurals are ambiguous between a generic reading and an existential (or indefinite) reading — as in example (24) which has the existential reading glossed in (25).

(24) Tigers are on the front lawn.

(25) Some tigers are on the front lawn.

(24) is thought to have an existential reading since (24) seems true in contexts where only a few tigers are on the said front lawn, whereas the generic reading would be false. The existential reading is introduced to account for the intuition that there is a true reading of (24) in such contexts.

Liebesman's approach to cases like (17a.), thus, is to explain the ambiguity on lexical grounds, rather than structural grounds. In particular, he argues that the two readings of (17a.) are explained by the ambiguity present in the bare plural *typhoons*.

Having noted the possibility of (17a.) having an existential interpretation, Liebesman goes on to argue that the (17c.) reading of (17a.) is in fact an existentially quantified sentence and not a generic one. This makes it possible for the proponent of kind-predication to predict the two interpretations and assign two logical forms — the one which they claim is generic, whose logical form is glossed in (23) above, and the one they claim is non-generic, whose reading is glossed in (26) and whose logical form is glossed in (27):

(26) Some typhoons arise in this part of the Pacific.

(27) $\exists x(\text{Typhoon}(x) \wedge \text{Arise-in-this-Part-of-the-Pacific}(x))$

Liebesman claims that some evidence that (17c.) is interpreted existentially is that it passes a standard test used to identify existentially interpreted noun phrases.[6] The test concerns whether we can replace the noun phrase in question, in a truth-preserving manner, with any noun phrase whose extension is more inclusive than the original noun phrase. Linguistic contexts in which such replacements are permitted are called *upward-entailing contexts*. Existentially interpreted noun phrases are upward-entailing. In (24), for example, we can replace *tigers* with more inclusive bare plurals like *animals* or *things*, and the corresponding

[6] Cf. Krifka et al. (1995, pp.13-4).

sentences are still true. By contrast, generically interpreted noun phrases are not generally upward entailing. Consider the generic (3), for example, we cannot replace *candy* with *food* or *things* and at the same time preserve truth. Applying this test to (17a.), Liebesman claims that (17a.) is existentially interpreted since it is upward-entailing — as evidenced by the fact that (28) and (29) are both true:

(28) Storms arise in this part of the Pacific.

(29) Heavy storms arise in this part of the Pacific.

The main issue with Liebesman's proposed account of cases like (17a.) is that there are cases just like (17a.), that exhibit the same ambiguity, which are not upward-entailing on the relevant reading (i.e., the analogue of (17c.)). Consider for example:

(30) a. Murderers roam the streets in this part of town.
b. Murderers are such that they (generally) roam the streets in this part of town.
c. This part of the town is (generally) such that murderers roam the streets.
d. People roam the streets in this part of town.
e. Mammals roam the streets in this part of town.

(31) a. Vicious dogs protect Ft. Knox gold.
b. Vicious dogs are such that they (generally) protect Ft. Knox gold.
c. Ft. Knox gold is (generally) such that vicious dogs protect it.
d. Mammals protect Ft. Knox gold.
e. Animals protect Ft. Knox gold.

(30a.) and (31a.) both exhibit the same ambiguity that is present in (17a.) — i.e., they are ambiguous between the two readings in (30b.) and (30c.), and (31b.) and (31c.) respectively. But (30a.) and (31a.) are not upward entailing on either reading: all the relevant readings of (30d.) and (30e.), and (31d.) and (31e.) are intuitively false despite the bare plurals in question being more inclusive substitutions.

Note that the fact that one might, in addition, hear intuitively true readings of (30d.), (30e.), (31d.) and (31e.) is not what is at issue here. What I am pointing out is that there at least two intuitively false ones: The quantificational theorist is equally committed to the lexical ambiguity of bare plurals that Liebesman appeals to (i.e., the existence of an existential interpretation of (17a.)). The quantificational theorist is claiming that there are at least two generic readings, whether

or not there is in addition an existential reading does not put her position in jeopardy.

Since (30a.) and (31a.) have at least two readings on which they are false, they are not upward-entailing on the relevant readings. Moreover, since all existentially interpreted sentences must be upward-entailing, such an analysis (i.e., treating the (c.) readings of the sentences in question as existentially interpreted) is not open to the kind-predication theorist in general. Treating the cases as generically interpreted, on the other hand, does not preclude the appearance of upward-entailingness in particular cases — generic readings can have the appearance of being upward-entailing, in such cases the facts just happen to be that way. As such, the quantificational theorist can claim that the apparent upward-entailingness of (17a.) is a mere accident of the example chosen: The relevant generic readings of (28) and (29) just happen to be intuitively true. Thus, the (17c.) reading of (17a.) in question should not be treated as existentially interpreted since this analysis cannot be extended to provide a fully general explanation of the ambiguity data at issue.

Still further, there is other evidence that the (17c.) reading of (17a.) should be treated as generically interpreted. Accompanying his original multiple readings argument, Carlson considers additional evidence for the generic interpretation:

> But is this reading in fact a generic reading of [(17a.)]? By all tests, it appears to be. Consider whether it expresses a regularity (it does); whether it is epistemologically determinate with regard to the present moment (it is not); whether it is stative (it is); whether it is based on a nongeneric (it is; there is an eventive reading, more salient in the past tense — *Hurricanes arose in that part of the Pacific* (so all boats avoided the area)). It is also intensional; consider substituting the phrase *where I am pointing my finger* for the phrase *this part of the Pacific*. The intersubstitution is not automatically licensed even if where I am pointing is the part of the Pacific I am talking about. (1989, p.171)

Thus, Carlson provides at least five additional reasons for treating (17c.) as a generic.

Finally, (17a.) fails another test used to preclude existentially interpreted noun phrases (Leslie, n.d.):

> The second test concerns interactions with adverbs of quantification: We cannot insert an adverb of quantification into a sentence receiving the existential interpretation without inducing a significant change

> in meaning. If a sentence receives the generic interpretation, however, then we can. *Usually* is a good candidate to be inserted into the first of the pairs above: *tigers are striped* and *tigers are usually striped* are very close in meaning. Some "troublesome" generics — to be discussed at length in coming chapters — become false with the insertion of *usually*. However, in those cases we can insert *sometimes* with a minimal change of meaning. *Tigers are sometimes on the front lawn*, however, has a distinctly different meaning from [24]. If we cannot insert any adverb of quantification into the sentence without changing the meaning significantly, this suggests the statement is not a generic. (n.d., p.7)

Note that it is important that we apply the test to locatives, as in the case of (17a.) and (24), as the test won't work for habitually interpreted predicates.[7] Applying the test and inserting *sometimes* into (17a.), we have:

(34) a. Typhoons sometimes arise in this part of the Pacific.
b. Sometimes typhoons arise in this part of the Pacific.

(26) Some typhoons arise in this part of the Pacific.

It is clear that both insertions of *sometimes* yield distinctly different meanings from (26). This suggests that (17c.) is a generic interpretation.

I now move on to Koslicki's proposal on behalf of the kind-predication theorist.

3.2.1.2 Koslicki on Multiple Readings

Koslilcki (1999) has another reply to the multiple readings argument on behalf of the kind-predication theorist. Her approach is to claim that a kind-predication analysis of (17a.) can account for both generic readings by means of structural ambiguity:

> While the difficulties raised by [(17a.)] are no doubt serious, they do not seem to me to present knock-down arguments against the subject-

[7] To see that the test doesn't work in general, take the existential reading of:

(32) A hedgehog visits my garden.

Adding an adverb does not significantly change the interpretation:

(33) A hedgehog sometimes visits my garden.

> predicate analysis of genericity. One possibility would be, for example, to account for the difference between [(17b.)] and [(17c.)] along the lines of [(35)] and [(36)]:
>
> (35) Arise-in-this-part-of-the-Pacific (hurricanes)
>
> (36) Hurricanes-arise-in-it (this part of the Pacific)
>
> These paraphrases suggest that the difference between [(17b.)] and [(17c.)] depends on which part of the sentence is to be analysed as the subject and which as the predicate. (1999, p.446)

The idea, then, is that the kind-predication theorist can account for both generic readings of (17a.) by saying that both readings are kind-predications — in the case of (17b.) the bare-plural *hurricanes* is interpreted as a kind, while in the case of (17c.), *this part of the Pacific* is treated as a kind-term.

The problem with Koslicki's proposal is that it is hard to motivate with independent evidence. In particular, direct kind-predications do not exhibit the same type of ambiguity — (12) and (13) are not ambiguous between the two readings in (37) and the two readings in (38) respectively.

(37) a. Dinosaurs are extinct.
b. Extinct things are dinosaurs.

(38) a. Mosquitoes are widespread.
b. Widespread things are mosquitoes.

If the kind-predications in (37a.) and (38a.) were ambiguous as in Koslicki's proposal, then they would expect each of them to have two kind-predication readings: One with its straightforward interpretation as in (37a.) and (38a.) and another where the other parts of sentence — i.e., the predicates *extinct* and *widespread* respectively — are treated as the kind-term, and yield the interpretations in (37b.) and (38b.). But (37a.) and (38a.) are quite obviously not ambiguous in the respect Koslicki proposes on behalf of the kind-predication theorist. Thus, I take it that this proposal fails.[8]

[8] If one insists are preserving the locative element, then the following examples I contend suffice to make the same point:

(39) a. Elephants are extinct in this part of Africa.
b. λx. Extinct-in-this-part-of-Africa(x)(elephants)
c. λx. Elephants-are-Extinct-in-it(x)(this part of Africa)
d. Elephants are extinct in this part of Africa.
e. This part of Africa is such that elephants are extinct in it.

Having demonstrated that the different options available to the kind-predication theorist are not adequate, I conclude that the multiple readings argument withholds scrutiny — we still have good reason to endorse a Quantificational Analysis of generics. Next, I will offer a novel argument in favour of *Gen*, to the effect that if one endorses covert structure more generally (which is a fairly standard endorsement in linguistics and philosophical circles), then one should endorse *Gen* as well.

3.2.2 Binding: The Binder Argument

In considering arguments in favour of the existence of *Gen*, it will be worthwhile to consider why the early Carlson (n.d., p.161), who endorsed kind-predication, felt the introduction of a predicate modifier *Gn* was needed as a necessary semantic addition to simple kind-predication. Carlson (n.d.) noted that simple kind-predication is met by difficulties when faced with generics such as (41a.) and (42a.):

(41) a. Cats clean themselves.
b. λx. Clean(x, x) (cat)
c. Gn$(\lambda x$. Clean$(x, x))$(cat)

(42) a. Goldfish like everyone who likes them.
b λx. [x likes everyone who likes x](goldfish).
c. Gn$(\lambda x$. [x likes everyone who likes x])(goldfish)

(41a.) and (42a.) both have readings where individual coreference is required to yield the given interpretation: In (41a.), the relevant reading is that each cat cleans itself, and in (42a.), the relevant reading is that each individual goldfish has the property of liking everyone who likes itself. A simple kind predication logical form for (41a.) like that given in (41b.) does not capture the appropriate truth-conditions for the desired readings since it says that the kind cat cleans the kind cat, and not that each member of the kind cat cleans itself. To deal with this, Carlson introduces *Gn* to take the object-level predicate to a kind-level predicate. He formalises (41a.) as (42c.), which is to be understood as true iff the object-level predicate λx. Clean(x, x) holds with sufficient regularity over the

(40) a. Mosquitoes are widespread in this part of Africa.
b. λx. Widespread-in-this-part-of-Africa(x)(mosquitoes)
c. λx. Mosquitoes-are-Widespread-in-it(x)(this part of Africa)
d. Mosquitoes are widespread in this part of Africa.
e. This part of Africa is such that mosquitoes are widespread in it.

realisations of the kind. (42a.) is handled in a similar manner, in this case *Gn* scopes over the lower quantifier.[9]

The early Carlson used such examples as evidence for the presence of a predicate modifier, *Gn*, in the logical form of generics. Note, however, that this type of data is equally good, perhaps even better, evidence for the presence of a quantifier expression, *Gen*, in the logical form of generics. In the rest of this section, I will argue that examples like (41a.) and (42a.) are evidence of binding and therefore, evidence for the existence of *Gen*.

One of the, by now, classic arguments for the existence of unpronounced syntactic structure is the so-called *binding argument* (cf. Stanley and Szabo (2000), Fintel (1997)). The binding argument proceeds from evidence of bindability to the existence of an unpronounced domain variable in the logical form of quantified sentences. There is a similar argument for the existence of a quantifier expression in the logical form of generics, except that this argument proceeds from evidence of bindability to the existence of an unpronounced quantifier expression — what I will call *the binder argument*. While the bind*ing* argument argues from bindability to the existence of an unpronounced variable which gets bound by an explicit quantifier, the bind*er* argument argues from bindability to the existence of an unpronounced binder which does the binding.

Some of the most telling evidence of bindability is given by generics like (41a.)-(42a.) and (43)-(44):

(43) Good students answer every question on their exam.

(44) Wine shops keep most bottles in the stockroom.

(43) expresses something close in meaning to *in general, if x is a good student, then x answers every question on x's exam.* Similarly, (44) expresses something like *in general, if x is a wine shop, then x keeps most bottles (in x) in the stockroom of x.* What is important to notice about examples (42a.) and (43)-(44) is that there is a variable under the scope of the lower quantifier in each case, *everyone that likes them* in (42a.), *every question on their exam* in (43) and *most bottles in the stockroom* in (44), as evidenced by the pronouns *them* and *their*; and crucially, the interpretation of this variable varies as a function of individuals introduced by the bare plurals *goldfish, good students* and *wine shops* respectively.

Similar examples, where variability is evidenced by a variability in the interpretation of a lower quantifier, are used as evidence of bindability in the binding argument:

[9] Cf. Carlson (n.d., pp.161-2) for three additional reasons for treating the pronouns in question as bound.

> ... quantifier domain restriction is due to the presence of domain variables in the actual syntactic structure of sentences containing quantified noun phrases. But syntactic structure cannot simply be postulated on semantic grounds. Rather, evidence of a syntactic sort must be available for the existence of domain variables. The main source of syntactic evidence comes from the fact that domain variables interact in binding relations with quantifiers. Here is the evidence from bindability. Consider the sentence:
>
> (45) Everyone answered every question.
>
> [(45)] can express the proposition that everyone x answered every question on x's exam. What this indicates is that there is a variable accessible to binding somewhere in the quantified phrase *every student*. ... On the assumption that binding is fundamentally a syntactic phenomenon, such examples provide evidence for a variable somewhere in the syntactic structure of quantified noun phrases. (Stanley, 2002b, p.369)

If examples like (45) are taken as evidence of domain variables in quantified sentences (Stanley and Szabo (2000), Fintel (n.d.), Partee (1989), among others), then examples like (42a.), (43) and (44) should be taken as evidence of a binder. The analogy with the binder argument is strong: We have precisely analogous evidence of a binding relation (or variability) — i.e., of how the interpretation of a lower quantifier phrase varies as a function of the the expressions found in higher positions. In the case of the binding argument this is evidence for a hidden variable in the lower quantifier phrase, in the binder argument this is evidence for a hidden higher quantifier expression, a binder.

If we countenance the existence of a quantifier expression in the logical forms of (42a.),(43) and (44), then we have an explanation for why the (42a.) and (46), (43) and (47), and (44) and (48) have similar interpretations respectively:

(46) Typically goldfish like everyone who likes them.

(47) Generally good students answer every question on their exam.

(48) Every wine shop keeps most bottles in the stockroom.

In a similar fashion, in (46)-(48) where there is an explicit higher quantifier, the interpretation of the variable varies as a function of the individuals introduced by the higher quantifier phrase, *typically gold fish* in (46), *generally good students*

in (47) and *every wine shop* in (48). (42a.), (43) and (44) have similar interpretations to (46)-(48) respectively, because they have similar logical forms, as exemplified by the logical form schemas in (16) and (17):

(49) Gen x [Restrictor(x)] [LQ y [Restrictor(x, y)] [Scope(x, y)]]

(50) HQ x [Restrictor(x)] [LQ y [Restrictor(x, y)] [Scope(x, y)]]

The natural explanation of the evidence of bindability exhibited by (42a.), (43) and (44), thus, is that there is an unpronounced quantifier expression which acts as the higher quantifier. In order to account for the binding facts exhibited by (42a.), (43) and (44), an unpronounced quantifier expression — i.e., a binder — is needed. *Gen* is part of the logical form of generics.

Related evidence of bindability comes from examples of generics containing gradable adjectives,[10] like the following:

(51) Europeans are tall.

(52) Sports teams have members that are old.

On some views, the interpretation of gradable adjectives involves the provision of a *standard* or *comparison class*. In Stanley (2002b) and Ludlow (1989), examples like (53) and (54) below are used as evidence that gradable adjectives contain a variable which takes as its value a standard (or comparison class) in the logical form of sentences containing gradable adjectives:

(53) Everyone in my family is tall.

(54) a. In general, my family members are tall.
b. My family members are generally tall.

In order for (53) to be true, it needn't be the case that everyone in my family has to meet some single contextually given standard of tallness; rather in order for (53) to be true it must be that everyone in my family is tall for the standard of tallness relevant to each individual member. My father must be tall by the standard for older adult males, my niece tall for female teenagers, etc. What accounts for this in the case of quantified sentences containing gradable adjectives (roughly) is that the quantifier phrase is binding a comparison class variable in the gradable predicate.

[10] Cf. Nickel (2012b).

Just like (53)-(54), in each of (51)-(52), the most salient reading is one in which the standard for the gradable adjectives can vary as a function of the individuals introduced by the bare plurals. For example, in order for (51) to be true, it needn't be the case that Europeans in general meet some single contextually supplied standard of tallness; rather (18) is true just in case (roughly) in general, if x is a European, then x is tall by the standard of tallness relevant to x. The adult male Europeans are tall by the standards for adult males, the female teenaged Europeans are tall by the standards for female teenagers, etc. Similarly, in the case of (52): Its truth in the envisioned scenario merely requires that in general, if y is a sports team, then there are members of y that are old by the standard for that sport.

If we countenance the existence of a quantifier expression in the logical forms of (51)-(52), then we have an explanation of why generics like (51)-(52) and quantified sentences containing gradable adjectives like (53)-(54) get analogous interpretations. They are analogous because they have similar logical forms, as exemplified by (55) and (56):

(55) Gen x [European(x)] [$\text{Deg}_{Height}(x) \leq f_i(x)$]

(56) $\forall x$ [Member-of-my-Family(x)] [$\text{Deg}_{Height}(x) \leq f_i(x)$]

Thus, a natural explanation of examples like (51)-(52) is that a comparison class variable in the logical form of generics containing gradable predicates is bound by some unpronounced quantifier expression — i.e., generics containing gradable adjectives are further evidence for the existence of a binder in the logical form of generics.

Next, I go on to defend the binder argument from objections made by its opponents: The advocates of kind-predication and opponents of the bind*ing* argument.

3.2.3 Objections and Replies: Binding, Kind-Predication and Truth-Conditional Pragmatics

Can the kind-predication analysis or opponents to the binding argument (i.e., advocates of *truth-conditional pragmatics*) account for the binding phenomena of the previous section? Here, I will argue that they have no plausible story to account for the binding data.

3.2.3.1 Objection 1: Kind-Predication and Binding

The Kind-Predication Analysis almost certainly cannot be extended to account for the binding phenomena (or for that matter, the other data I present in the Appendix). Consider again some of the evidence of bindability presented with the binder argument above:

(41) a. Cats clean themselves.

(43) Good students answer every question on their exam.

On the reading under consideration, (41a.) says something similar to *in general, if x is a cat, then x cleans x*. Here, the evidence of bindability is that the denotation of *themselves* varies as a function of the individuals introduced by *cats*. On the relevant reading of (43), it says something like *in general, if x is a good student, then x answers every question on x's exam.* In the case of (43), the evidence of bindability is that the variable present in the lower quantifier expression varies its interpretation as a function of individuals introduced by *good students*.

The question is whether the linguistic resources that Liebesman (2011) has at his disposal (kind predication, the existential reading of the bare plural and other semantic or pragmatic properties of the given sentence) can simultaneously account for bindability and the readings we are after for (41a.) and (43). I will consider two options for basic logical forms of (41a.) and (43), which I take as exhaustive given Liebesman's limited linguistic resources: On option one, the bare plurals in (41a.) and (43) are interpreted as kind-denoting and hence, (41a.) and (43) are of subject-predicate form. On option two, the bare plurals are interpreted as indefinites and so (41a.) and (43) are (loosely speaking) existentially quantified sentences.

Option one: Option one, assigns (41a.) and (43) a subject-predicate form and sees the subject terms as kind-denoting. An immediate issue arises, however: How should the seemingly bound pronouns, *themselves* in (41a.) and *their* in (43), be interpreted? They need to be interpreted in such a way that accounts for the variability. But on a kind-predication story, they cannot be interpreted as bound. Liebesman does not endorse Carlson's *Gn*, and even if he did, we already saw (in section I) why *Gn* doesn't do justice to other existing linguistic evidence; so Liebesman cannot appeal to anything like *Gn*.

If *cats* and *good students* are akin to names for kinds, then how can the interpretation of themselves in (41a.) and their in (43) vary as a function of the members of the kinds? If the pronouns in (41a.) and (43) are not bound, then

they must be interpreted in some other way. The most plausible way to go, modulo Carlson's *Gn* strategy, seems to be that the variable gets an interpretation on which it is kind-denoting:

(57) λx. Clean(x, x)(cat)

(58) λx.[x answers every question on x's exam](good students)

It seems plausible that Liebesman would endorse such a logical form for (41a.) and (43): There is no reason to think (41) and (43) are not generics and he thinks generics are kind-predications. The trouble with this proposal for the logical forms of (41a.) and (43) arises, however, when considering what sort of metaphysical story he would have to tell in order for (41a.) and (43) to be generics on his theory.

On Liebesman's theory a special metaphysical notion of *inheritance* accounts for the sense of variability in the case of simple generics and so he seems committed to there being a similar metaphysical story to tell, which can account for the sense of variability present in (41a.) and (43) — i.e., for complex generics as well. Indeed, inheritance is the central notion in Liebesman's theory of generics:

> ... kinds can inherit properties from their members in much the same way that ordinary objects inherit properties from their parts ...
>
> Generics, then, do form something of a natural class. It is necessary for a sentence to be a generic that it is a kind-predication that ascribes a property that can be inherited from the kind's members. (2011, p.421)

The properties expressed by the predicates in (57) and (58) are not, as is, the appropriate properties for inheritance. If (41) is to have its intended generic interpretation, then Liebesman's story about inheritance needs to be that the kind cat inherits the object-level property of cleaning itself. But instead the kind cat inherits the property of cleaning the kind cat. This latter property is an object level property, but not the intended one. Things get worse with (43) since the property doesn't even seem like a genuine object-level property: The property Liebesman is committed to if he assigns (43) the logical form in (58) is the property of answering every question on the kind good students' exam — this does not seem to be a property of members of the kind good students.

Thus, when (41a.) and (43) are given a subject-predicate form, one cannot account for the sense that the pronouns in (41a.) and (43) are bound and still

preserve their intended interpretation. Consequently, (41a.) and (43) are not kind-predications and hence, not generics by Liebesman's account. The proponent of Liebesman's view will need to consider a different option as the basic logical form of (41a.) and (43).

Option two: Option two treats (41a.) and (43) as, loosely speaking, existentially quantified sentences since the bare plurals in each are assigned their indefinite interpretations on this proposal. This option is *prima facie* appealing since it at least provides a binder in the logical forms of (41a.) and (43) that binds the pronouns *themselves* in (41a.) and *them* in (43). On this interpretation of the bare plural, *themselves* in (41a.) and *them* in (43) are bound by means of existential closure, and (41a.) and (43) receive the logical forms in (59) and (60):

(59) $\exists x(\text{Cat}(x) \wedge \text{Clean}(x, x))$

(60) $\exists x(\text{GoodStudent}(x) \wedge \forall y(\text{Question-on-Exam-of}(y, x) \rightarrow \text{Answer}(y, x)))$

The first thing to note at this point is that now (41a.) and (43) assert something much weaker than they appear to — they say something like:

(61) Some cats clean themselves.

(62) Some good students answer every question on their exam.

With such a weak semantic interpretation of (41a.) and (43), we must rely on pragmatics to somehow supplement the content of (41a.) and (43) to yield a stronger interpretation, that at least resembles the generic force apparent in (61) and (62). However, I don't see how that can be done. Liebesman appeals to the Gricean maxim of relevance in discussing the classic multiple readings example, (17a.); he claims:

> Importantly speaker 2's utterance of [(17a.)] also has pragmatic properties that may explain why the utterance feels generic or law-like. Specifically, speaker 2's utterance of [(17a.)] would violate the maxim of relevance if it didn't communicate something interestingly non-accidental about the part of the Pacific under discussion. The obvious relevant fact is that the part of the Pacific is hospitable to tornadoes. (Liebesman, 2011, pp.427-8)

But this pragmatic strategy clearly won't work in general: Why should we expect all generics with a free variable under a lower quantifier to communicate something "interestingly non-accidental" about their subjects? This seems implausible.

The second major problem for this view, as argued above in relation to the multiple readings argument, is that (41a.) and (43) should be upward-entailing in their subject position since they are existentially quantified sentences. This is problematic since clearly (41a.) and (43) are not upward entailing in subject-position — consider:

(63) a. Mammals clean themselves.
b. Animals clean themselves.

(64) a. Students answer every question on their exam.
b. People who are good at something answer every question on their exam.

(41a.) clearly does not entail (63a.), nor (63b.); and similarly, (43) does not entail (64a.), nor (64b.). Thus, any approach on which the bare plurals in (41a.) and (43) are interpreted existentially fails to account for the data.

Thus, (41a.) and (43) cannot be given readings on which their bare plural is interpreted existentially. The proponent of Liebesman's approach has exhausted the basic options, he will have to appeal to something more radical to account for variability.

3.2.3.2 Objection 2: Kind Predication and Truth-Conditional Pragmatics

There is a lively debate in philosophy of language and linguistics as to whether binding data of the sort presented above suffices to establish covert structure in the semantically derived logical form of the sentence in question. Stanley and Szabo (2000) and others represent one side of the debate, however there are a whole slue of theorists that are not convinced that binding data warrants positing covert structure (cf. Recanati (2004, 2010, 2002), Sperber and Wilson (1995), Carston (2000, 2002), Bach (2003), among others).

Advocates of truth-conditional pragmatics will not be convinced by the pro *Gen* data presented above — i.e., that *Gen* is part of the logical form of generics. They will argue that the linguistic data in question only indicates that *Gen* can enter into the interpretation of generics at some stage of interpretation; no more, no less — and that stage might be semantic, but it might also be pragmatic.

Cohen (2012) has recently advocated just such a position. According to Cohen, generics share a semantically derived logical form with kind-predications, and the quantificational form of generics involving *Gen* is only communicated via characteristically pragmatic means.

> The generic quantifier is not "there" in the input sentence: it is introduced by reinterpretation... Generics are a case of predicate transfer... which is clearly a pragmatic process. Hence, the current work can be seen as evidence for views that logical form is not determined by semantics alone. Rather, semantics interacts with some pragmatic process (what Recanati (2004) calls *modulation*), and they establish logical form together. (Cohen, 2012)

On his view, the literal meaning (or the semantically derived content) of a given generic is the content of the corresponding kind-predication — which is only at a later stage enriched into a quantificational interpretation containing *Gen* by pragmatic means.

In particular, Cohen thinks generics are analogous to examples like the following:

(65) I am parked out back.

(66) Noam is inside today's newspaper.

In these examples, it is obviously not the literal meaning which is communicated, but rather some reinterpreted meaning. According to Cohen, following Nunberg (1995), acceptable logical forms for (65) and (66) are pragmatically derived via a process called *predicate transfer*. In (65) for example, "the predicate *parked out back* contributes a property of persons, the property they possess in virtue of the locations of their cars" (1995, p.111), but this cannot be the semantically derived contribution of *parked out back* and so a pragmatic alteration of its contribution to the truth-conditions of what is said is needed — the alteration in these cases is yielded by predicate transfer. Similarly, the predicate is *in today's newspaper* in (66) contributes a property of objects, the property they possess in virtue of there being stories about them in today's newspaper.

Predicate transfer is a process which applies to make sense of the speaker's utterance, it is triggered when the literal meaning of the utterance doesn't make sense. Further, the transferred property is meant to offer "a useful way of classifying its bearer relative to the immediate conversational interests" (1995, p.114). Cohen (2012)

On Cohen's proposal, generics, like (67) and (68), are cases of predicate transfer.

(67) Tigers are striped.

(68) Mosquitoes are irritating.

(67) and (68) are initially treated as kind predications, but since the relevant kind-predications don't make sense, predicate transfer applies to make sense of the speaker's utterance. The distinctively generic interpretation involving *Gen* is derived from the sentence expressing direct kind predication. According to Cohen, the predicate (from the kind-predication), *striped* and *irritating* in our examples (67) and (68), changes meaning so as to apply to instances of the kind rather than the kind itself, and this is what allows and yields the generic interpretation.

Reply to Objection 2: Here I respond to the analogy between generics and cases of predicate transfer brought up by Cohen and the threat of truth-conditional pragmatics for generics construed as a semantically unified class. I bring up worries in addition to those rehearsed by allies of Stanley and Szabo (2000). There are at least three big worries in applying such strategies to generics.

Unavailability of Literal Content: First, it is a worry that the literal semantic content Cohen proposes (i.e., what is expressed by the corresponding kind-predication) seems entirely unavailable to ordinary speakers. To see this, observe that dialogues like those in the introduction, (14) and (15), and like the following seem bad:

(69) A: Tigers are striped.
B: ?? That's absurd; the kind tiger doesn't have stripes.

(70) A: Mosquitoes are irritating.
B: ?? That's absurd; the kind mosquito can't be irritating.

Cohen predicts that (69B) and (70B) should be fine since the literal content of the generic uttered by A in each case should be an interpretation which is available to B. By contrast, consider such dialogues with Nunberg's cases (65) and (66):

(71) A: I am parked out back.
B: That's absurd; you can't be parked, only your car can.

(72) A: Noam is inside today's newspaper.
B: That's crazy; Noam can't be inside the newspaper, only a story about him can.

In Nunberg's cases, the literal interpretation is available to B.[11]

Generics as Category Mistakes? On one plausible description, Nunberg's cases are instances of category mistakes: A person is the wrong kind of thing to be parked or to be in a newspaper — just as in classic cases of category mistakes, rocks are the wrong kind of thing to predicate sleep of.[12]

If this is a plausible description of Nunberg cases, it is problematic, to say the least, to classify generic statements in general as a species of these types of cases. To do so would be to claim that all generic statements involve making a category mistake. That is a very radical view: Generic utterances are ubiquitous in everyday communication and defeasible reasoning, not to mention in parental speech and scientific discourse. This obviously is not a conclusive consideration against treating generics as category mistakes — it might be that we are making this type of mistake all the time, but it is worthwhile to draw attention to the fact Cohen's analogy between cases like (65)-(66) and generics is not such an innocent one.

Perhaps a better analogy for Cohen to rely on is with predications like the following — in each, an ordinary object is ascribed a property that it instantiates in virtue of its parts:[13]

(73) The chair is wooden.

(74) The kettle is red.

(75) The table is touching the wall.

In order for (73) to come out as true, it is sufficient that a majority or almost all of the chair's parts are wooden. The existence of a fabric seat or metal armrests doesn't jeopardise the truth of (73), depending on context of course. In the case of (74), it seems sufficient that some whole visible part is red for the predication to come out as true. For (75), it is very often sufficient that only some small part of the table is touching the wall. Just as the quantity and salience of the parts that are required to make (73)-(75) true are unsystematic and difficult, the quantity and salience of members of a kind that are required to make generics, like

[11] As a further point, note that it is possible to conjure up contexts in which (65) and (66) are interpreted with their literal content: Consider a context where a wizard wraps Chomsky into a box made of today's newspaper. In such a context, we can interpret the sentence quite literally. By contrast, however, thinking up such contexts in the case of generics seems impossible.

[12] Cf. Magidor (2009) for why category mistakes can be meaningful.

[13] Liebesman (2011) argues that such cases are analogous to generics, and hence that generics are predications.

(1)-(3), true is similarly unsystematic and difficult. Since we have no problem, attributing properties to objects based on their parts, we should have no problem attributing properties to kinds based on their members, or so the advocate of kind-predication would claim.

Note, however, that sentences like (73)-(75) can be conceived of in two distinct ways: As above (i.e., where the subject is enriched so that the predicate applies to the parts of the object, for example), or as the property predicated being enriched so that it becomes a property that is ascribed to the object in virtue of it being a property instantiated by the object itself. Readings of (73)-(75) construed as above are given in the (a.) cases below and readings construed as just described in the (b.) cases:

(76) a. The legs of the chair are wooden.
b. The chair is predominantly wooden.

(77) a. The outside of the kettle is red.
b. The kettle is red on the outside.

(78) a. The side of the table is touching the wall.
b. The table is touching the wall with its side.

Supposing that it is the legs of the chair which are wooden in the context of an utterance of (73). If the subject is enriched so that the predicate *is wooden* appropriately applies to some sufficiently relevant part of the object in (76a.), we end up with an interpretation along the lines of (76a.). Whereas, if the predicate *is wooden* is enriched so as to apply to the entire chair, then we end up with an interpretation along the lines of (76b.). Similarly, in cases (77)-(78).

If generics are given similarly enriched logical forms, then we get interpretations like the following given the two different types of enrichment.[14]

(79) a. The instances of the kind tiger have stripes.
b. Tigers predominantly have stripes.

(80) a. The instances of the kind mosquito are irritating.
b. Mosquitoes predominantly irritating.

If this analogy is meant to be taken all the way along, then I think it supports a more radical pragmatic free-enrichment proposal, one in which there is nothing which unifies generics into a distinctive semantic class. Just as there is nothing

[14]It is interesting to note that the (b.) interpretations could be one way of giving a pragmatic version of the *Gn* verb-phrase operator of Carlson (n.d.), discussed earlier in this chapter.

which unifies cases like (73)-(75) into a distinctive semantic class. But this takes us into an entirely different territory to what Cohen intended. Since for the purposes of this paper, we are assuming that generics form a distinctive semantic class, I leave this option to other work. Further, if the analogy is not taken all the way, then we are again left with the problem that the literal predications in (73)-(75) can perfectly well be made sense of, while those of generics cannot.

The Optionality Criterion: The second big related worry for free-enrichment based proposals is that it fails, as far as I can judge, the test put forward by Recanati (2002) for the applicability of such pragmatic processes (what he calls the *optionality criterion*). Recanati's test can be understood from the following quote:

> ... what characterizes genuine [enriched] constituents is the fact that their contextual provision is not mandatory — it is not required in virtue of a linguistic convention governing the use of a particular construction (or class of constructions). In context, it may be that the [enriched] constituent is "required"; but then it is required in virtue of features of the context, not in virtue of linguistic properties of the expression type. A constituent is mandatory in the relevant sense only if in every context such a constituent has to be provided (precisely because the need for completion is not a contextual matter, but a context independent property of the expression type). This, then, is the criterion we must use when testing for (genuine) [enrichment]: Can we imagine a context in which the same words are used normally, and a truth-evaluable statement is made, yet no such constituent is provided? If we can imagine such a context, then the relevant constituent is indeed [enriched]; if we cannot, it is articulated, at some [semantic] level of linguistic analysis. (2002, p.316)

According to the optionality criterion, in order for genericity to be genuinely yielded via the processes of free-enrichment, for any given generic, we must be able to imagine a context in which "the same words are used normally, and a truth-evaluable statement is made, yet no such constituent is provided." For a given generic (choose your favourite one), the test is not easy to implement. All I can do is try to imagine such contexts and report that I fail. I know of no context where these sentences are used "normally", and there is no generic claim made. Any such context will require using the terms in a non-normal way (e.g., metaphorically, ironically or enriched with some other quantifier expression). So I conclude: Generics fail Recanati's test — or there is at least no evidence that

they pass it — since free-enrichment is fairly controversial in the first place, then the fact that generics fail one of its lead proponent's test is a serious problem. A generic interpretation is the normal interpretation of the sentences under consideration. Moreover, a generic interpretation is obligatory for these sentences, in all normal circumstances.

3.3 Defending Gen

3.3.1 Objection 1: Adverbs of Quantification, the Indefinite Singular and Generics

Recall that in the introduction to this chapter I claimed that the Quantificational Analysis is more precisely described as an A-quantificational Analysis. Well, there is a worry with the A-quantificational account and the general strategy I will employ below (in my reply to the linguistic evidence for kind-predication), brought out by considerations of Cohen (2001a). The basic problem is this: It is possible that this strategy, and the A-quantification theorist more generally, falsely predict the distribution facts of generics. If *Gen* is a covert adverb of quantification according to the A-quantification theorist, then it should be possible (modulo some details) to get generic readings whenever a reading of a corresponding sentence with an overt adverb of quantification is available. If there are sentences which do not yield generic readings when the overt adverb of quantification is deleted, then the *Gen*-theorist and the strategy of my reply above appear to be in trouble.

Cohen on Indefinite Singular Generics: Cohen (2001a) argues that there are empirical differences between overt A-quantifiers and *Gen*, and that such differences are sufficient to doubt that *Gen* should be classified as the same semantic type as overt A-quantifiers. In particular, Cohen argues that overt quantifiers like *usually* or *normally*, are always acceptable with the indefinite singular, while generics are often unacceptable. This is a problem for the A-quantification theorist since she is committed to A-quantified sentences and generics being acceptable in the same linguistic environments.

Cohen offers the following as examples — the generics in (81a.) and (82a.) are, according to Cohen (cf. Chapter 5 of this dissertation for further references and examples), unacceptable while the overtly quantified sentences (81b.) and (82b.) are fine:

(81) a. A madrigal is popular. (?)
b. A madrigal is usually popular.

(82) a. A king is generous. (*?*)
b. A king is usually generous.

This data is intended to be evidence that explicit A-quantifiers have different semantic properties from generics — we cannot simply claim that *Gen* is an A-quantifier if it cannot appear in the same sorts of linguistic environments as explicit A-quantifiers.

Reply to Objection 1: Note, however, that it is not that (81a.) and (82a.) are unacceptable *per se*, rather it is that the generic reading is not as easily available since the existential readings of (81a.) and (82a.) are so strongly preferred. In the case where there is an explicit adverb of quantification, there is no issue for the interpreter as to whether the correct interpretation in that context is the generic or existential reading.[15] Thus, we shouldn't expect that the distribution of overt A-quantifiers and *Gen* match by one-to-one correspondence. There are many instances where this won't be the case — for instance, if the generic is obviously false, then the existential reading might be preferred; or if the speaker has previously been talking about kings, for example, in general terms, then the generic reading will be preferred.

Thus, I conclude that small differences in the distribution facts of A-quantifiers and generics can be accounted for by the fact that the former are explicit quantifiers while the latter is not. These small differences in distribution pose no genuine worry to the claim that generics are A-quantified sentences.

3.3.2 Objection 2: Linguistic Arguments for Kind-Predication

A number of influential arguments have been adduced with the conclusion that generics are kind-predications. These arguments use linguistic evidence that kind-predications and generics share a logical form, and that the subject terms of generics directly refer to kinds to argue that generics are a species of kind-predication. Such arguments have been offered most notably by Carlson (n.d.) and more recently by Liebesman (2011); both claim that linguistic evidence of this sort entails that generics are kind-predications of a particular sort. As a corollary of this, they claim that the logical form of generics is not quantificational — i.e., that there is no *Gen* in the logical form of generics.

In what follows, I will briefly present the two main arguments of Carlson and Liebesman — the *argument from uniformity* and the *argument from direct reference*. I

[15] Chapter 5 presents further considerations to this effect.

will then go on to present my reply on behalf of the Quantificational Analysis.

The Argument from Uniformity: The primary argument of Carlson and Liebesman[16] is the argument from uniformity. The argument goes roughly like this, according to Liebesman:

> Kind-predications, such as [(12)] and [(13)], are superficially similar to generics, though they uncontroversially don't have a tripartite quantificational structure. The argument from uniformity proceeds from this observation. The idea is that there is significant pressure to give a uniform semantics to generics and kind-predications. [Since] kind-predications have a simple bipartite structure [and given] considerations from uniformity, we can conclude that generics have the same structure. (2011, p.412)

The "pressure to give a uniform semantics" that Liebesman alludes to in the quote, comes from two related linguistic phenomena: the co-occurrence of kind predication and genericity, and the inferences patterns in which kind-predications and generics jointly participate. To illustrate the former phenomenon,[17] consider the following examples involving complex predication (as in (84)) and non-restrictive relative clauses (as in (85)):

(84) Mosquitoes are widespread and irritating.

[16] It is worth noting that Liebesman weakens Carlson's original position of uniformity in a way that directly undermines the argument from uniformity. Liebesman endorses an ambiguity view of bare nominals (this is evident from commitments he goes in for on p.425) so that he has a quantificational semantics of existential readings. The problem is that the argument from uniformity applies equally well to the existential readings as it does to the generic readings of the relevant sentences: We have all the same co-occurrence, joint inference and direct reference data with existential readings (cf. Krifka et al. (1995, p.116)). Carlson, on the other hand, was motivated by providing a unambiguous semantics for bare nominals, and provided a non-quantificational semantics to the existential and generic readings of the relevant sentences.

[17] To illustrate the latter phenomena, the existence of inference patterns in which kind-predications and generics jointly participate, consider the inference from the kind-predication in (13) and the generic in (83) to (84):

(13) Mosquitoes are widespread.

(83) Mosquitoes are irritating.

(84) Mosquitoes are widespread and irritating.

The inference pattern encompassed by (13), (83) and (84) seems valid: It seems that in general, we can infer from truth of a kind predication and the truth of a generic to the truth of the co-occurrence sentence. The same issue as with co-occurrence arises for the *Gen*-theorist, but perhaps slightly more acutely: The *Gen*-theorist, who is committed to a non-uniform account, needs a story to tell about how to ground the inference without a common semantic element to do so. By contrast, a uniform account needn't provide any such story.

(85) Mosquitoes, who are irritating, are widespread.

The property of being widespread is a property of kinds. As such, the predicate that denotes this property, *widespread*, is said to be kind-selecting. At least prima facie, generics do not contain predicates which are kind-selecting, but rather predicates which select for individuals (members of kinds) — for example, the property of being irritating is a property of individual mosquitoes and the predicate *irritating* combines with the bare-plural *mosquitoes* to yield a generic interpretation, as in the generic:

(83) Mosquitoes are irritating.

What is interesting about (84) and (85) is that they have a co-occurrence of kind-predication and genericity. The problem for the *Gen*-theorist, according to Carlson and Liebesman, is that, on a non-uniform account, the single occurrence of *mosquitoes* must perform two distinct semantic functions: It must contribute a referent for the kind-predication and a domain to restrict the generic generalisation. Since the *Gen*-theorist has to go non-uniform (all theorists agree that widespread is kind-selecting and therefore, generates a kind-predication), then she must give a story about how the single subject, *mosquitoes*, can perform both semantic functions. On the other hand, a uniform account, which the advocate of kind-predication can provide, has a much easier explanation of the phenomena of co-occurrence, since *mosquitoes* only performs the single semantic function. Thus, examples of co-occurrence like (84) and (85), according to Carlson and Liebesman, are evidence that kind-predications and generics share a semantics and hence, a basic logical form (that of kind-predications).

The argument from uniformity and the corresponding phenomena of co-occurrence and joint inferences seems to be a genuine challenge for the *Gen*-theorist. But not one that cannot be overcome: In a moment, we will see some solutions and a general reply to this sort of data. But first, I will put the *Gen*-theorist in an even worse position by considering Carlson and Liebesman's second argument, the argument from direct reference, and the seemingly problematic linguistic evidence brought out by it.

The Argument from Direct Reference: Recall from my introductory remarks that generics come in a variety of syntactic forms. The primary insight of the argument from direct reference exploits this fact by appealing to sentences which appear to express generic contents and nonetheless contain paradigmatically directly referential expressions in subject position. The argument can be motivated by reasons of unity and it goes something like this:

1. Assume for reductio that generics are given a quantificational analysis.

2. Assume further, that we are seeking a unified treatment of the logical form of generics (i.e., the general content expressed by generics comes from a single source).

3. Some generics have paradigmatically directly referential expressions (that refer to kinds) in subject position.

4. Such expressions should not be given a quantificational analysis (they are paradigmatically directly referential after all).

5. Therefore, the *Gen*-theorist must concede that some generics express kind-predications.

6. But this is a contradiction, since we are seeking a unified treatment of the logical form of generics.

As evidence for premise 3., the advocate of kind-predication invites us to consider sentences like the following which have paradigmatically directly referential expressions in subject position:

(86) This kind of animal is a mammal.

(87) This kind of animal barks.

Examples like (86) and (87) are generic since their predicates typically give rise to generic interpretations when combined with bare plurals — as in, for example:

(88) Dogs are mammals.

(1) Dogs bark.

Moreover, (86) and (87) seem to share the same or relevantly similar truth-conditions of their counterparts (88) and (1) — in the sense that (86) and (88), and (87) and (1) seem to tolerate the same exceptions and require that the same prevalence of dogs possess the predicated property. Thus, (86) and (87) are generics. But the complex demonstrative *this kind of F* is thought by most theorists to be paradigmatically directly referential. And if it refers to anything, it refers to a kind — after all it says so. Thus, since a unified analysis of the logical form of generics is desirable, the subject terms of generics must refer directly to

kinds, which leaves no room for quantification, and no room *Gen*.[18,19]

Reply to Objection 2: How is the *Gen*-theorist — i.e., those that endorse the quantificational analysis — to deal with the argument from uniformity and the argument from direct reference? There are several straightforward solutions for the *Gen*-theorist to appeal to in order to deal with the linguistic data in favour of kind-predication, but it would be nice to have a more general reply.[20, 21] That

[18]Liebesman (2011, pp.417-18) points out that the data also works for simple demonstratives — consider:

(89) This barks. [Demonstrating the kind dog.]

[19]Further evidence for co-occurrence and direct reference comes from examples involving anaphora and the *so-called* construction. Consider the following examples, from Liebesman (2011) and Carlson (n.d.) respectively:

(90) Mosquitoes are widespread. They're also irritating.

(91) Cardinals are so-called because they are red.

The first sentence in (90) is a genuine case of kind predication by most accounts: The property of being widespread is something that applies to the kind as a whole and not individual mosquitoes. The second sentence is generic for analogous reasons as (84) and (85) are; and the occurrence of *they* is anaphoric on the kind referring subject. Since anaphoric reference is thought to directly pick out the referent of its antecedent (i.e., the kind-denoting bare plural *mosquitoes*, this is evidence that the subject of the second sentence in (90) is kind denoting as well. But (90) is generic, thus we have evidence that the subjects of generics directly refer to kinds, which again leaves no room for quantification and *Gen*.

In the case of (91), the *so-called* construction signals that the bare plural *cardinals* names a kind. As in the case of (90), *they* is thought to be anaphoric on a kind-referring subject. Thus, the kind-predication theorist maintains, the subjects of generics directly refer to kinds, which based on considerations of uniformity precludes a quantificational treatment of their logical form involving *Gen*.

[20]To consider just one such solution for the co-occurrence data, let us again look at (84) and (85):

(84) Mosquitoes are widespread and irritating.

(85) Mosquitoes, which are irritating, are widespread.

One solution for the *Gen*-theorist is claim that there is not a single occurrence of mosquitoes in the logical forms of (84) and (85) since they are both elliptical — i.e., (84) and (85) are, for example, in fact elliptical for (92):

(92) Mosquitoes are widespread and mosquitoes are irritating.

This seems plausible enough given the intuitive truth-conditions of (84) and (85). But a defence of such solution is not needed given the more general reply I give in this section.

[21]Chierchia (1998) has provided a very elegant and independently motivated account of how expressions, which can act as the subjects of generics, can be given a directly referential semantics, while generics themselves are analysed as quantificational. His account demonstrates that the issue of direct reference to kinds is independent of whether or not generics are quantificational. The account uses the sophisticated machinery of semantic type-shifting (cf. Partee (1991)) to empirically motivate the mechanism Chierchia calls *derived kind-predication*. This mechanism allows a semantics

is what I will provide here. The data presented by Carlson and Liebesman is not evidence for the non-presence of quantificational structure — provided it is construed as A-quantificational structure. The main insight of the reply is that overt A-quantifiers can occur in all the co-occurrence and direct reference examples put forward by Carlson and Liebesman — observe:

(93) a. Mosquitoes are widespread and irritating.
b. Mosquitoes are widespread and generally irritating.

(94) a. Mosquitoes, which are irritating, are widespread.
b. Mosquitoes, which are normally irritating, are widespread.

(95) a. This kind of animal is a mammal.
b. In general, this kind of animal is a mammal.

(96) a. This kind of animal barks.
b. This kind of animal normally barks.

With this data in tow, my argument in reply to the argument from uniformity, the argument from direct reference and this type of linguistic data more generally is as follows:

1. The overt adverbs of quantification as they appear in (93b.)-(96b.) are not semantically redundant.

2. Overt adverbs of quantification uncontrovertially give rise to quantificational structure.

3. The presence of a kind referring expression in (93b.)-(96b.) is compatible with the presence of an overt adverb of quantification.

4. If an overt adverb of quantification can occur non-redundantly with kind referring expressions, then non-overt ones can too.

5. The presence of a kind referring expression in (93a.)-(96a.) is compatible with the presence of a non-overt quantifier in the structure of (93a.)-(96a.). (I.e., a generic reading of (93a.)-(96a.) does not show the absence of a non-overt adverb of quantification in (93a.)-(96a.).)

for bare nominals that directly refers to kinds, and also allows generics to contain a quantifier expression *Gen*. I won't elaborate here, simply note that the argument from direct reference and the corresponding linguistic data needn't pose any dire consequences for the *Gen* theorist provided she accepts Chierchia's theory (and it turns out to be empirically correct).

Since the adverbs of quantification in (93b.)-(96b.) are non-redundant, they have to interact in some way with the corresponding kind referring expression. Call the correct account of that interaction T (since the A-quantifiers are non-redundant, we know there is some such T, even prior to fully spelling it out). We can assume that whatever T is, it will also account for the interaction between the non-overt *Gen* in (93a.)-(96a.) and the occurrence of the corresponding kind referring expressions in (93a.)-(96a.).

Thus, the linguistic evidence for kind-predication is compatible with the presence of A-quantificational structure and hence, with *Gen* in the logical form of generics.[22] But the kind-predication theorist might insist that his account of the data is still simpler — why think we need *Gen*? The positive arguments in favour of the existence of *Gen* — like the binder argument — tell us why we need it.

3.3.3 Objection 3: Unpronounced

Perhaps the strongest reason to doubt the existence of *Gen* is that it is never pronounced. Not only is it unpronounced in English, but according to a handful of theorists (Krifka et al., 1995), it is not pronounced in any natural language. Liebesman (2011) claims that "this flies in the face" of a general principle he calls Pronounced:

Pronounced: If English contains an unpronounced, semantically significant adverb of quantification,[23] then we would expect that some other natural language contains a synonymous adverb of quantification that is, or at least can be, pronounced. (2011, p.415)

Indeed, if this empirical fact about natural languages is true (though I know of no extensive investigation on the issue), then the universal absence of an articulated A-quantifier in generic sentences across languages is striking. It could be a merely accidental feature of natural languages or it might warrant explanation. However, no more explanation than the universal absence of other expressions

[22] Note that the data in the previous footnotes is also accounted for:

(97) Normally, this barks. [Demonstrating the kind dog.]

(98) Mosquitoes are widespread. Generally, they're also irritating.

(99) Cardinals are so-called because they are normally red.

Carlson (n.d., pp.123-4) was aware of data of this sort and argued for a non-quantificational semantics for A-quantification. Such a view is not generally accepted.

[23] In Liebesman's original formulation, "adverb of quantification" is "operator". Since the evidence indicates that *Gen* belongs to the category of adverb of quantification, this is a more suitable statement of Pronounced.

which occur in logical form, or so I will argue.

Reply to Objection 3: First, note that the principle PRONOUNCED is surely more general: It would be blatantly *ad hoc* to limit its applicability to A-quantifiers or some limited group of expressions. In order to qualify as a principle at all it needs to apply to expressions in general. In other words, PRONOUNCED is merely an instance of the more general principle PRONOUNCED*:

PRONOUNCED*: If English contains an unpronounced, semantically significant expression, then we would expect that some other natural language contains a synonymous expression (or phrase) that is, or at least can be, pronounced.

PRONOUNCED* is false for the simple reason that it has counterexamples. There are notable expressions that are universally absent from surface structure, and are no less "semantically significant" for being so. Some examples are as follows: quantifier domain variables, implicit argument places, *PRO* and existential closure. Further, the fact that there are no synonymous expressions (or phrases) substantiating their existence doesn't make them any less real.

Domain variables, implicit argument places, *PRO* and existential closure all have the features which are purported to warrant doubting the existence of *Gen*: In each of the languages in which these expressions occur, they are not unpronounced. Further, there is no more evidence for an overt expression or phrase which is synonymous with *Gen* and is invariant across contexts of utterance, as there is for any of these expressions. Since we don't doubt the existence of such expressions, we have no reason to doubt the existence of *Gen*. PRONOUNCED* is false.

More generally, many things we find in a fully fleshed out logical form are unpronounced, but they are no less syntactically and semantically significant — for example, they have compositional roles and language wouldn't work without them. We don't have any ready explanation for why such bits of logical form are unpronounced — no one really knows, at least at this stage, why they are not.

The proponent of PRONOUNCED* might respond that expressions like domain variables, *PRO* and existential closure are disanalogous to *Gen* because such expressions are not "semantically significant" in the right way. These expressions differ from *Gen* in some important way. This presents a challenge to the proponent of PRONOUNCED* to find a way w of being semantically significant such that *Gen* is semantically significant in way w and domain variables, implicit argument places, *PRO* and existential closure are not. In other words, the proponent of

PRONOUNCED* needs to articulate:

PRONOUNCED$^*_w$: If English contains an unpronounced expression which is semantically significant in way w, then we would expect that some other natural language contains a synonymous expression (or phrase) that is, or at least can be, pronounced.

There are ways the proponent of PRONOUNCED* can go in spelling out w, but these won't be explored here. I will only remark that domain variables, implicit argument places, *PRO* and existential closure are very important to the syntax and semantics of natural language, and no one has a good story about why they are not pronounced.

It is also worth noting that there might be some potential independently motivated explanations of why *Gen* is never pronounced. Leslie (2008, 2007) offers one such explanation; she proposes a theory of generics on which *Gen* expresses our default mode of generalisation which is invoked in the absence of other modes, and as a result our language does not need a pronounced expression *Gen*.[24]

3.4 Conclusion

This chapter is devoted to defending the existence of an unpronounced quantifier expression, *Gen* in which the logical form of generics . The Chapter defends the original argument of Carlson (1989). In addition, this chapter provides a new argument in favour of treating the logical form of generics as quantificational — what I call the *binder argument*. The Chapter also responds to various objections. In particular, it responds to the linguistic evidence offered in favour of a kind-predication analysis, and some recent arguments in favour of treating generics as kind-predications which are enriched by pragmatic mechanisms.

3.5 Appendix: Additional Evidence in Favour of Gen

3.5.1 Scope Ambiguities

Not only do generics on their own exhibit ambiguities as a result of which material gets mapped to the restrictor or scope of the generic, which is an argument for at

[24] As a final note which I will not discuss in detail, the evidence offered in favour of *Gen* is at least as substantial as the evidence offered in favour of the existence of domain variables, implicit argument places, *PRO* and existential closure, and thus, either we are sceptical of the existence of all such expressions, or we allow that such evidence is enough to establish quantificational structure.

least the quantificational structure of generics. They also exhibit a full range of scope ambiguities with negation, quantifiers, operators, *de re/ de dicto*, etc. Such scope ambiguities are perhaps the most direct evidence for the existence of an A-quantifier in generics at the level of logical form. There cannot be a scope ambiguity with only one operator/quantifier present at logical form.

Consider the following examples:

(100) Canadian academics are supported by a single granting agency. (Schubert & Pelletier, 1987)

(101) Dogs must wear a leash. (adapted from Krifka et al. (1995) who cite examples from Halliday, (1970))

Example (100) is purported to exhibit scope ambiguity between the indefinite *a single granting agency* and *Gen*.

The two readings of (100) can be paraphrased as (102a) and (102c), and whose logical forms are construed as (102b) and (102d) respectively:

(102) a. In general, if x is a Canadian academic, then there is a single granting agency y such that x is supported by y.
b. Gen x [Canadian(x)] [$\exists y$ Single-Granting-Agency(y) $\wedge$ Support(y, x)]
c. There is a single granting agency y such that in general, if x is a Canadian academic, then y supports x.
d.$\exists y$ (Single-Granting-Agency(y) $\wedge$ Gen x [Canadian (x)] [Support(y, x)])

Similarly, the two readings of (101) can be paraphrased as (103a.) and (103c.), and whose logical forms are construed as (103b.) and (103d.) respectively:

(103) a. In general, if x is a dog, then it must wear a leash.
b. Gen x [Dog(x)] [Must(Wear-a-Leash(x))]
c. It must be the case that in general, dogs wear a leash.
d. Must (Gen x [Dog(x)] [Wear-a-Leash(x)])

3.5.2 Embedding

Pelletier and Asher (1997), in a different context, notice that generics can exhibit embedding behaviour. It is hard to see how generics could exhibit such behaviour if they were not a A-quantifier in their logical form.

Consider the following examples adapted from Pelletier and Asher (1997):

(104) People who work late wake up late.

(105) A cat is healthy if it chases an object when it is moved in front of its eyes.

The embedded reading of (104) can be paraphrased as (106a) and its logical form as something like (106b):

(106) a. In general, if x is a person who works late, then in general, x wakes up late.
b. Gen x [Person(x) $\wedge$ Gen y [$x = y$] [Work-Late(y)]] [Wake-Up-Late(x)]

Similarly, the embedded reading of (105) can be paraphrased as (107a) and its logical form as something like (105b):

(107) a. In general, if x is a cat and in general, x chases objects when they are moved in front of x's eyes, then x is healthy.
b. Gen x [Cat(x) $\wedge$ Gen y, s [Object(y) $\wedge$ Move-in-Front-of-Eyes(y, x, s)] [Chase(x, y, s)]] [Healthy(x)]

4 Generics and Context

4.1 Introduction and Landscape

When speakers intend to communicate *something general* about the world, they often use sentences like the following:

(1) Candy is bad for your teeth.

(2) A raven is black.

(3) The tiger has stripes.

(4) Dogs bark.

(5) This kind of animal barks. [Uttered while demonstrating a dog.]

These are examples of generic sentences (or simply *generics*). Generics come in a variety of forms, as evidenced by (1)-(5), all of which have non-generic uses as well. Moreover, notice that generics seem to communicate this something general without any overt or pronounced element which is responsible for communicating this general content. What characterises the class of generic sentences, then, is that each member of the class communicates some sort of general content without any obvious overt element which is responsible for communicating this content. For example, (1) is a generic since it communicates something general about candy, let's say a generalisation akin to *in general, candy is bad for your teeth*, but does not contain an explicit expression like *in general* which is responsible for communicating the general content of (1). Similarly for (2)-(5).

I have specified this characterisation of generics since it leaves open at least four important, yet mostly neglected, dimensions along which to investigate this class of sentences. The first concerns whether or not generics have a uniform logical form. For example, since (1) seems to be close in meaning to the explicitly quantified sentence *in general, candy is bad for your teeth*, we might suppose that the logical form of (1) is quantificational, and hence, that the logical form of generics more generally is quantificational. It is universally assumed that generics have a uniform logical form — i.e., that they are all quantified sentences or all (kind-)predications — but one could have a view on which some generics (or, alternatively, generic utterances) are quantified sentences, some are kind-predications and some are some other plausible logical form — call this position *scepticism about uniform syntactic structure*.

The second dimension concerns whether or not the general content communicated by a given generic is semantically expressed or pragmatically conveyed, and further, whether or not the mechanism by which they are communicated is uniformly semantic or uniformly pragmatic. To illustrate, consider (1) again: There is a choice point about whether or not we treat the *in general* content as semantically or pragmatically conveyed, just as there is a choice point as to whether we treat quantifier domain variables or the standards or comparison classes of gradable adjectives as semantically or pragmatically conveyed. With the exception of some of Cohen's recent work (cf. Cohen (2012)), I know of no quantificational theory that doesn't treat the general content communicated by generics as semantically expressed and uniformly so. But one could envisage a theory on which for some generics (or, alternatively, generic utterances) the general content is semantically expressed and for others it is pragmatically communicated — call such a position *scepticism about a distinctive mechanism*.

The third dimension concerns whether or not the general content communicated is in any way uniform (or stable) across different generics (or alternatively generic utterances). For example, one can contend that the general content communicated by each of (1)-(5) is uniformly that expressed by *in general*; alternatively, one could contend that the general content expressed by each is different (i.e., it is the content expressed by *in general* for (1) and it is the content expressed by *all normal* for (2), etc.). It is widely assumed that the content communicated is uniform, especially amongst those that endorse a quantificational analysis of generics (e.g., Nickel (2008), Pelletier and Asher (1997); Asher and Pelletier (2012) and Asher and Morreau (1995)). But, again, it is not hard to imagine a view on which the communicated content of each of (1)-(5) is different — call this position *scepticism about a distinctive meaning*.

Finally, the fourth dimension concerns whether or not any part of the general content communicated by a given generic varies across different contexts of utterance. For example, does (1) communicate the general content expressed by *in general* in one context of utterance and communicate the entirely different content of *all normal* in another? Or, for example, if (1) expresses the same *all normal* in all contexts, is there some component of this content that is context-sensitive — e.g., *normal*? Some authors endorse the latter form of context-sensitivity in some form or another (e.g., Cohen (2008), Nickel (2008), Pelletier and Asher (1997); Asher and Pelletier (2012)), but no has yet to endorse full out indexicality of the former type. Other authors do not endorse either form of context-sensitivity (e.g., Leslie (2007), Liebesman (2011)). Call the position which asserts the claim that a given generic can communicate different general contents in different contexts of utterance, *scepticism about contextual stability*.

In Chapter 3, I argue, contra the sceptic about uniform syntactic structure, that generics have a uniform logical form and that this form is quantificational — i.e., involves a covert, unpronounced quantifier expression *Gen*.[1] In particular, I argue that *Gen* is a quantificational adverb (i.e., an *A-quantifier*) like *always* or *in general*,[2] as opposed to a quantificational determiner (i.e., a *D-quantifier*) like *every* or *some*. This allows me to say that generics communicate A-quantificational contents — in other words, that the general content communicated by generics is A-quantificational. I take it that all A-quantified contents are general contents in virtue of being contents that are genuine *generalisations*, and hence, I will characterise the general content communicated by generics as a generalisation. Thus, from this stand-point, I can claim that generics communicate genuine generalisations.

In Chapter 3, I, in addition, argue that the content communicated by *Gen* is semantically expressed and, contra the sceptic about a distinctive mechanism, that this content is uniformly expressed via a semantic mechanism.[3] This allows me to say that generics semantically express generalisations. From this stand point, we can begin to investigate what type(s) of generalisation(s) generics express, and whether or not linguistic context-sensitivity has any role to play.

This chapter has three main goals: The first will be to argue for a form of scepticism about a distinctive meaning for generics. In other words, I will argue that different generics express different generalisations — there is no distinctively "generic" generalisation. The second goal will be to argue for a form of scepticism about contextual stability. In other words, I will argue that a given

[1] Cf. Chapter 3.

[2] I take this class to include quantificational operators.

[3] Cf. Chapter 3.

generic sentence expresses different generalisations in different contexts of utterance. Finally, given that generics are context-sensitive, I will argue that the context-sensitivity involved is indexicality. That is, I will defend the claim that generics are context-sensitive in the non-trivial sense that they can express wholly different generalisations in different contexts of utterance — i.e., their context-sensitivity is not merely due to contextual domain restriction, the presence of some context-sensitive expression or some other context-sensitivity exhibited by A-quantified sentences more generally. I offer an account in which *Gen* is treated as an indexical and generics express different generalisations in virtue of this.

This chapter is structured as follows: In Section 4.2, I motivate the form of scepticism about a distinctive meaning for generics that I wish to defend, and in Section 4.3, I argue that generics are context-sensitive. Indexicality is argued for in Section 4.4. Section 4.5 provides a semantics for generics which treats *Gen* as an indexical A-quantifier. The indexical approach has numerous virtues which I explain in Section 4.6. In Section 4.7, I reply to objections, and in Section 4.8, I briefly conclude.

4.2 Scepticism about a Distinctive Meaning

Here is the thesis of scepticism about a distinctive meaning for generics which I endorse:

Scepticism about a Distinctive Meaning (Sdm): In general, different generics express different generalisations.[4,5]

My main argument for Sdm will be via induction. I will list a sample of generics, and observe that each of them has distinct intuitive truth-conditions — i.e., that each of them, intuitively, express different generalisations . Since intuitive truth-conditions are our best guide to semantic content, we have an argument for Sdm. To this end, consider the following sample of generics and ask yourself what their most salient readings are:

(6) Philosophers give the Locke lectures.

(7) Tigers are fierce.

[4]There are many other ways to put this: One could for instance, change the quantifier or partition the set of generics in some informative way.

[5]This formulation presupposes that generally, generics do express generalisations. I have no problem with endorsing this presupposition for reasons I explained above.

(8) Dogs have four legs.

(9) Supreme Court Justices have even social security numbers.

(10) Books are paperbacks.

(11) Mosquitoes carry the West Nile virus.

(12) Germans wear Lederhosen.

If (6)-(12) express generalisations, then they express very different generalisations. To get a sense of this, consider the paraphrases in (13)-(19) which reflect common judgments about what (6)-(12) communicate on their most salient readings:

(13) Traditionally, only philosophers give the Locke lectures.

(14) In general, tigers are disposed to be fierce.

(15) All dogs without defects have four legs.

(16) All Supreme Court Justices have even social security numbers because they are Supreme Court Justices.

(17) All normal books are paperback books.

(18) Mosquitoes sometimes carry the West Nile virus.

(19) It is a distinctive of a traditional German that he wear Lederhosen.

(6) is intuitively true because traditionally the Locke Lectures are given only by philosophers; this is captured by the presence of *traditionally* in the paraphrase (13). The truth of (14) seems to require that a majority of tigers are fierce and that it is something about the essence of individual tigers that makes them so; the truth of (14) does allow for the existence of some gentle tigers, however. (8) is intuitively true because all dogs except those with birth defects or who've had an accident have four legs. (9) wouldn't be true even if it happened that all Supreme Court Justices had even social security numbers which seems to indicate that the generalisation communicated by (9) introduces some sort of intensional component akin to that expressed by the paraphrase in (16). (10) is intuitively false despite the fact that, generally, books are paperbacks; one way to reflect this is by giving conditions which expresses an even stronger generalisation than *in general* — akin to (17) which is intuitively false. (11) is intuitively true despite the

fact that only a tiny minority of mosquitoes actually carry the West Nile virus at any given time; one way to capture this is to give a paraphrase in terms of a quite weak generalisation like (18). Finally, many generics are intuitively true despite only a minority of the kind in question satisfying the predicated property because the property in question is some sort of traditional and/or distinctive property of the kind, like (12). In these cases the generalisation communicated seems to include such features as part of its content — this is reflected in the paraphrase given in (19).

Many of these cases play a prominent role in the literature on generics (in constructing theories, etc.). Thus, I take them as a representative sample. Further, it is easy to tack on generic after generic on this list and continue to yield different generalisations as intuitive paraphrases. Thus, I conclude, by induction, that SDM.

The astute reader might have noticed that it is hard to pin-down a paraphrase that seems to accurately, uniquely and stably express the generic in question. This is not an objection, but merely fuel for the fire for those that are concerned about the project of specifying a distinctive meaning for *Gen*. A primary source of the concern is that it is hard to see what sort of meaning to assign *Gen* which is invariant across all these different generics.

4.2.1 Variants of SDM

It is worth noting that many authors come close to endorsing SDM — in the sense that they would be untroubled by endorsing, at least in some substantive sense, my diagnosis of the data.

Liebesman (2011), for instance, accepts that generics have varied intuitive truth-conditions and uses this to argue that *Gen* is not part of the logical form of generics; instead, generics are kind-predications, and therefore, do not express generalisations at all:

> The relationship between kinds and their members is... unsystematic. The distinct truth-conditions of [(4)], [(7)], and [(11)] attests to this, as do myriad proposals for the semantics of generics. Each of these proposals tries to capture the force of generics as generalisations. On the theory of simple generics, the mistake is revealed. Generics ascribe properties to kinds and, given the multiplicity of properties and the multiplicity of ways kinds inherit properties from their members, no fully general account of inheritance will be forthcoming. This is why *Gen* has proven intractable. (2011, p.420)

Another example is Leslie (2007, 2008) who accepts that generics express generalisations by means of *Gen*, but uses the seeming intractability of finding a suitable content for *Gen* to argue for a disquotational semantics — the seemingly complex truth-conditions of generics are by Leslie's lights a specification of the correctness conditions for judgments issued by the mind's "primitive cognitive mechanism of generalisation", these types of specifications are not suitable as the semantic content of *Gen* and hence, should be pushed into the metalanguage:

> I am inclined to view the above [complex] truth specifications as worldly descriptions rather than semantically derived truth-conditions. I propose instead that the best semantic truth-conditions we can supply for generics will expose their logical forms by means of a tripartite structure, but will simply use the generic operator *Gen*... Just as we are best off using *is red* in the truth-conditions of *Bob is red*, we are best off using *Gen* in the semantic truth-conditions of generics. When it comes to providing semantic truth- conditions for generics, we may well be best off disquoting the generic operator *Gen*. (2008, p.44)

Still further, other authors make observations which are germane to SDM without actually endorsing it — they hold on to the idea that we can somehow specify a distinctively generic content; it is just a matter of isolating all the noise and interaction effects (e.g., pragmatic effects, the interaction of *Gen* with gradable adjectives, etc.) which are blocking our ability to do so. Asher and Pelletier (2012), recognise many difficulties in interpreting generics uniformly, but think that their modal conditional approach has the resources to handle variability by appealing to the interaction of these two components.[6] Nickel (2008, 2012a) also recognises that there is variability in the interpretation of generics, but claims that it is still possible to isolate a distinctively generic content: Once the interaction effects which complicate matters are isolated, then *Gen* can be given a distinctively generic content — on his proposal, very roughly *all normal*. For Nickel, the complexities are pushed into isolating linguistic interaction effects and specifying the complexities of what *normality* amounts to.

Thus, I take it that the fact that many contemporary authors are converging on something germane to SDM as evidence that the scepticism I am advocating here is not as radical as it seems. There is something in the air and that something is SDM, even if no one has yet had the courage to endorse it.

[6] The semantics of *Gen* which Asher and Pelletier (2012) provide quantifies over all elements of a constant domain, but the consequent of the conditional is evaluated relative to each individual x only in those worlds where x is assumed to be normal.

4.3 Generics and Context

So far I have briefly outlined some data which supports SDM and some considerations in favour of it. I now turn to the stronger thesis I would like to defend:

SCEPTICISM ABOUT CONTEXTUAL STABILITY: In general, for a given generic, the generalisation expressed by that generic varies across different contexts of utterance.

It is worth noting that SDM does not imply SCS — i.e., SCS is a stronger claim: It could be that the truth-conditions for each of (1)-(12) are different, but that all utterances of each of (1)-(12) have the same content in every context (i.e., that for no generic does its truth-conditions vary between contexts of utterance). I believe SDM is true and SCS is true, and the fact that different generics express different generalisations is explained in part by the fact that generics are context-sensitive. That is, the fact that different generics express different generalisations is in part an instance of the fact that generics are context-sensitive. Of course, there will be other contributing factors (e.g., any part of their interpretation which is driven by the content of the restrictor or scope). But at the very least, we shouldn't be surprised by the fact that SDM is the case if SCS is the case. Since SCS is the stronger claim, I will support it by providing five cases which I take to provide evidence for a distinctive form of context-sensitivity exhibited by the generics.

My claim is that *Gen* is a context-sensitive expression and in particular, an indexical expression since, I claim, it is the general content of the generalisation itself which varies — this is evidenced by the very different paraphrases given for (1)-(12): For example in (1) the general content was akin to *in general* and in (10) the general content was akin to *all normal*; two entirely different relations between the restrictor and the scope. In Section 4.4, I will provide further considerations in favour of this, but for now it is ambitious enough to establish a distinctive form of context-sensitivity for generics. For each of the five cases below, I take a single generic and provide two contexts in which the intuitive truth-values of the utterances of this generic sentence differ; I then argue that the difference in truth-value is not a result of some alternative source of context-sensitivity which might arise from the fact that generics are A-quantified sentences or the presence of some other context-sensitive expression which is part of the sentence in question. In this sense, there is a distinctive form of context-sensitivity exhibited by generics. I take it that there are three main alternative sources which could be responsible, in whole or in part, for the context-sensitivity exhibited by the cases

below, it is useful to outline these before we begin:

Alternative Source 1: The presence of a context-sensitive predicate (e.g., a gradable or a vague predicate) in the restrictor or scope of the generic might explain, in whole or in part, the context-sensitivity at issue.

To see an example, consider the following explicitly A-quantified sentence in the two contexts given below:

(20) a. Typically, purple turtles are fast.
b. **Context 1:** Suppose there are two kinds of turtles — purple turtles and yellow turtles. The purple ones are much faster walkers than the yellow ones. Suppose further, that the context is such that the speakers are talking about turtle walking.
c. **Context 2:** Again, suppose there are two kinds of turtles — purple and yellow. The purple ones are much faster walkers than the yellow ones. But this time the context is such that the speakers are talking about fast animals — antelopes, leopards, etc.

In the first context, (20a.) is intuitively true while in the second context (20a.) is intuitively false. The context-sensitivity here seems to be directly as a result of the predicate *fast* picking out different standards of fastness which are salient in the two contexts.

Alternative Source 2: The presence of a quantifier domain variable which forms part of the restrictor of the generic might explain, in whole or in part, the context-sensitivity at issue.

There are two standard ways in which this domain variable can be resolved for A-quantified sentences (cf. Beaver and Clark (2008) on *free association with focus*): Their domains are sometimes contextually restricted to be a domain which is *salient* in the context of utterance; they are also sometimes contextually restricted to a domain that is derived from the relevant alternatives of the scope. To illustrate I'll provide examples of the former and the latter in turn:

(21) a. Typically, cabs are yellow.
b. **Context:** Uttered while talking about cabs in New York City.

It should be clear how the content of an utterance of (21a.) can differ as a result of the salient domain outlined in (21b.). To see an example of the latter, consider:

(22) Bartenders always serve alcohol.

(22) does not quantify over any situation or event whatsoever, otherwise it would never be true; rather, there is an implicit restriction to situations or events in which something is served (or in which drinks are served).

Alternative Source 3: The general feature of so-called question-sensitivity exhibited by A-quantifiers might explain, in whole or in part, the context-sensitivity at issue.

This type of context-sensitivity is evidenced by examples like the following adapted from Schaffer and Szabo (n.d.):

(23) a. Dogs always bite the mailman.
b. **Context 1:** Each morning, some dog, as opposed to cats, mice, etc. bite the mailman. Ann is wondering what kind of animal bites the mailman when he arrives each morning, and Ann learns that there have been mailman bites and milkman bites, but that some dog has been the one doing the biting every time. So Ann says (23a.).
c. **Context 2:** . Each morning, some dog, as opposed to cats, mice, etc. bite the mailman. Ann is wondering what each dog has bitten, and Ann learns that there have been mailman bites and milkman bites, but that a dog has been the biter every time. So Ann says (23a.).

In Context 1, (23a.) seems true, but in Context 2, (23a.) seems false. The difference in truth-value seems to be a result of the different inquiries taking place in the two contexts.

In each of the cases I provide below, I will eliminate these alternative possibilities (where relevant). A useful heuristic for doing so, is to check whether explicitly A-quantified sentences vary their truth-value across the given contexts. If there is no difference in truth-value, then the alternative sources are not responsible for the contextual variability, and hence, this is evidence that the context-sensitivity at issue is distinctive. I will make use of this heuristic throughout the cases.

I now turn to the cases. Again, later I will provide further arguments that *Gen* not only displays a distinctive form of context-sensitivity, but I will argue that *Gen* is an indexical A-quantifier and it is in virtue of this that generics display the distinctive form of context-sensitivity exhibited; further, this is what sets *Gen* apart from regular A-quantifiers. For now, the cases of distinctive context-sensitivity:

Case 1: The only case (as far as I am aware) arguing for the context-sensitivity of generics (not arising from the alternative sources listed above) is from Nickel (2008):

> Consider [(24)].
>
> (24) Dobermans have floppy ears.
>
> The important fact about dobermans is that they are born with floppy ears that breeders then cut to given them the pointy shape we are familiar with. In the context of evolutionary biology, [(24)] is true. The text [(25)] certainly sounds acceptable.
>
> (25) Some breeds of dogs have evolved to focus on their hearing. These breeds have pointy ears. Dobermans, however, mostly rely on their sense of smell, which is why Dobermans have floppy ears.
>
> However, in the context of a discussion of dog breeding, [(24)] seems clearly false, as the text [(26)] illustrates.
>
> (26) While Labradors and golden retrievers have floppy ears, dobermans don't. Dobermans have pointy ears.
>
> (2008, p.644)

Nickel's argument is familiar from other places in which the context-sensitivity of a given expression is argued for. The same sentence is uttered in different contexts and it is claimed that the truth-conditions differ. Nickel describes the two contexts as "the context of evolutionary biology" and "the context of a discussion of dog breeding".

Let's try the heuristic outlined above with several different explicit A-quantifiers, using the same two contexts:

(27) a. Typically dobermans have floppy ears.
b. Generally dobermans have floppy ears.
c. Normally dobermans have floppy ears.

(28) Some breeds of dogs have evolved to focus on their hearing. These breeds have pointy ears. Dobermans, however, mostly rely on their sense of smell, which is why typically / generally / normally dobermans have floppy ears.

(29) While Labradors and golden retrievers have floppy ears, typically / generally / normally dobermans don't have floppy ears. Typically / Generally / Normally dobermans have pointy ears.

There is no difference in intuitive truth-value for each of (27a.)-(27c.) in the two contexts given by (28) and (29): In both contexts, I hear (27a.)-(27c.) as false.[7] This is evidence that the difference in truth-value for (24) between the two contexts is not a result of the gradable predicate *floppy*, the relevance or salience of the content of the preceedeing text in (28) and (29) in settling any contextual domain restriction or the fact that *Gen* is an A-quantifier. If any of these were causing the contextual variability at issue, then they would presumably be causing contextual variability for (27a.)-(27c.), but they do not, thus we can conclude that there is a distinctive type of contextual variability which is exhibited by (24).

Case 2: This is a variation on a case that has been discussed quite a bit in the literature (e.g., Cohen (2011, 2001b)). Suppose that as a matter of fact, all Supreme Court judges happen to have even social security numbers. This, it seems, does not suffice to make the generic, (9) true:

(9) Supreme Court judges have even social security numbers.

Various diagnoses have been made of this, but I want to use the case in support of Scs. In effect, my claim will be that in certain contexts, when certain questions are under discussion or the speakers have certain practical interests, we would, under such circumstances judge (9) to be true. I'm going to assume it is easy to come up with contexts in which it is false, because the standard assumption in the literature is that it is false (even in contexts where all Supreme Court judges have even social security numbers). I will provide two different contexts in which it is intuitively true. To this end, consider the following two contexts:

Suppose the context is such that two friends are planning a party where they want all the party guests' social security numbers to add to an even number. In hopes of providing information that will help compile the list of guests, one of the friends says (9). In a context where all supreme court judges have even social security numbers, I hear (9) as true.

[7] One of my informants was unsure about his judgments in this case. If my reader is unsure about their judgments, recall that the heuristic is merely a diagnostic and this lacuna is likely a peculiarity of the case: There might be lexical reasons for a difference in truth-value for any of the given A-quantifiers. Moreover, my case for a distinctive form of context-sensitivity does not rely solely on this case.

Suppose the context is such that a journalist, Jim, is searching for evidence that Supreme Court judges are involved in some US government conspiracy. Jim observes that all Supreme Court Judges have even social security numbers. He, then says (9) to his colleague as evidence of their involvement in the conspiracy. In such a context, I hear (9) as true.

To get evidence that this context-sensitivity is not in virtue of the fact that *Gen* is an A-quantifier, let's consider the heuristic applied to these cases. If the truth-value of the explicitly quantified utterances is preserved, then we have support for the the claim that the context-sensitivity in the case of (9) is distinctive. To this end, consider utterances of *typically / normally / generally, Supreme Court Judges have even social security numbers*, in the contexts specified above (i.e., the null context, and the two contexts outlined). Since there is no difference in the intuitive truth-value between contexts (I hear them as true across all three contexts), we indeed have evidence of a distinctive form of context-sensitivity.

Case 3: Consider (30):

(30) Cabs are yellow.

Uttered about New York City cabs, (30) is intuitively true, however there are contexts in which it is intuitively false — take the following:

> **Context:** In New York City more or less all cabs are yellow as a result of a city council regulation. Consider a discussion in a New York City Council meeting where the participants are deliberating over changes to the regulations governing the appearance and colour of cabs. Suppose there is an emerging consensus on changing the regulations to allow pink cabs as well.

In such a context, uttering (30) is plausibly false. To bolster this: Note that we are imagining that this is a world where the yellow or pink regulation wins out and becomes law, but no cab driver ever chooses to drive a pink cab, so as a matter of fact all present and future cabs remain yellow.

Our heuristic applied to this case again yields a supporting result: Consider *typically / normally / generally, cabs are yellow*, on its most natural reading it remains true even after the emerging consensus is a change of regulations.

Case 4: Consider (31):

(31) The French eat horsemeat.

It is often noted that when the distinctive properties of the French population are salient, (31) sounds intuitively true. In such a context, (31) seems to express a generalisation along the lines of:

(32) It is a distinctive of several traditional French people that they eat horsemeat.

Whereas, in a context where a group of nutritionists is querying the unhealthy eating patterns of the French population, (31) seems intuitively false:

(33) The French eat croissants and baguettes. They don't eat traditional food, like horsemeat and grains.

In such a context, the negation of (31) seems to express something along the lines of (34):

(34) Generally, the French don't eat horsemeat.

Again, the heuristic provides the result we want: Consider *typically / normally / generally, the French eat horsemeat*, this is arguably false in both contexts.

Case 5: Consider (35):

(35) Mammals lay eggs.

In a context where a biologist, Suzy, is discussing birds, and their relationship to other species, she utters the following:

(36) Birds lay eggs. Mammals lay eggs too.

In this context, (35) is intuitively true. However, in the context of a mother teaching her child the properties of mammals (35) is intuitively false.

Moreover, the explicitly A-quantified *typically / normally / generally, mammals lay eggs* is false in both contexts.

4.3.1 More Arguments for Context-Sensitivity

In *Relativism and Monadic Truth*, Cappelen and Hawthorne present a series of tests which if passed, are meant to provide evidence of context-sensitivity. One of these is the following:

THE AGREEMENT TEST: Let u be a sincere utterance of a sentence S by a speaker A in a context C, and u' be a sincere utterance of *not-S* by a speaker B in a context C'. If from a third context C'', A and B cannot be correctly reported by *A and B disagree whether S*, then S is semantically context-sensitive. Meanwhile, if from a third context C'', A and B can be correctly reported by *A and B disagree whether S*, then this is evidence that S is semantically invariant across C, C' and C''. (2009, p.54-55)

Now if we apply the Agreement Test to a sample of the examples just discussed, we have additional evidence for SCS. Consider again two utterances of (24) in Case 1:

(25) Some breeds of dogs have evolved to focus on their hearing. These breeds have pointy ears. Dobermans, however, mostly rely on their sense of smell, which is why Dobermans have floppy ears.

(26) While Labradors and golden retrievers have floppy ears, dobermans don't. Dobermans have pointy ears.

Knowing about the difference in context, we would be hard pressed to classify the two agents as disagreeing over whether (24).

Similarly, if we apply the Agreement Test to Case 4. Consider (31) embedded in the following text where the speaker is pondering and listing the weird cultural habits of different nations:

(37) People eat lots of weird things. Koreans eat dog meat. Frenchmen eat horsemeat. The Scottish eat haggis.

Meanwhile, a context where the speakers are considering the typical diets of different nations and utters the negation of (31):

(38) The French eat baguettes and croissants. They don't eat horsemeat.

Again, applying the Agreement Test, we have evidence that (31) is context-sensitive since we don't get the sense that the speakers in the two contexts are disagreeing over a particular generalisation about the eating habits of the French.

The Agreement Test is also perspicuously applied to Case 5. Consider a context where a biologist is discussing different species that lay eggs:

(36) Birds lay eggs. Mammals lay eggs too.

Now consider a mother teaching her child about the different characteristics of different species — she utters:

(39) Birds lay eggs, but mammals don't. Mammals give birth to live young.

The judgement here is easy — we do not take the biologist and the mother to be disagreeing over whether or not mammals lay eggs.

Again, the explicit A-quantifier heuristic is useful here since it provide confirmation that the context-sensitivity at issue is distinctive. If we substitute explicit A-quantifiers into the appropriate spots in the examples above, we do get the sense the parties disagree.[8] The Agreement Test, thus, gives further evidence for a distinctive type of context-sensitivity exhibited by generics.

In the next section, I go on to argue that *Gen* is not only context-sensitive, *Gen* is an indexical.

4.4 Generics, Indexicality and Covert Structure

4.4.1 Inductive Argument for Indexicality

As with SDM above, we can provide at least a preliminary argument for the claim that *Gen* is an indexical via induction. In this case, I take a single generic, and observe that in different contexts the general content expressed by the generic is different — and different in the strong sense that the general content expressed is a wholly different relation between the restrictor and scope. To this end, consider (40a.) and (41a.) and what their intuitive truth-conditions are in the given contexts:

(40) a. Dobermans have floppy ears.
b. **Context 1:** The context of evolutionary biology from Case 1.
c. **Context 2:** The context of dog breeding from Case 1.

[8] One of my informants was sketchy about this claim as applied to (25), so I will spell this example out explicitly so my reader can decide for herself:

(25) Some breeds of dogs have evolved to focus on their hearing. These breeds have pointy ears. Dobermans, however, mostly rely on their sense of smell, which is why Dobermans normally have floppy ears.

(26) While Labradors and golden retrievers have floppy ears, dobermans normally don't have floppy ears. Dobermans have pointy ears.

If my reader does get this sense, recall that the Agreement Test is merely a diagnostic used as additional evidence to back up the more substantive arguments above, and this uneasiness is likely a peculiarity of the case.

(41) a. The French eat horse meat.
b. **Context 1:** The context where salient properties of the French population under discussion from Case 4.
c. **Context 2:** The context of an inquiry into the eating habit of the French population from Case 4.

(42) a. Mammals lay eggs.
b. **Context 1:** The context where salient properties of the birds are under discussion from Case 5.
c. **Context 2:** The context of a mother teaching her child about mammals from Case 5.

If the utterances of (40a.) and (41a.) express generalisations, then they express very different generalisations in the given contexts. To get a sense of this, consider the paraphrases in (42b.) and (42c.), and (43b.) and (43c.) which reflect common judgments about what (40a.) and (41a.) communicate on their most salient readings in the respective contexts:

(43) a. Dobermans have floppy ears.
b. All normal dobermans in biologically uniform worlds have floppy ears.
c. In general, dobermans have floppy ears.

(44) a. The French eat horse meat.
b. Many stereotypical French people have the distinctive property of eating horsemeat.
c. In general, the French eat horse meat.

(45) a. Mammals lay eggs.
b. Some mammals lay eggs.
c. In general, mammals who reproduce in some way lay eggs.

Note that the general content expressed in each context is an entirely different generalisation. The general content of (43a.) in Context 1 is something akin to the quantifier-phrase *all normal situations in biologically uniform worlds*, while in Context 2 it is something akin to *in general*. Similarly, the general content of (44a.) and (43a.) in the respective contexts are intuitively very different as evidence by the distinct paraphrases. The fact that the intuitive general content is so distinct between utterances of a single generic sentence, is evidence that the expression responsible for expressing this content, *Gen*, is not merely context-sensitive, but an indexical.

What we've observed is that these generics can express very different generalisations in different contexts. But given that *Gen* is an A-quantifier, what does it mean for an expression of this sort to express different generalisations in different contexts? With a better characterisation of what is meant by this, we can construct a more precise semantic theory and more precise considerations in favour of the claim of indexicality, that takes us beyond merely relying on an inductive argument based on intuitive truth-conditions, like the one provided above.

In the case of explicit A-quantifiers, in addition to the context-sensitive components of their meanings, there is a stable semantic contribution made — a stable generalisation or relation between the restrictor and the scope of the quantified sentence which is expressed in every context of utterance. In the case of A-quantifiers, that semantic contribution is represented in logical form by two lexical components:

(i) That which contributes the quantificational force of the generalisation, and;

(ii) that which contributes any (stable) lexical restriction on the domain of quantification.

It is useful to consider a few examples to see what is meant by the stable semantic contributions of (i) and (ii). *Sometimes* and *always* are the easiest to illustrate:

(46) a. Sometimes dobermans have floppy ears.
b. Dobermans always have floppy ears.
c. **Context 1:** The context of evolutionary biology from Case 1.
d. **Context 2:** The context of dog breeding from Case 1.

The explicit A-quantifier *sometimes* contributes a quantificational force which is existential and it contributes this same existential force across all contexts of utterance. As a concrete example consider (46a.) in (46d.) and (46e.): In both contexts, (46a.) says something like some dobermans have floppy ears. Further, *sometimes* it contributes a lexical restriction to actual situations and this lexical restriction is stable across all contexts of utterance. We can observe that in a similar fashion *always* has a stable quantificational component and a stable lexical restriction over the domain of quantification: The quantificational force is universal and the lexical restriction is again over actual situations. Consider, for example, (46b.) in (46d.).

If *Gen* is an indexical A-quantifier, then at least one of these semantic components is not stable, but rather varies between contexts. Thus, there are three types of indexical accounts of *Gen*:

(Indexicality 1) The quantificational force of *Gen* varies with context and the lexical restrictor remains fixed.

(Indexicality 2) The lexical restrictor of *Gen* varies with context and the quantificational force remains fixed.

(Indexicality 3) Both the lexical restrictor and the quantificational force of *Gen* vary with context.

Though the strongest position and hence, most difficult to defend, I prefer views on which (Indexicality 3) is the case, not because I am particularly enthusiastic about variability, but rather because I think not only does the data confirm variability of both components, but the virtues of understanding *Gen* as a pure indexical of this sort are amplified by this form of pure variability. Views which endorse (Indexicality 3) furnish *Gen* with greater expressive power which makes *Gen* a particularly special and useful linguistic device. Section 4.6 deals with what I see as the virtues of the indexical approach in further detail.

In the remainder of section 4.4, I will outline in more detail what (Indexicality 1) to (Indexicality 3) amount to and then, provide some further considerations in favour of treating *Gen* on the model of indexicals.[9]

4.4.2 The Variable Quantificational Force of Generics

Here I want to consider the claim of (Indexicality 1) – that is, the claim that the quantificational force of generics vary. All theorists as far as I am aware treat the quantificational force of generics as stable across all generics. Most theorists think *Gen* has universal force (cf. Pelletier and Asher (1997); Asher and Pelletier (2012), Asher and Morreau (1995), Asher and Pelletier (2012), Nickel (2008, 2012a)), while some think that it has the force of *most* (cf. (Cohen, n.d., 2011)). There is also at least some reason to think that *Gen* could be treated as having existential force (cf. for example Nickel (2012a) and Leslie (n.d., 2008, 2007), and footnote 16 in Chapter 3 of this thesis).

Those that think that *Gen* has universal force appeal to cases such as the following:

(47) a. Primes are odd.
b. Two and two equals four.

[9] If one finds that *Gen* seems like a *prima facie* unorthodox expression to be an indexical, then it is worthwhile to note that several seemingly unorthodox expressions have been given convincing indexical semantics in the literature on context-sensitivity — cf. for example, Rothschild and Segal (2009).

c. A spinster is an old, never-married woman.
d. Books are paperbacks.
e. Humans are Asian.
f. Sharks never attack bathers.

Each of (47a.)-(47f.) not only sound as though they have universal force ((47b.) is the best to illustrate this), but their truth-conditions also seem to require it: For example, (47a.) is intuitively false even though all but one prime is odd; (47c.) is seems to express a linguistic necessity; and (47d.)-(47f.) are each false despite the fact that most members of the domain have the given property.

While those that think that *Gen* has the force of *most* appeal to cases like the following:

(48) a. Tigers have stripes.
b. Peacocks have brightly coloured tail-feathers.
c. Ducks lay eggs.
d. Lions have manes.

Each of (48a.)-(48d.) do not sound as though they have quantificational force that is as strong as a universal, and moreover, their truth-conditions don't seem to require it. (48a.)-(48d.) each have many instances of the domain which do not possess the given property.

The *universal*-camp doesn't have tidy accounts of (48a.)-(48d.), not to mention the existential cases mentioned above. They might for instance appeal to domain restriction to get the proportion of instances right, but this does nothing to explain the fact that their quantificational force just simply does not sound as strong as a universal.

Similarly, the *most*-camp doesn't have tidy accounts of (47a.)-(47f.). For instance, Cohen (2004) and Leslie (n.d., 2008, 2007) place constraints on the entire domain (instances and counter-instances), to account for the intuitively false ones like (47d.)-(47f.).[10] Not only do these accounts face many counterexamples, but they also fail to account for the fact that (47a.)-(47f.) sound like they have universal force.

My bold conjecture is that what accounts for the difference in force between the (47a.)-(47f.) cases and the (48a.)-(48d.) cases is simply that the quantificational force of *Gen* varies. Moreover, not only does the force vary between generics, but also it varies with the context of utterance given a single generic:

[10]Cohen's *homogeneity constraint* says that the property must hold homogeneously (in a certain sense of the term) across any salient partition of the domain. Leslie's contrastive approach says that the counter-instances must lack a salient alternative property

(49) a. Lottery tickets are losers.
b. **Context 1:** The null context.
c. **Context 2:** While warning a friend to not waste time and money on playing the lottery.

(50) a. Cats are black.
b. **Context 1:** The null context.
c. **Context 2:** A teacher asks the class: What colours are cats?

In normal contexts, (49a.) sounds intuitively false since it seems to require that all lottery tickets are losers for it to be true. By contrast, if a speaker utters (49a.) while warning a friend not to be wasteful and foolish, (49a.) seems intuitively true since most lottery tickets are indeed losers. The intuitive quantificational force of (50a.) varies as well between the two contexts given. In regular contexts, (50a.) sounds like a claim about all cats, whereas in other contexts, like Context 2, (50a.) sounds like a claim about merely some cats.

4.4.3 The Variable Lexical Restrictor of Generics

Here I want to consider the claim of (Indexicality 2) – that is, the claim that the lexical restrictor of generics varies with the context of utterance. Nearly all theorists treat the lexical restrictor of generics as stable across all generics.[11] The two most plausible candidates which have been proposed to act as the lexical restrictor of *Gen* are:

(Uniform) The lexical restriction is to *uniform* members of the domain.[12]

(Normal) The lexical restriction is to *normal* members of the domain.[13]

I will consider each of these candidate lexical restrictors in turn and argue that for each, there are, systematically, generics for which that lexical restrictor does is not adequate. Thus, none of the candidate lexical restrictors is empirically adequate. Since these are the best candidates and none of them is adequate, I conclude that the lexical restrictor of *Gen* varies as a function of the context of utterance.

[11] There are a couple of notable exceptions: the relevant quantification approach of Declerck (1991) and constraints on situations approach of Barwise (1993), Cavedon and Glasbey (1994) and Meulen (1986).

[12] Cf. Cohen (2011).

[13] Cf. Pelletier and Asher (1997); Asher and Pelletier (2012), Asher and Morreau (1995), Nickel (2008, 2012a).

UNIFORM: Cohen (2011) provides an account of the stable lexical restrictor of generics as a restriction to *uniform* members of the domain. He characterises the uniform worlds as:

> ... worlds among the metaphysical alternatives to the actual world that are preferred by a uniform ordering source. Such worlds share the history of the actual world up to the reference time. But from this point on, no significant change occurs: the future resembles the past. Hence, uniform worlds retain the stable properties of the actual world: Non-stable properties will change in the actual world, but will remain the same in a uniform world. Cohen (2011)

To simply discussion, if we understand quantification to be over a domain of situations (i.e., parts of possible worlds), then the uniform situations are situations which are parts of uniform possible worlds. Uniform worlds are metaphysical alternatives to the actual world whose future most resembles the past, relative to some reference time t (typically the time of utterance or some salient reference time). Thus, *Gen* quantifies over a domain of situations which are parts of worlds whose future most resembles the past (of the actual world), relative to the time of utterance (typically). According to Cohen, the uniform situations preserve the stable properties of the actual world. He provides the following example to support his claim:

(51) a. A computer computes the daily weather forecast.
b. A computer computes the main news item.

Example (51a.) is originally discussed in Carlson (1989) who provides the following description of its most salient interpretation:

> ... *the daily weather forecast* requires an intensional interpretation, where its meaning cannot be taken as rigidly referring to the present weather forecast, e.g. the one appearing in today's copy of the Times predicting light rain and highs in the upper thirties. (1989, p.179)

Cohen reiterates that supposing today's weather forecast predicts a severe storm and is accordingly the main news item, we cannot then conclude (51b.). His diagnosis is that we accept the generic in (51a.) as true because computing the daily weather forecast is a stable property of situations in which computer forecasting takes place, whereas computing the main news item is not a stable property of situations in which computer forecasting occur. But other examples just don't fit Cohen's description — consider again (9) for example:

(9) a. Supreme Court judges have even social security numbers.
b. **Context:** Two friends are planning a party which will take place in three days time where they want all the party guests' social security numbers to add to an even number. Suppose further that all the Supreme Court judges will be changed in four days time and that many of them will not have even social security numbers.

In the given context, (9a.) is intuitively true, and yet having an even social security number is certainly not a stable property of the Supreme Court judges. According to Cohen, (9a.) should be false; he does not have an account of utterances of generic sentences like (9a.) in the the given context. Since these are perfectly normal generic utterances, I am left to conclude that restriction to a uniform lexical restrictor will not adequately provide an account of the lexical restriction of generics in full generality. On the indexical approach, a uniform lexical restriction is only one of many ways in which the lexical restrictor of generic utterances can be specified in context.

NORMAL: The dominant view amongst theorists is that the lexical restrictor of generics is a restriction to *normal* members of the domain. Here I will argue that a restriction to normal instances is not empirically adequate: Though a restriction to normal instances might reflect the content expressed in some contexts, it doesn't suffice for all generic utterances.

Consider again what Nickel (2008) says about example (24):

> (24) Dobermans have floppy ears.
>
> The important fact about dobermans is that they are born with floppy ears that breeders then cut to given them the pointy shape we are familiar with. In the context of evolutionary biology, [(24)] is true. The text [(25)] certainly sounds acceptable.
>
> (25) Some breeds of dogs have evolved to focus on their hearing. These breeds have pointy ears. Dobermans, however, mostly rely on their sense of smell, which is why Dobermans have floppy ears.
>
> However, in the context of a discussion of dog breeding, [(24)] seems clearly false, as the text [(26)] illustrates.
>
> (26) While Labradors and golden retrievers have floppy ears, dobermans don't. Dobermans have pointy ears.

(2008, p.644)

Nickel (2008) provides an explanation of the variance:

> On the view I am proposing, whether or not [(24)] expresses a truth depends on the kinds of influences we countenance and which we abstract from, specifically how we treat the interventions of dog breeders. That, in turn, depends on the kind of inquiry we are pursuing. Each of the texts makes a certain inquiry salient, which is why [(24)] expresses different propositions when embedded in each.

According to Nickel, the truth of (24), and generics more generally, depends on certain features of a given inquiry which is salient in the context of utterance. Very roughly, a generic *Ks are F* is true if:

(N) All normal Ks are F.

Moreover, crucially, what constitutes being a normal K is determined by features of the inquiry which is salient in the context of utterance. In the example, the idea is that the salient inquiry determines different domains of normal Dobermans over which to quantify — for (25), it is the biologically normal Dobermans, and for (26), it is the normally bred Dobermans.

Nickel's account provides the closest candidate for a distinctive and stable lexical restrictor across contexts. Supposing Nickel is correct at least for some generic utterances — i.e., that the lexical restrictor is in some cases a restriction to the normal instances. Then this counts as one way in which the lexical restriction of generics can be specified in context. But restriction to the normal instances is not the only way.

To me, normality, even a context-dependent form of normality, is not a stable component of the generalisations expressed by generics — to see this, consider the following discourse involving an explicit question (inquiry):

(52) A: What's the difference between humans and monkeys?
B: Humans kill themselves.

(53) A: What's the difference between humans and monkeys?
B: All normal humans kill themselves.

(52B) sounds acceptable and yet it seems strange to say that normality plays a role in whatever generalisation is expressed. As evidenced by (53) — there seems no way to repair (53B) so that it sounds acceptable.

Both uniformity and normality do not suffice to account for the full range of available lexical restrictions for utterances of generic sentences. Since these are the best candidates, I conclude that the lexical restrictor is not a stable component of the general content expressed by generics and rather the lexical restrictor varies with the context of utterance. In the rest of this section, I go on to give some considerations in support of treating *Gen* as an indexical.

4.4.4 Unpronounced

Several authors (cf. Krifka et al. (1995), Leslie (n.d., 2008, 2007) and Liebesman (2011)), I think probably rightly, make a big deal out of the need to explain why *Gen* is unpronounced:

> The unarticulated nature of *Gen* is puzzling, and has not received adequate attention in the literature ... it is worth noting that this indicates that *Gen* is not just another everyday quantifier/determiner/adverb, which just happens not to be articulated in English. Were this some accidental fact about English, we would expect to find various other languages in which *Gen* was phonologically realised on a par with other quantifiers, determiners, or adverbs. That the non-articulation of *Gen* appears to be a linguistic universal is an interesting fact that should not be neglected. (n.d., p.27)

> Standard glosses of *Gen* such as *typically*, *usually* and *generally* fail in cases like [(11)] and [(48c.)], as well as a host of others. The absence of any pronounced English adverb synonymous with *Gen* is striking, though not nearly as striking as the cross-linguistic data. According to a number of theorists, no known language has a pronounced *Gen* operator. (2011, p.414)

Leslie explains the unpronounced nature of *Gen* by providing a theory of *Gen* on which natural languages have no need for a pronounced generic A-quantifier — in particular, she claims that *Gen* expresses our default mode of generalization, which is invoked in the absence of explicit modes. Liebesman on the other hand, uses the fact that *Gen* is never pronounced to argue that it doesn't exist — on his view generics are kind-predications. On the view I will be proposing, *Gen* is unpronounced because it is an instance of a certain type of expression in natural language which has a tendency to be unpronounced: *Gen* is an indexical which is represented as a free variable at the level of logical form, and many

free variables are never pronounced. Such variables have important and semantically significant roles and therefore, are represented as covert structure in logical form despite the fact that they are never unpronounced. There are several concrete examples: Quantifier domain variables, implicit argument places and *PRO*-constructions. Whatever explanation(s) apply in these cases (if any), might very well apply in the case of *Gen* understood as a free variable. An explanation for why *Gen* is unpronounced can be understood as part of a more general explanation of why covert expressions of this sort have a tendency to be unpronounced.

The fact that there exists an explanation of such a tendency is a consideration in favour of the indexical approach. Other quantificational approaches do not have an explanation: If *Gen* has a stable content, why wouldn't speakers utter that content just like they do for other quantificational expressions? Moreover, giving generics a disquotational analysis (as Leslie does) or a kind-predication analysis is empirically inadequate.

4.4.5 Underdetermination

The second consideration in favour of the indexical approach is that it supplies an explanation of a form of underdetermination that is displayed by generics. If one thinks about how, in a given context, a speaker determines what the extension of *Gen* is — i.e., what generalisation *Gen* expresses, one quickly realises that there isn't an easy answer to give. There is often times more than one plausible candidate for the extension of *Gen*, for a single generic utterance[14] Take the following examples in the given contexts:

(54) a. Norwegians have blue eyes.
b. **Context:** A speaker asks *What do Norwegians look like?* and someone responds with the given generic.
c. Characteristically actual genetically original Norwegians have blue eyes.
d. Typically actual biologically original Norwegians have blue eyes.
e. Significantly many Norwegians have blue eyes.

(55) a. Mammals lay eggs.
b. **Context:** A biologist is discussing birds, and their relationship to other species, she utters *Birds lay eggs. Mammals lay eggs too.*
c. There is a homogeneous subset of mammals such that all of them lay eggs.

[14] Never mind the additional levels of underdetermination of generic content between contexts (as evidenced by examples (43)-(45)) and between different generic sentences (as evidenced by examples (54) and (55)).

d. Several mammals that reproduce in some way lay eggs.
e. Many mammals that have reproductive capacities lay eggs.

If a speaker utters the generic (54a.) in the context specified in (54b.), then there are many available interpretations for what the speaker said, some examples are given under (54c.)-(54e.). Similarly for (55).

Quantifier domain variables, implicit argument places and *PRO*-constructions display a similar type of underdetermination (cf. for example REFS). As an example consider the following from Buchanan (2010):

> An hour before the party is to begin, Tim asks Chet 'Are we ready to rage?', 'So bro', Chet responds, 'We are totally ready. The living room totally looks like a pirate ship. The strobe lights are up. Every beer is in the bucket. I just need to find an eye patch to wear with this pirate hat.' Consider [the following]:
>
> (56) Every beer is in the bucket.
>
> ...
> Even though the case is somewhat under-described it is clear that there are many equally good, yet non-equivalent, candidates for what Chet said in uttering (56). For example, the following possibilities spring to mind:
>
> > Every beer we bought at the bodega is in the bucket in the backyard.
> > Every beer we will serve at the party is in the bucket decorated in pirate motif.
> > Every beer for our guests is in the bucket filled with ice.
> > Every beer at the apartment is in the bucket next to the hot tub.
> > ...
>
> As well as combinations of the foregoing:
>
> > Every beer we bought at the bodega is in the bucket next to the hot tub.
> > Every beer at the apartment is in the bucket in the backyard.
> > ...
>
> (2010, p.349)

The answer of how the extension of such expressions gets fixed in a given context is a foundational problem of context dependence (cf. for example, Stanley and Szabo (2000) for elaboration); a problem for the metasemantics — the descriptive semantic theory might be able to specify some constraints on how to fix the values of such expressions in context, but there is likely nothing beyond that a semantic theory can and should contribute. The indexical approach allows us to delegate the explanation of the underdetermination displayed by generics to the foundations of context dependence in natural language, or the metasemantics of natural language if you like. This, I think, is the right place to be searching for answers to the many puzzling features of generics. It is worth noting that other theories have not offered any explanation of the type of underdetermination exhibited by (54a.) and (55a.), or why *Gen* should differ from other quantifiers in respect of the underdetermination of their quantificational force and/or the underdetermination of their lexical restrictor.

4.5 Gen: The Indexical A-Quantifier

Theorists often think generics are special in some way — its's just hard to pinpoint what their specialness is. On the indexical approach generics are special because *Gen* is the only quantifier of its kind. It is standardly agreed that most quantifiers have context dependent components — e.g., contextual domain restrictions, standards for vague quantifiers, etc. But on my proposal *Gen* is unique, a very special kind of A-quantifier, an *indexical A-quantifier*.

To be clear about what I mean by this it will be useful to go through an example of why even contextually based models of vague quantifiers, like *often* or *seldom*, are not indexical quantifiers like *Gen*. Take *often*; on a standard contextual analysis,[15] the sentence *often Ks are F* or *Ks are often F* has the quantificational force of the vague D-quantifier *many* and a fixed lexical restriction to actual members of the domain. The indexical component of *often* is a contextually supplied standard "oftenness". An important difference between vague A-quantifiers and *Gen*, understood as an indexical A-quantifier, then is that though a vague A-quantifier, like *often*, has an indexical component supplied by the fixed D-quantifier *many*, *Gen* has no such fixed D-quantifier. The quantificational force supplied by the D-quantifier in the case of *Gen* can vary, in the case of *often*, the force is given by the force of *many* which varies with the standard, but is nonetheless always the force of the D-quantifier *many*. *Gen* can have the vague force of *many*, but it can also have the force of non-vague quantifiers, like *all* or *most*, depending on context.

[15] Cf. for example, Cohen (2001b).

Vague quantifiers are special devices of language, but *Gen* is a different type of device — it has a specialness all of its own.

Section 4.6 will present some of the ways in which this specialness manifests itself; in this section I present the outline of a semantics.

I propose that we give a semantics for generics which combines a standard semantics for A-quantifiers with the idea that *Gen* is a free variable which changes its extension across contexts. In doing so, we can say that *Gen* behaves in some ways like a simple indexical, akin to *I*, *she*, *he*, *here*, *now*, *this* or *that*. My goal is not to try to defend a specific proposal about how the semantics is implemented, but rather to show that there is a plausible linguistic implementation along these lines.

On a situation-theoretic semantics for A-quantifiers, my proposal is that both A-quantifiers — e.g., *always*, *usually*, *typically* and *normally* — and *Gen* involve quantification over a contextually variable domain of situations. As such, they both express quantitative relations between a domain and a scope. Further, both A-quantifiers and *Gen* are restricted lexically, contextually and sometimes explicitly to a domain of situations; and their scope is invariably given by an overt clause. But A-quantifiers and *Gen* differ in how their lexical meanings are assigned: While the quantificational force and the lexical restriction for a given A-quantifier is fixed, the quantificational force and the lexical restriction for *Gen* is not — the quantificational force and the lexical restriction for *Gen* varies across contexts. As such, *Gen* will be represented by a variable at the level of logical form. It is due to this final point that I call *Gen* an *indexical quantifier*. This proposal is in some ways standard, in that it preserves the idea that generics are A-quantified sentences, and in some ways innovative, in that no A-quantifier has ever been conceived of as pure indexical and further, no theory of generics has been proposed which endorses straight-out contextualism about genericity. In what follows, I attempt to provide what might be conceived of as a *character* for *Gen*, in the Kaplanian sense, though the constraints proposed below might also be conceived of as *semantic presuppositions* or *presupposition triggers*, or indeed in some other way. I simply call them *metasemantic constraints*.

4.5.1 Metasemantic Constraints on Gen

The indexical approach is at this stage merely an approach: I've so far said nothing about how speakers resolve a semantic value for *Gen* – i.e., get from an utterance of a generic in a particular context to a particular semantic value of *Gen*. One might worry that the approach is simply too unconstrained and that this renders the task too cumbersome for speakers.

It is worth noting that despite the various sources of context-sensitivity, there are several metasemantic constraints imposed already: *Gen* is an A-quantifier and has A-quantificational structure, like other A-quantifiers there are constraints on what counts as the C-restrictor — for instance, as mentioned above, the relevance of the domain and a dependence of the domain on the question under discussion are both metasemantic constraints on the selection of the C-domain in context. Further, in a similar fashion, I think there are dedicated metasemantic constraints on the selection of the quantificational force and the lexical restrictor of *Gen*.

It is also worth noting that other indexicals rely solely on metasemantic constraints in fixing their semantic values in context. There is lots of flexibility as to how the semantic values of these indexicals get fixed, but there are still some constraints (perhaps quite complicated) on what values are allowable. Consider the demonstrative *that* for example. There is no limit on what a speaker can demonstrate — anything from the solar system to a number, but in a particular context, an utterance of *that* cannot denote something that the agent did not demonstrate or did not intend to demonstrate. Just how to spell out the metasemantic constraints for *that* (i.e., just what it is to demonstrate and to intend to demonstrate) is complicated. My proposal is that the kind of flexibility *Gen* encompasses and the nature of metasemantic constraints on *Gen* are like those for *that* and other similar indexicals.

I outline some constraints below, however, views that endorse the indexical approach needn't endorse any one of these constraints. There is a lot of flexibility here and one can distinguish different views based on the kinds of semantic properties one would like generics to have. The below constraints merely gesture at some properties they might plausibly possess.

As an alternative to the below, one could, for instance, endorse a cognitively based picture similar to what Leslie (2008, 2007) proposes and build in constraints based on the types of cognitive features Leslie argues are central to generic generalisation — as an example, Leslie suggests that (11) is apt as a generalisation since the correctness conditions of our primitive cognitive mechanism of generalisation is in part sensitive to properties we judge as striking properties of the kind in question. As I see it, the relevant cognitive feature which judges strikingness, could play a role in resolving the semantic value of *Gen* as Leslie's proposed correctness conditions suggest: If the property in question — in the case of (11), carrying the West Nile virus — is striking, then the semantic value of *Gen* is the general content expressed in the clause *there exists some mosquitoes which carry the West Nile virus and those that do not are disposed to carry the West Nile.*

Ultimately, if there is anything general to say about these types of metaseman-

tic constraints, I think it is an empirical matter just what sort of constraints are the correct ones.

As a matter of methodology, I think it is useful to consider the different sources of variability outlined in Section 4.4 and to search for constraints imposed on these different individual sources of variability; taken as a whole, these will provide the constraints imposed on *Gen*. Despite the difficulty in saying something general about such metasemantic constraints, I propose some here.

Constraint 1 - Relevance of the Lexical Restriction: Recall that I argued above that the lexical restrictor varies with context. The first constraint on the fixing of the generalisation expressed by *Gen* in a given context is that the lexical restrictor must be *relevant*. Like in the case of the contextual restrictor, I think the QUD will play a role in understanding what counts as a relevant lexical restrictor.

This departs from other proposals which take the lexical restrictor to be context-sensitive — for example, Nickel (2008) who takes the lexical restrictor to be to normal instances, but thinks that what counts as normal varies with context; and Declerck (1991) who takes the lexical restrictor to be to relevant instances, but thinks that what counts as relevant varies with context. The constraint I have in mind, rather, takes relevance to determine, at least in part, what lexical restrictor is expressed. Sometimes the relevant lexical restrictor might be to normal situations, other times to relevant situations and still other times to typical situations.

In so far as relevance is modelled as deriving from representing discourses as structured salient inquiries (e.g., stacks of QUDs)[16], the notion of higher order notion of relevance advocated here can be modelled in a similar way. I have no clearer picture to offer about just how the relevant lexical restrictor is determined by the context. Note, however, that the picture offered here is no less clear than those offered for other context-dependent elements — it is no less clear than those offered by theories of quantifier domain restriction (i.e., how contextual restrictors get settled) or how the degrees or standards relevant to the interpretation of gradable adjectives get fixed.

Constraint 2 - Relevance of the Quantificational Force: The examples presented in Section 4.4.2. were evidence that the quantificational force of *Gen* varies with context. The second constraint I'd like to propose is that not only does relevance play a role in the selection of the domain (i.e., the selection of the value of the contextual restrictor and the lexical restrictor), but I think relevance plays a role in the selection of the quantificational force. In some contexts, what will

[16]There are several such models, my favoured can be found in Roberts (1998).

be relevant is that all instances of the domain are related to the scope, in some it will be most, and in some what will be relevant is that 75% of the domain is related to the scope, etc.

In this way, generics are analogous to certain other utterances involving quantity or a standards of precision. Consider for example:

(57) The train arrived at 3pm.

Someone might sincerely and felicitously utter (57) even though the train arrived at 3:02. Some standard is relevant to the interpretation of and evaluation of that claim. Just how that standard enters into interpretation and semantics is a disputed topic (cf. for example, Wilson and Sperber (2002)), but not one I mean to take a stand on here. What is uncontroversial is that somehow or other the context of utterance is such that a standard is selected and figures essentially in the interpretation and evaluation of the claim. My suggestion is that an analogous mechanism is involved in the interpretation and evaluation of generics. Something similar, I think is going on with the relevant standard required. What counts as the relevant standard for the generalisation for generics will vary from context to context — sometimes it will be all instances of the domain that need to satisfy the scope, other times simply most, and other times it might be 75%.

Similarly, on contextualist accounts of vague quantifiers, like *many* and *often*, some standard of "manyness" or "ofteness" relevant to the evaluation of the utterance is supplied by the context of utterance. As outlined above (cf. the introductory remarks of this section), *Gen* is even more flexible than vague quantifiers, but that's not to say that some of the same metasemantic resources are not used. I think the same types of constraints that apply to the interpretation of vague quantifiers can also be perspicuously applied in the case of resolving a semantic value for *Gen* in context.

I have no clearer story to offer about just how the relevant quantificational force of *Gen* is settled in context, again, the story on offer here is only as clear as those on offer elsewhere.

Constraint 3 - Generic Generalisation Requires a Majority: It is sometimes assumed that generic generalisation does not require a majority because of sentences like (11), (6) and (7). The thought is that these generics can be true even though only a minority of the domain satisfies the scopal property. I argue in Chapter 2, that this conclusion is mistaken for this class of generics where implicit cognitive biases are likely playing a role in mistaken attributions of truth-values. We should not take the intuition that these cases are true at face value.

On further reflection, it becomes clear that this initial intuition is mistaken, (11), (6) and (7) are in fact false. I won't repeat those arguments here. In Chapter 2, I argue for this conclusion based on examples like the following:

(58) a. Mosquitoes carry the WNV, but typically they don't.
b. Generally humans don't kill themselves, but humans kill themselves.

The conjunction in (58a.) and (58b.) sound contradictory or at least bad, but they could turn out to be true on an account which allows minority readings of generics. This is undesirable and good reason to accept that generics require a majority of the domain satisfy the scope.

The structure of my view leaves open the possibility of allowing minority readings, but for reasons outlined in Chapter 2, I think this is wrong for these cases — there may be other minority cases which warrant allowing minority generalisations, for example, the *distinctive property* interpretation discussed in Chapter 2 or *existential generics* discussed in Chapter 5. On the constraint proposed here however, any generalisation expressed by a generic requires that a majority of the domain satisfy the scopal property, but just what majority varies with the context of utterance (see Constraint 2 above).

Constraint 4 - Evidence and Explanatory Strategies: One important aspect of the whole business of making generalisations is from what grounds we can say a particular generalisation holds, given facts about individual instances. There is the issue of what is required to hold of individual instances in order for the generalisation to be true; and there are also in the case of generics especially, the confounding and frustrating issue of knowing what is sufficient knowledge of the instances to warrant asserting a generic generalisation or under what circumstances we accept them as true. Chapter 2 of this dissertation makes it clear that drawing the line on what subjective factors to allow into the content of a given generic utterance is particularly difficult.

What is clear is that some generic generalisations seem to require quite subjective and weak evidence for their truth/acceptance while others by contrast seem to require quite objective and strong evidence for their truth/acceptance — contrast (47a.) with (31). What does seem invariant across utterances of generic generalisations, however, is that the truth of the utterance depends on some contextually negotiated evidential grounds, speakers interests and their stardards of precision. I think a complete metasemantic story should impose constraints on the selection of a suitable generalisation expressed by *Gen* based on whatever ev-

idential grounds, interests and standards of precision are relevant and salient in the context of utterance.

What evidence speakers have, what their interests are and what standards of precision they have play a role in selecting an appropriate generalisation expressed by *Gen*. Speakers with weak grounds will typically intend to express weak generalisations, speakers with only superficial interests will intend to communicate weak generalisations and speakers with low standards of precision will intend to express weak generalisations. By contrast speakers with strong grounds will typically intend to express strong and precise generalisations. All these factors play a role in determining what generalisation is expressed by a given generic utterance — what is expressed needs to be at least in the ball park of true, before any truth-evaluation takes place.

The evidence in a particular context which is relevant to the interpretation and evaluation of the truth of a generic utterance could be a shared property which causes the instances to have the property F (cf. Greenberg (2007) on in-virtue-of generalisations), a shared disposition to have the property F (cf. Leslie (n.d., 2008, 2007, 2013) on "Type B troublesome generics"), a shared capacity to have the property F (cf. Nickel (2010c)), a shared function or purpose which is to F (cf. Krifka et al. (1995)), or whatever sets of propositions adequately satisfy *because* clauses.

A constraint which incorporates evidential grounds might explain some otherwise puzzling features of generics. For instance, it might explain why speakers sometimes disagree with generic claims by appeal to reasons like that the corresponding disposition or capacity attribution is not true. Take for instance:

(59) A: Sharks attack bathers.
B: No they don't. Sharks are only disposed to attack bathers if they can smell blood in the water.

(60) A: Peugeots go 200mph.
B: No, they can only go 180mph.

If the generalisation expressed by generic in (59A) has as part of its content that something that entails shared disposition, then this permits B to disagree with the claim by appeal to the fact that the sharks in the domain of the generalisation do not have the shared disposition in question. Similarly in the case of (60), but in this case it is the capacity attribution which underlies the disagreement. Note that it is not in general possible to disagree with a given generic claim with the corresponding disposition attribution, capacity attribution or the like, it is merely

possible some of the time — the fact that evidential grounds and the generalisation expressed vary explains these puzzling tid-bits of data.

I again do not have a fully fleshed out theory of how this constraint works — that goes well beyond the scope of this paper. Nickel (2012a) in his discussion of a context-dependent form of normality appeals to what he calls *explanatory strategies* to explain some of the issues discussed above. It is worth noting that a theory of explanatory strategies could be used as the basis of the types of metasemantic constraints I have in mind.

My proposal has several limitations. The semantics does not explain how, in a given context, a speaker determines what the extension of *Gen* is. On my proposal this is an instance of indexical resolution. *Gen* is like *that*, in that it is a non-trivial task to explain how the term acquires an extension in a given context. However, what I have done is attempt to give four semantic constraints on the resolution of the extension of *Gen* in a given context. The indexical approach does not stand or fall with these constraints, nonetheless, I think constraints exist and the constraints proposed here are plausible.

4.6 Some Virtues of the Indexical Approach

The two main components of the indexical approach, A-quantification and representing *Gen* as a covert free variable (i.e., a covert indexical), are explanatorily powerful along many dimensions. First, the A-quantificational component allows *Gen* to be put to a tremendous amount of theoretical work and allows the explanation of the many features generics share with explicitly A-quantified sentences: Focus-sensitivity, question-sensitivity, and distribution facts to name a few. Second, the indexical component allows the explanation of SDM and SCS, and a host of properties which generics are often thought to have.

Here I will focus on the latter — i.e., on how treating *Gen* as an indexical can be used to explain some of the properties generics are often said to possess. The explanatory virtues of treating *Gen* as an A-quantifier are well-known and so I merely refer the reader to Krifka et al. (1995), Rooth (n.d.), Beaver and Clark (2008) and Cohen (2001b). Further, it should be obvious how treating *Gen* as an indexical helps explain SDM and SCS.

I will outline four explanatory virtues of the indexical approach:

- How treating *Gen* as an indexical is useful in explaining how generic communication and cognition play a role in basic information gathering and inductive inference.

- How understanding *Gen* as an indexical might justify the ubiquity of generics in our thought and talk.
- How the indexical approach solves the main problem associated with finding an adequate theory of genericity — the *Problem of Variability*.
- How understanding *Gen* as an indexical provides a route to explaining certain empirical results on generics and the mind — in particular, the *Paradox of Acquisition* and results on the recollection of quantified statements.

Elucidating how *Gen* is explanatorily powerful is wonderful, but of course does not constitute a knock-down argument for the view. The arguments for the view were presented in Section 4.4; in this section, I present additional considerations in favour of the view and point towards ways of spelling out the interfaces of generics with a host of other subject-matters. To this end, I will consider each from the list immediately above in turn.

4.6.1 Basic Information Gathering, Generalisation and Generics

Some of the considerations mentioned under Constraint 4 outlined above, I conjecture, have an important role to play in understanding generics as a basic tool for information gathering and as a basis for inductive generalisations.[17] This constraint is roughly that whatever generalisation the generic utterance at issue happens to express in a given context, that generalisation is dependent on information and evidence that is available in the context — that information could be evidence, statistical or otherwise, or is often explanatory in nature, especially in circumstances of incomplete information or uncertainty.

Further, if *Gen* really is the major device for expressing basic general information, then it doesn't seem natural to assume that given the wide variety of kinds (or groups of objects) and properties we inquire about, that the basic relation we seek to identify in our inquiries is a single relation across all kinds and properties.[18] Still further, given the wide variety of inquiries we undertake, and the wide variety of settings and circumstances these inquiries take place, it seems unnatural to assume that a single relation, i.e., generalisation, is what we seek to identify in all inquiries and circumstances. The indexical approach allows basic information gathering to seek to identify a any number of relations (i.e., generalisations)

[17] Another way to think of generics as basic is that they are *hypotheses*.

[18] Leslie (n.d., 2008, 2007) gives disjunctive "worldly" truth-conditions for generics. On her proposal there are four disjuncts, which indicates she would think there are (at least) four possible basic generalisations expressed by generics.

between a given kind, property, and over different inquiries: *Gen* expresses different relations across different contexts depending on a wide variety factors — e.g., the evidence available in the context, the inquiry being undertaken, the interests of the speakers, the level of precision relevant to the inquiry, the explanatory strategies in play, etc.

4.6.2 Ubiquity

Not only does the indexical approach have an explanation of the above, but it also has an explanation of the more general property of the ubiquity of the use of generics in everyday speech, communication and reasoning.

Generics are ubiquitous because they are useful. One can easily imagine how useful a tool *Gen* is, when understood as an indexical which expresses different generalisations in different contexts. First, one often knows that some generalisation within a certain range holds without knowing the precise generalisation which does hold. It may be that as a community we simply do not or cannot know which generalisation within this range holds precisely, but it is still useful to know that some one of these generalisation holds — this information allows us to make inferences that we wouldn't otherwise be able to make. Second, it is often impractical or irrelevant to state a precise generalisation, and in such instances it is useful to have the means to express that some generalisation holds we just don't care which one — generics allow us to do this. Finally, sometimes speakers simply do not want to be committed to a particular generalisation and prefer to leave their commitments open — generics allow speakers to be committed to some generalisation holding, without being committed to a particular one.

Having said all this though, I do not mean to commit myself to the claim that generics do not express particular generalisations in particular contexts — they do. As I tried to stress above, generics are like other context-sensitive expressions. Take *smart* for example: It is not always the case that the speaker knows precisely what the relevant degree or standard of smartness is. These matters are settled by metasemantic and pragmatic mechanisms that need not be in the speakers grasp. Nonetheless, in context, *smart* and, I claim, *Gen* can express specific contents.[19]

[19]The issues raised here are strongly related to underdetermination which is discussed above. I do not wish to take provide an account of underdetermination here, only to point out some potential benefits of having it.

4.6.3 The Problem of Variability

Endorsement of SDM gives an automatic solution to what most authors see as the main problem of generics — what I will call the *Problem of Variability*. The least committal way of expressing the problem is probably as follows: How can we specify the distinctive content (or a distinctive set of contents) that generics like (1)-(12) are meant to share while at the same time accounting for the variety of ways in which their intuitive truth-conditions vary? I think such a formulation of the problem of variability leaves the source of the variability open. But the problem has been articulated in numerous forms, most often as problem of *variable exception toleration* — for example, Pelletier and Asher (1997) characterise it as follows:[20]

(61) a. A potato contains vitamin C and amino acid.
b. Potatoes contains protein.
c. The potato contains thiamin.

...

> A crucial feature of characterizing sentences — indeed, perhaps their most interesting semantic feature and that feature which brings them to the attention of this volume — is that they "admit exceptions". That is, [the] sentences in [(61)] are true even though there are potatoes which do not contain those nutrients... (1997, p.1128-9)

In this sub-Section, I will discuss why I think the Problem of Variability is not merely a problem of variable exception toleration (i.e., domain variability).

The first argument to this effect was offered in Section 4.4.2. There I maintained that the quantificational force of generics varies between generics, but also between contexts of utterance. A corollary of this is that the Problem of Variability is not merely a problem of exception toleration.

The second consideration, pace Leslie (2007), is that there seem to be many cases in which there is no suitable domain to restrict to the generalisation, or that even when restricted we do not get the desired interpretation or intuitive truth-value. We discussed some of those cases in Chapter 2. Good examples are the puzzling *distinguishing property* cases like (12). Domain restriction doesn't work seem to work for generics like (12) since if we consider the tiny number of Germans that actually wear lederhosen, the only thing that distinguishes them from the non-lederhosen-wearers is the lederhosen-wearing itself. There is no

[20] Cf. for example, Pelletier and Asher (1997), Cohen (2004), Greenberg (2007) and Leslie (2008, 2007).

systematic or significantly predictive set of features that pick out this subset. For that reason it is very hard to see how to limit the domain of quantification in (45) with some property F so that *all F-normal Scots wear kilts* or *generally, F-Scots wear kilts* comes out true or interpretable in the desired way.

The main puzzle of generics is better understood as a problem of variable truth-conditions, rather than merely a problem of variable exception toleration. The indexical approach has an easy solution to the Problem of Variability: On the indexical approach, variability is not a problem or a puzzle, rather variability is embraced wholeheartedly.

4.6.4 Generics and the Mind: Acquisition and Recollection

Before embarking on a direct discussion of how the indexical Gen fits in, I'd like to make clear an association between certain aspects of the indexical approach and the observations of some authors about the role of generics in our cognition. Thinking about *Gen* as an indexical can be understood as a way to explain another connection which several authors have postulated in order to resolve some of the puzzling features of generics. Carlson (2009), Leslie (2008, 2007), Cohen (2004) and others each postulate some connection between the mind and natural language generics:

> …taking a very broad and hazy perspective, there appears to be a striking confluence of interests between the study of concepts in the psychological and cognitive science literature, and the study of certain types of sentences in the formal semantics literature. Despite these initial confluences, the relationship between the meanings of generic sentences, and the study of conceptual structures remains unclear. (Carlson, 2009, p.17)

> … this suggests that our goal of understanding the quirky behavior of generics might be furthered by looking carefully into some of the relevant psychological literature. If the relation between objective statistical frequencies and the truth-values of generics is in some way mediated by psychological considerations, it would be useful to determine exactly what role these considerations are playing. (Leslie, 2008, p.16)

> …the generalisations that generic sentences express correspond to the cognitive system's most primitive, default generalisations. (Leslie, 2007, p.381)

> Findings of psychological experiments about the ways in which different domains are represented, and the factors affecting such representations, account for judgments of generic sentences, facts which cannot be explained by linguistics alone. … Generics express default rules, and psychological findings have shown that, the more homogeneous the domain, the easier it is for subjects to infer rules about it. Thus, cognitive results form a crucial part of a comprehensive account of the meaning of a linguistic expression, [*Gen*]. (Cohen, 2004, p.529)

All these authors think there are notable and interesting connections between genericity and the mind. Here I will consider how understanding *Gen* as an indexical can help us understand these connections.

If *Gen* is represented as an indexical, then in order to fix the semantic value of *Gen* in a given context, the metasemantic and pragmatic mechanisms of indexical resolution need to take place. This allows room for the psychological factors which concern these authors to enter into the picture. These psychological factors are involved in the metasemantic and pragmatic mechanisms at play in indexical resolution and hence, play a role in fixing the truth-conditions of generic utterances; as such, we can see the impact of such psychological factors on the truth-conditions of generics. Not only that, but it also provides a route by which to address cognitively based questions about generics.

I don't mean here to commit myself to a particular view of how indexical resolution takes place, there are numerous options: Sperber and Wilson (1995), Ludlow (2000), Recanati (2003), among others. I merely mean to point out that there is room to accommodate what many see as important observations in the literature connecting generics and the mind.

Acquisition: In Leslie (2008), one of the considerations motivating her cognitively based theory, and at least some form of endorsement of SDM, is what she calls the *Paradox of Acquisition*, which she characterises as follows:

> A puzzling question now arises: how does a language learner ever come to master generics? Not only is the interpretation of *Gen* rather complicated, the operator is not even phonologically realized… To make matters all the more puzzling, it happens that generics are acquired quite early on. Children start using generics by two years of age, which is significantly earlier than explicit quantifiers (Gelman 2003; Roeper, Strauss, and Pearson 2006). That children ever mas-

> ter generics is perplexing; that children master them more readily than explicit quantifiers borders on the paradoxical. This is a phenomenon that demands explanation. (2008, p.19)

On the indexical approach, there is a simple explanation for the Paradox of Acquisition. Children do not need to learn what *Gen* means since it has no fixed meaning. *Gen* is like other context-sensitive expressions — for example, *that*, *smart* or *tall*, or quantifier domain variables. The acquisition and understanding of gradable predicates like *smart* and *tall* are a matter of having the relevant metasemantic and pragmatic skills which govern the use of the relevant context-sensitive expression. Moreover, though theorising about the acquisition of such skills is a quite unsettled topic, it is natural to expect that metasemantic and pragmatic skills are acquired very early on (e.g., children use *that* and gradable predicates quite early on). Many theorists think that these types of skills are already in play in early word acquisition (cf. for example Tomasello (2006, 2001)).

To get a sense of what would be required of a child language learner, we can compare what Leslie (2008) proposes is explicitly required with what the indexical approach might explicitly require. According to Leslie:

> The child's innate language endowment would then only need to provide the learner with three principles for the acquisition of generics to proceed:
>
> 1. All variables must be bound for an LF to be interpretable.
> 2. Variables that are free in the Restrictor are bound by a default operator.
> 3. This default operator invokes the conceptual system's default means of generalising.
>
> If the acquisition of generics proceeded even roughly along these lines, we can begin to see why generics are so easy to acquire; they correspond most closely to what the child already knows how to do. (2008, p.28-9)

On the indexical approach, acquisition would need something along the following lines in order for the acquisition of generics to proceed:

1. All variables must be bound for an LF to be interpretable.
2. Variables that are free in the Restrictor are bound by a variable of the semantic type of an A-quantifier.

3. This variable invokes a process of indexical resolution which employs broadly metasemantic, pragmatic and conceptual mechanisms to arrive at a generalisation expressed.

Notice that just as in the case of Leslie's proposal, on the indexical approach we can see why generics are easier than explicitly quantified sentences to acquire — no explicit semantic knowledge of quantifiers is needed for the child to be competent with generics.

One might have the following worry, however; in particular, with regards to (3.) above and about the analogy between *Gen* and other indexicals (or indexical containing expressions), like *that*, *smart* and *tall*: Couldn't it just be that the metasemantic and pragmatic skills needed to acquire *Gen* are different from these expressions, and that the relevant skills relevant are acquired much later in the case of *Gen*? It is obviously impossible to say at this point, though I see no reason to think that so long as the grammatical and conceptual apparatus of A-quantification is available, why the resolution of variables would be any problem. There is at least no principled reason to think that so long as the grammatical and concept of A-quantification is there, resolution of a variable of this sort couldn't take place. Moreover, there is empirical evidence that the grammatical and conceptual apparatus of A-quantification is in place: Many theorists think the acquisition path of D-quantifiers proceeds, grammatically, via A-quantification. The evidence for this is that children often interpret D-quantified sentences as if they were A-quantified — cf. the phenomena of *quantifier spreading*. This would indicate that the grammatical (and perhaps conceptual) apparatus of A-quantification is acquired before that of D-quantification. Since it is the grammatical and conceptual aspects A-quantification that is needed to acquire *Gen*, this at least indicates that the grammatical and conceptual resources needed on the indexical approach are available. Some relevant references about the phenomena of quantifier spreading are: Roeper, Strauss, and Zurer Pearson (2006), Roeper and Matthei (1974) and Roeper and Villiers (1993).

Since it seems the grammatical and conceptual resources are available, so long as the metasemantic and pragmatics skills required to resolve variables of the appropriate semantic type are available, the indexical approach has a straightforward solution to the Paradox of Acquisition. We have no reason to think such capacities are not there. Hence, on the indexical approach competence with generics is achieved through possessing the relevant metasemantic, pragmatic and conceptual competencies which are linked to resolving the values of free variables.

Recollection: In a recent cognitive psychology article, Leslie and Gelman (2012) present four studies with 136 preschool children and 118 adults, where the adults and preschoolers were both found to have the tendency to recall quantified statements as generics. They suggest that this tendency derives from the fact that explicitly quantified statements are more cognitively taxing than generics. Any theory of generics which can explain why adults and children have this tendency is better off. Leslie (2008, 2007) and Gelman (2009) have both proposed theories on which generics express cognitively default or fundamental generalisations, thereby explaining these recollection facts. The indexical approach can also explain these recollection facts by appeal to cognitively less taxing metasemantic, pragmatic and conceptual resources required for indexical resolution. It is natural to suppose that if a subject is unsure of or cannot recall an explicitly quantified statement, then that subject might rely on features of the context to supply the intended generalisation. Moreover, we would expect that the subject will not want to commit him or herself to an explicitly quantified statement. By uttering a generic, the subject can negotiate with his or her audience about what generalisation is suitable. Understanding *Gen* as an indexical allows the subject to rely on context and her audience; this makes generics less cognitively taxing and thus, explains the tendency for quantified statements to be recalled as generics. Note that the indexical approach, unlike Leslie's and Gelman's theories, does not commit itself to any claims about the existence of cognitively default or fundamental generalisations or that generics express such generalisations.

4.7 Objections and Replies

I've outlined several virtues of an indexical approach to generics, but there are also several worries that need to be addressed. In this section, I attempt to defend the view against objections.

4.7.1 Overgeneration

There is a genuine worry that the theory is just not constrained enough in that it radically overgenerates true readings of generics. If a given generic is allowed to express just about any generalisation in context, then doesn't it predict that many intuitively false generics are in fact true. A version of this objection is given in Krifka et al. (1995) against the account of Declerck (1991) where the domain of generics is fixed by relevance:

> Declerck adopted a principle which says that when a statement is made of a "set", the hearer will use his or her world knowledge to

> restrict the statement to just those members of the "set" to which it can be applied in a suitable way. ... One obvious problem with this approach is that the principle, as it stands, can easily justify all kinds of generic sentences — it is easy to find restrictions which would make any quantification as true. (1995, pp.45-6)

Of course, the objection is even more acute for the indexical approach since both the domain and the force of generics is allowed to vary — aren't we relying on context too much. To illustrate the objection, consider:

(62) a. Primes are odd.
b. Generally, primes are odd.

(62a.) is intuitively false in many contexts and yet it is intuitively true when we add an explicit A-quantifier like *generally*. How does the indexical approach ensure that the generalisation in (62b.) is not expressed and that an appropriate one is?[21]

First, it is worthwhile to point out that there are contexts in which (62a.) does express intuitively true generalisations. Consider, for example, a context where a student is looking at a blackboard with numbers on it, and that student is looking for primes amongst the numbers on the board. A helpful onlooker remarks (62a.). In such a context, (62a.) is intuitively true. Thus, (62) might indeed express something akin to (62b.) in some contexts. The question is rather how context makes sure the right generalisations are the output of indexical resolution in a given context. This is a metasemantic question.

Second, the metasemantic explanation of what constrains the choice of generalisation expressed that makes a given counterexample false in a given context will vary from counterexample to counterexample. In the case of (62a.) a plausible explanation is that in most contexts in which prime numbers are being discussed are quite precise mathematical contexts. In such contexts, the relevant standards which play a role in fixing the quantificational force (see discussion in Section 4.5) will be quite rigorous; as such the quantificational force will be

[21] Leslie (2007) discusses a different classes of examples (in arguing against the domain restriction of Pelletier and Asher (1997)):

(63) a. Dogs are poodles.
b. Mammals are cows.

All the same points below can be made for these cases, but I think these a bit less convincing to bring about worries of overgeneration: Where the kinds in question are well-defined and well understood, I think these cases just sound silly and false in normal contexts, particularly the null context.

very strong — akin to *every* — so that (62a.) comes out as false on conceivable specifications of the lexical restrictor.

Finally, as mentioned above, the semantics proposed in this Chapter leaves many questions unanswered. I have offered some constraints, but it is not the job of a descriptive semantics of indexicals proper to say how, in a given context, the extension of that indexical gets fixed — i.e., what it is to be the generalisation expressed by the measures of the context. On my account, this is a task for the metasemantics and pragmatics of indexical resolution — it is not a task of the semantics. As mentioned above, I see this as a virtue and not a vice of the proposed account — I have managed to isolate the difficult complexities of generics from the semantics and find the right home for investigating them. How indexical resolution is accomplished is an extremely complex matter — many theories propose general principles about how this takes place (cf. Sperber and Wilson (1995), Ludlow (2000), Recanati (2003), among others, as noted above), however, endorsing one of these goes well beyond the scope of this Chapter. What I can, again, say is that the resolution of *Gen* seems in no way different from the resolution of other indexical expressions, like domain variables, implicit argument places, the force of vague quantifiers, etc.

4.7.2 Binding Gen

On the view defended here *Gen* is a variable and thus, one worry is that there should be the possibility of direct binding of *Gen*. The indexical expression *she*, for example, has both unbound and bound uses as exemplified by (64) and (65) respectively:

(64) She is sitting next to Nora.

(65) Everytime Jim sits next to a woman, she smirks.

But there are not bound uses of *Gen* in the same way as there are for other indexicals, like she. This looks like a problem for the indexical approach.

It is not a problem, however. The simple reason *Gen* cannot generally be bound, but indexicals like *she* can, is because *Gen* is a variable of a significantly different semantic type. *Gen* is variable of the same type as an A-quantifier — i.e., *Gen* is a variable which takes as its value a relation between two propositions (according to the semantics given above). I think one should not be surprised that there are not sentences in which *Gen* is directly bound by a quantifier expression — there are very few expressions which can be directly bound by quantifiers.

4.7.3 Gen Doesn't Feel Like an Indexical

A more vague and yet common argument against my kind of indexical proposal[22] (and contextualist proposals more generally) is the claim that *Gen* does not "feel" like an indexical. More precisely, the worry is that *Gen* does not share certain essential features with simple and indisputably indexical expressions. There are numerous such essential features on offer — for instance, that it is somehow part of the meaning of the expression that it refers to an element (or parameter) of the concrete utterance situation (e.g., the speaker in the case of *I*, or the time of utterance in the case of *now*), that the speaker's intentions somehow play a role in the determination of the semantic value of the expression, that the expression serves a certain function or purpose (cf. for example Neale (2007)), or that there is some algorithm which determines the semantic value of the expression in context (i.e., the expression has a character in Kaplanian terms). I am not convinced that any such feature is essentially part of what it is to be an indexical expression — all of these claims I think are on empirically dubious grounds. There are many expressions which are simple and indisputably indexicals, but which do not have the features in question. Consider the examples cited above — *that*, *smart*, *often*. The only genuine such feature that I think indexicals share is that their extension varies with context and that it is the extension as a whole which varies. These are the features I've argued *Gen* has. Thus, I am unmoved by such worries.

Over and above the foregoing, even if some such feature is essentially part of what it is to be an indexical, this by no means shows that *Gen* does not have the relevant feature.

4.7.4 Other Views Consistent with SCS

Another way to treat the context-sensitivity of generics is as some sort of pragmatic phenomena. There are many ways to one could go about giving a pragmatic account. On any such account, a given generic communicates or conveys a generalisation, but does not literally express it. The most plausible such approach, I think, would make use of the paradigm of truth-conditional pragmatics in some form or another. Proponents of truth-conditional pragmatics would not represent *Gen* as a free variable at the level of logical form, thus, treat *Gen* as part of the semantically derived logical form of generics. Rather, *Gen* would be treated as an *unarticulated constituent* or as *freely enriched* (or *modulated*, whatever one's favourite terminology is) and its content derived by some purely pragmatic mechanism, rather than a metasemantic mechanism.

[22] Cf. Rothschild and Segal (2009) for a more extensive treatment of this type of objection.

To illustrate, consider a paradigm example of truth-conditional pragmatics in action — the phenomena of *deferred reference* or *metonymy*, consider:

(66) The ham sandwich wants his bill now.

Here is the diagnosis of (66) provided in Stanley (2005):

> Not only is there is a strong intuition that the deferred meaning is part of the intuitive truth-conditions, but the deferred meaning enters into certain linguistic processes, such as anaphora and ellipsis. For example, the natural reading of [(66)] is one in which the anaphoric element *his* receives its value not from the "literal" content of *the ham sandwich*, but from its deferred meaning. (2005, p.9999)

In the case of generics, the idea would be that though the content expressed by *Gen* is part of the intuitive truth-conditions of generics, it only enters into interpretation after some pragmatic process has taken place. Just as the content of *the ham sandwich* is part of the intuitive truth-conditions, but is pragmatically derived. We saw one view like this in Chapter 3.

I won't take a stand here against such approaches: I think both the indexical approach and such pragmatic approaches are interesting and novel proposals for the treatment of generics, and are worthy of attention. Note, however, that I have given some considerations in in favour of a primarily metasemantic picture in Chapter 3. There are also other considerations in favour of the indexical approach: It preserves the truth-conditional compositionality and also the idea that our intuitions about the truth-conditions of utterances provide genuine and sound data for semantic theory. Leslie (n.d., 2009), by contrast, thinks that semantics and semantic cognition need drastic revision in light of the data of generics. Other considerations in favour of the approach here are rehearsed by opponents of these pragmatic views (cf. for example, the overgeneration worry found in Stanley (2002a)).

4.8 Conclusion

In this chapter I've tried to defend the positive view that the unpronounced quantifier expression *Gen* is an indexical. The chapter argues that a given generic sentence expresses wholly different generalisations in different contexts of utterance, and a semantics which treats *Gen* as an indexical A-quantifier was given. The view advocated has many virtues which are outlined in the chapter and objections to the view are also addressed. The indexical approach gives a radically

new spin on the investigation of generics that I hope will inspire future work on generics along the variety of dimensions outlined in the chapter. For further consideration of the view, I direct the reader to the heading entitled "Further Considerations" below, there I provide new avenues to address two classes of generics not considered in this chapter. I also attempt to outline the connection to the topic of prejudice discussed in the "Consequences and Conclusions" section of Chapter 2.

5 Further Issues

This section provides extensive discussions of two important classes of generics — *existential generics* and *indefinite singular generics*. I don't have a settled view on either class of data and thus, have not included the content therein as main chapters. They are important topics to address, however, for anyone interested in providing a theory of generics. The goal of this final chapter is to outline the puzzle cases, say something about standard accounts, and sketch how the indexical approach in Chapter 4 relates to the data. I think a unified theory of genericity is to be preferred over one that deals with problem cases in piecemeal fashion, so my long-term goal is to deal with both classes of data.

Further Issues I: Existential Generics, Information Structure and Indexical Quantification

In first part of this chapter, I will showcase an interesting class of data put forward by von Fintel (1997) and Cohen (2003), called *existential generics*. I will develop more data of this kind, briefly go through Cohen's account of existential generics and how Cohen derives what he calls the *quasi-existential* interpretation of generics from their regular *quasi-universal* interpretation. After that, I will present some data that Cohen's account has trouble dealing with and show how even a more general account of the sort Cohen develops (which makes use of features of *information structure* and is in large part inspired by Beaver and Clark (2008)) still has trouble dealing with this new data. Finally, I will discuss how treating *Gen* as an indexical can help deal with the puzzling data.

5.1 Existential Generics

Cohen (2003) provides examples of linguistic environments where generics receive quasi-existential interpretations — consider the following examples where the generic is embedded under *even* and *only* respectively:

(1) Even [mammals]$_F$ lay eggs.

(2) Only [mammals]$_F$ give birth to live young.

It seems that the mere existence of some platypuses is sufficient for the truth of (1). But this is puzzling since *even* in (1) requires that its generic prejacent is true and it's not the case that mammals, in general, lay eggs. The generic is getting

what Cohen calls a quasi-existential reading. Similarly, in (2) the mere existence of some reptiles and fish that give birth to live young is sufficient to falsify (2). Despite the fact that *only* requires the falsity of *all* the generic alternatives to its prejacent (e.g., the falsity of *birds give birth to live young* and *fish give birth to live young*).

Cohen also considers cases involving *too* and *also*:

(3) Birds lay eggs. Mammals lay eggs too.

(4) Birds lay eggs. Mammals also lay eggs.

As with (1) and (2), it's sufficient for the truth of (3) and (4) that some mammals lay eggs.

There are some important assumptions that I will follow Cohen (2003) and Fintel (1997) in making about these cases. It is worth flagging these assumptions from the beginning. First, I am assuming that the relevant expressions (i.e., *even*, *too*, etc.) are assumed to be propositional operators, which in these examples are operating on a generic prejacent. For example, in the compositional semantics of:

(5) Only [professors]$_F$ are confident.

The expression *only* combines with the proposition expressed by the generic *professors are confident*. Second, I am assuming that In addition to whatever their semantics requires, each of the relevant expressions conveys that their (generic) prejacent is true. For instance, in our example, it is (somehow) conveyed that the prejacent *professors are confident* is true; so in (ordinary) contexts where (5) is true, *professors are confident* is also true.[1]

One can confirm the hypothesis that it is focus-sensitive particles which cause quasi-existential interpretations, by testing the interaction of generics with these other categories of focus-sensitive expressions. Here I provide further examples of quasi-existential interpretations using *particularisers*, *intensives*, other *additives*, *exclusives* and *downtoners*.

PARTICULARISERS: Further linguistic environments where generics receive quasi-existential interpretations are yielded by considering particularisers — consider the following examples:

(6) For example, [bodybuilders]$_F$ become important politicians.

[1] The problem evaporates without these assumptions, but there are good reasons for each — cf. von Fintel (1997).

(7) In particular, [stand up comedians]$_F$ become senators. Consider Al Franken's rise into politics.

In conversations about American politics, it seems the existence of Arnold Schwarzinager and Al Franken is sufficient to render (6) and (7) true, despite the generic prejacents on their own being false.

INTENSIVES: Consider further cases involving intensives and a generic prejacent:

(8) Significantly, [monkeys]$_F$ use tools.

(9) Significantly, [mammals]$_F$ lay eggs.

In a conversation where a scientist is discussing experimental results indicating biological capabilities, (8) sounds intuitively true, despite only a few monkeys in captivity displaying the relevant characteristic. Similarly, in a conversation comparing birds, reptiles and mammals, (9) is intuitively true, despite the fact that platypuses and echidnas are the only egg-laying mammals.

ADDITIVES: There are more cases involving *additives*:

(10) Birds lay eggs. [mammals]$_F$ lay eggs as well.

(11) Spas have jacuzzis. Likewise, [private jets]$_F$ have jacuzzis.

Just as with Cohen's examples (3) and (4), (10) and (11) only require the existence of a few egg-laying mammals and private jets with jacuzzis.

EXCLUSIVES: Exclusives also yield still further cases:

(12) Just [mammals]$_F$ bear live young.

(13) Exclusively [enrolled students]$_F$ use the university swimming pool.

(14) Solely [Argentinian winemakers]$_F$ make good malbecs.

Like *only*, exclusives like *just*, *exclusively* and *solely* yield existential generics.

DOWNTONERS: Finally, downtoners give rise to quasi-existential interpretations as well:

(15) Philosophers hardly smoke after [dinner]$_F$. They usually smoke in the mornings.

Downtoners like *hardly* when accompanied by a generic prejacent can also result in existential generics. (15)

5.2 Cohen's Account of Existential Generics

The basics of Cohen's account are as follows. Cohen uses a notion of *alternatives* as a primary ingredient necessary for determining the domain *Gen* quantifies over. As an example, consider:

(16) Mammals give birth to live young.

The domain of (16) is not all mammals: Male mammals, as well as females that are too young or too old, are irrelevant. Cohen accounts for this fact by making it part of the truth-conditions of generics that their domain is restricted by the disjunction of the alternatives to the predicated property. For example, the alternatives to giving birth to live young might be: giving birth to live young, laying eggs, undergoing mitosis, etc. As a result, the mammals that form the domain of (16) are only those mammals which satisfy at least one of the alternatives to bearing live young — that is, only mammals that are responsible for procreation in some way or another. Consequently, (16) is true since mammals that procreate in some way or another, in general, give birth to live young. A simplified version of Cohen's truth-conditions, then, are roughly as follows:[2]

> The generic *Ks are F* is true iff in general, if $K(x)$ and $\bigvee Alt_F(x)$, then $F(x)$.[3]
> (Where Alt_F is the set of alternatives to F.)

Both the selection of ALT_F and what belongs to ALT_F for a given F is influenced highly by context and features of information structure.

Cohen derives existential generics from their regular interpretation. He argues that in the cases at issue, alternatives are not induced. Not by focus and not by any other means. For example, in (1), where the generic *mammals lay eggs* appears embedded under *even*, a singleton set of alternatives containing only the property being predicated itself is restricting the domain of *Gen*. This makes the sentence trivially true whenever some instances satisfy the predicated property — i.e., it has an quasi-existential interpretation:

(17) a. $\text{ALT}_{lay\ eggs} = \{lay\ eggs\}$
b. In general, if x is a mammal and x lays eggs, then x lays eggs.

[2] The truth-conditions here do not reflect Cohen's actual semantics for *Gen* (he offers a probabilistic approach), rather these are an impartial heuristic which preserve the features which are relevant to the discussion.

[3] $Alt_F(x)$ is just the set of alternatives to $F(x)$ and $\bigvee Alt_F(x)$ is just the disjunction of all the alternatives in $Alt_F(x)$.

The focus in the sentence induces alternatives that are associating with *even* rather than *Gen*. This makes sense, Cohen argues, since the focus-sensitivity of *even* is lexically encoded and hence, association is mandatory. Generics, on the other hand, do not have to associate with focus — such association is optional. Moreover, the reason generics differ from explicit adverbs of quantification in receiving quasi-existential interpretations is that association is mandatory for the latter: In particular, focus is required to determine the logical form of adverbs of quantification. Focus is not required to determine the logical form of generics since these are initially interpreted as cases of direct kind predication. When this fails, *Gen* is derived with the kind-subect appearing as the restrictor automatically.

5.3 An Objection to Cohen's Account of Existential Generics

Here is what I take to be the primary claim of Cohen (2003):

(CLAIM) "[A common characteristic of all types of existential generic,] I propose [is] that in all these sentences, no alternatives are introduced. This means that alternatives are not introduced overtly, and, moreover, either there is no (narrow) focus, or the focused part is not associated with the generic quantifier." (2003, p.158)

In this subsection, I'll argue against (CLAIM) by providing data which indicates that (CLAIM) overpredicts quasi-existential interpretations.

Cohen's explanation predicts that we should get quasi-existential readings for the full range of focus-sensitive particles (whose focus-sensitivity is lexically encoded) whenever it is possible to combine them with generic prejacents; and further, that we should get quasi-existential interpretations no matter which generic the focus-sensitive particle is combined with. But, as I will demonstrate below, there are focus-sensitive particles that when combined with generic prejacents give rise to both quasi-existential and quasi-universal readings depending on which generic they are combined with, or in what context they appear.

Consider the following examples involving particularizers:

(18) a. Argentineans, in particular, [are vegans]$_F$.
b. ALT$_{vegan} = \{vegan\}$
c. In general, if x is Argentinean and x is vegan, then x is vegan.

(19) a. For example, policemen [are bisexual]$_F$.
b. ALT$_{bisexual} = \{bisexual\}$
c. In general, if x is a policeman and x is bisexual, then x is bisexual.

The sentences in (18a.) and (19a.) are intuitively false (in ordinary contexts) despite the fact that some policemen are gay and some Argentineans are vegans. In order for (18a.) and (19a.) to come out as true, the requirement seems to be that in general, to be gay is to be a policemen and in general, vegans are Argentineans, respectively — i.e., that the generics in question get a quasi-universal interpretation. Another example involving particularisers is as follows:

(20) a. For example / For instance / In particular / Specifically, [pet owners]$_F$ enjoy the company of dogs.
b. ALT$_{enjoy\ the\ company\ of\ dogs} = \{enjoy\ the\ company\ of\ dogs\}$
c. In general, if x is a pet owner and x enjoys the company of dogs, then x enjoys the company of dogs.

Here again, according to Cohen, (20a.) is receives the interpretation in (20c.) and hence is trivially true. But despite the particulariser, (20a.) receives a quasi-universal interpretation: The sentence is intuitively true, but not because a few pet owners enjoy the company of dogs, but rather because many pet owners enjoy the company of dogs.

Further, consider the following examples involving intensives:

(21) a. Especially / Importantly / Significantly / Well [cats]$_F$ enjoy the company of dogs.
b. ALT$_{enjoy\ the\ company\ of\ dogs} = \{enjoy\ the\ company\ of\ dogs\}$
c. In general, if x is a cat and x enjoys the company of dogs, then x enjoys the company of dogs.

(22) a. Especially / Importantly / Significantly / Well [boxers]$_F$ are smart.
b. ALT$_{smart} = \{smart\}$
c. In general, if x is a boxer and x is smart, then x is smart.

In (21a.) and (22a.), the intensive combines with the generic prejacent and the reading yielded is clearly false: cats are well known not to enjoy the company of dogs, even though there are some that do, and boxers aren't known to be a profession with the smartest of people, though some boxers are indeed very smart. On Cohen's predicted interpretations of (21a) and (22a), namely (21c) and (22c), they would incorrectly be labeled as true. This is surely not satisfactory.

There are various possible responses to the data one could take on behalf of Cohen — he could for instance respond that:

1. The data is just evidence that the focus sensitivity of such particles is not lexically encoded. In other words, association is optional.

2. The semantics of such particles is responsible for the false readings.

3. The semantics of GEN is responsible for the false readings.

4. The feature markings (or LFs) indicated in these examples is incorrect:

   – The focus is on the particle itself (and the associated element is a contrastive topic).

   – There are markings missing, that induce alternatives.

5. Association is optional.

None of these responses is sufficient, however. The bullet points explain why each response above is insufficient, respectively:

1. If the focus sensitivity of such particles were not lexically encoded, then why would quasi-existential readings be available the first place? Furthermore, the focus-sensitivity of such particles is well supported by linguistic evidence (cf. Beaver and Clark (2008), pp. 68-77). Moreover, there are examples involving the particles Cohen explicitly discusses — consider:

   (23) People enjoy the company of dogs. Cats enjoy the company of dogs too.

   In the same way that the examples above are false even though there are some instances of the kind that satisfy the predicated property in the quasi-existential way, (23) is false.

2. It cannot be that the semantics of such particles is systematically responsible for the false readings. To see this, take particularisers as an example. Particularisers merely indicate to the speaker that she has already given reason to believe that the prejacent is true[4] — this has no bearing on whether or not the generic receives a quasi-existential or quasi-universal interpretation. Moreover, the example (23) above involves the same particles Cohen discusses; and as demonstrated, all the particles at issue can receive quasi-existential readings in addition to the quasi-universal ones, their semantics is fixed across these cases, and thus, there is no reason to think that their semantics is responsible for the falsity in these cases.

[4] Cf. Beaver and Clark (2008, p.74)

3. This response seems incorrect on basic grounds. The issue is to determine how the two readings arise. In the counterexamples at issue a quasi-universal reading is preferred and hence, the falsity in itself is irrelevant. The meaning of *Gen* is stable across two contexts and hence, it cannot be the meaning of *Gen* which is playing a role.

4. This response interpreted appropriately is the most interesting and requires more consideration. In the next section, I take up the issue of complicating the features of information structure needed to get the different readings. I argue that such an approach won't work.

5. This response just seems *ad hoc*. And again, it doesn't seem to address the issue. Moreover, mere consideration of the alternatives does not always seem to be what is going on. At least in some cases, it is plausible that the alternatives are playing a role only in addition to other features of the context: The cases seem to be about what context we're in, what we assume is the case in the context and how specific the question under discussion is. As we shall see below (Section 5.5), it is possible to get what could be considered quasi-existential readings of generics without the presence of focus sensitive particles or any other obvious features of information structure.

These responses are rather brief, but the reader should still get the sense that none is viable as a solution for Cohen. The data presented above is robust and these band-aids won't work to address the overall problem adequately. Next I will consider what information structure more generally can offer to address the problem. In the last section, I will propose that the indexical approach has the resources (and more) to address the full range of data.

5.4 Information Structure and Existential Generics

In this subsection, I'll consider whether a more general approach, using information structure as a basis (as Cohen does), can explain the full range of data. The account I will expound here is inspired largely by a footnote in Beaver and Clark (2008):

> Note that typically the restrictor set of a generic is not identified with an already salient set, but with the maximal set that satisfies the [local] existential requirement... Note, relatedly, that while it has been argued that presuppositions in the scope of a quantifier cannot be accommodated directly into the restrictor (Beaver, 2001), it is not

> clear that the same arguments apply to generics. That is, in the case of generics, there may be reason to say that presuppositions directly constrain the domain of quantification. (2008, p.61)

The proposal I have in mind to account for the interaction of generics with other focus-sensitive particles is as follows:

(B&C) The different interpretations — quasi-existential and quasi-universal — are due to the sentences at issue answering different questions under discussion. A quasi-existential interpretation can result when the question under discussion is not congruent to the induced $\textsc{Alt}_F$. Otherwise a quasi-universal interpretation results.

There are three ingredients that accompany the (B&C) proposal:

- The domain of *Gen* is restricted by $\bigvee \textsc{Alt}_F$ pretty much as Cohen claims.
- The domain of *Gen* is also restricted by the maximal set that satisfies the existential requirement introduced by the (global, current, possibly implicit) question under discussion.
- The existential generics discussed here are cases of *multiple focus*. (Cf. Krifka (2006).)

This is a rather brief description of the proposal, so it will be useful to go through a concrete examples. Consider again:

(1) Even [mammals]$_F$ lay eggs.

Notice that primary accent is on *mammals* and so this focus associates with *even* and introduces the question under discussion:

(24) What animals lay eggs?

The existential requirement introduced by (24) is:

(25) x is an animal that lays eggs.

(25) restricts the domain of the generic. In addition, there is a second occurrence or default focus on *lay eggs* in (1), so that the domain is also restricted by $\bigvee \textsc{Alt}_{lay\ eggs}$. So that (1) says something like:

(26) You wouldn't expect this as an answer to (24), but in general, if x is a mammal that reproduces and lays eggs, then x lays eggs.

Another concrete example, where there is a quasi-universal reading, goes like this:

(27) Importantly boxers [are smart]$_F$.

In this example, primary accent is on *smart*. This focus associates with *importantly* and introduces the question under discussion:

(28) What property do boxers have?

The existential requirement introduced by (28) is:

(29) x is a boxer with some property.

(29) restricts the domain of the generic.

(30) *In general, if x is a boxer, then x is smart* is a notable answer to (28).

These are examples where the proposal could at least potentially work, but it doesn't work in general.

The proposed (B&C) solution doesn't work for cases where focus isn't congruent with the question under discussion and we still get a quasi-universal interpretation, or it is congruent and we get quasi-existential ones. For example, in the following example, we have focus associating with the generic and the focus is not congruent to the question under discussion, and we don't get a quasi-existential reading:

(31) Importantly, [boxers]$_F$ are smart.

There are of course at least some smart boxers, but (31) is nonetheless intuitively false in a context where the question under discussion is something like *what professions have smart people?* and is not congruent to the question under discussion.

The (B&C) proposal is of course just one way of spelling out the connection between existential generics and information structure. To my lights, the strategy proposed here is the most promising account for deriving the two readings from features of information structure. But as demonstrated above, the account doesn't work. Thus, I conclude that appealing purely to features of information structure lack support. It doesn't seem like information structure can do all the work in deriving the appropriate readings. Though I think information structure has some role to play, I do not think it is not enough on its own.

5.5 Indexical Quantification and Existential Generics

As discussed in Chapter 4, the overall issue and data of generics receiving different interpretations is much more general than the few cases discussed here. Existential generics are just an instance of the more general observation that generics can vary their interpretation depending on numerous features of the utterance in question. On the indexical approach outlined in Chapter 4, we would perhaps even expect that generics receive interpretations like those discussed by von Fintel and Cohen.

Further Issues II: Indefinite Singular Generics and Indexical Quantification

The standard view is that paradigm generics are *descriptive generalisations* of some sort (i.e., they express generalisations about the world) and that they are *quantificational* — akin to sentences like *all normal dogs bark* or *generally dogs bark*. However, many theorists think that there are several distinct generalisations which generics express aside from their paradigm majority-based reading — i.e., that generics are ambiguous and there are other generalisations that generics can express in addition to their standard interpretation. Some examples of the different interpretations are as follows:

(32) Scots wear kilts. (*Distinguishing property interpretation*)

(33) Frenchmen eat horse meat. (*Relative interpretation*)

(34) Peugeots go 120mph. (*Capacity reading*)

(35) Bishops move diagonally. (*Rules-and-regulations* reading)

(36) Bachelors are unmarried males. (*Definitional reading*)

This appendix will address the latter two interpretations. I will begin by discussing the observations found in the literature on these readings, and presenting the different views that have been put forward to account for the observations. I will, then go on to provide my own diagnosis of the data and end with some remarks about how the indexical approach outlined in Chapter 4 might be used to provide an account of the data.

5.6 Indefinite Singular Generics

Recall:

(35) Bishops move diagonally. (*Rules-and-regulations* reading)

(36) Bachelors are unmarried males. (*Definitional reading*)

The availability of these two interpretations is said to be evidenced by formal differences between bare-plural (BP) generics and indefinite singular (IS) generics. In particular, many theorists think that the availability of these interpretations explains why the distribution of IS generics is more limited than BP generics. To illustrate, consider:

(37) a. Waltzs are popular.
b. A waltz is popular. *(?)* or *existential interpretation*

(38) a. Kings are generous.
b. A king is generous. *(?)* or *existential interpretation*

While the (a.) cases receive a generic interpretation, the (b.) cases, according to these theorists, do not. It is important to emphasise that the (b.) cases are not false, it is simply that their generic interpretation is either unacceptable, uninterpretable, infelicitous, unavailable or much less preferred. There are many different proposals which attempt to describe the data (i.e., cases like (37) and (38)), what each of them have in common is that they all either postulate a distinct *rules-and-regulations* reading or a *definitional* reading of the relevant sentences.

More on the Data: One of the main reasons for adopting a distinct *rules-and-regulations* reading or a *definitional* reading of the relevant sentences, is the old-school observation that IS generics are restricted to properties that are somehow *essential* or *necessary* or *inherent* or *analytic* — cf. for example, Lawler (1973) and Burton-Roberts (1977). To see this, contrast:

(39) a. A bachelor is an unmarried man.
b. Bachelors are popular.
c. A bachelor is popular. *(?)* or *existential interpretation*

(40) a. A madrigal is polyphonic.
b. Madrigals are popular.
c. A madrigal is popular. *(?)* or *existential interpretation*

(41) a. A king is a ruler.
b. Kings are generous.
c. A king is generous. *(?)* or *existential interpretation*

(42) a. A gentleman opens doors for ladies.
b. Gentleman wake up early.
c. A gentleman wakes up early. *(?)* or *existential interpretation*

The IS sentences in (39a.)-(42a.) each predicate properties that are in some way *essential* or *necessary* or *inherent* or *analytic* of the subject, and they all receive what seems to be, at least *prima facie*, a generic interpretation, whereas the IS sentences in (39c.)-(42c.) do not predicate such properties and they do not receive a generic interpretation despite the fact that their BP counterparts do.

Early theorists thought that it was this distinction that accounted for the difference in distribution facts. Note, however, that there are many counterinstances to the old-school observation — i.e., there are many examples where non-essential properties yield indefinite generics (cf. Cohen (2001a) and Greenberg (1998)):

(43) A car costs no less than $5000.

(44) A madrigal is a popular song.

(45) A king is a generous ruler.

(46) A waltz is easy to dance to.

(47) A waltz puts people in a good mood.

None of (43)-(47), for example, predicate properties that are essential in any of the relevant senses, but nonetheless each of them receives what seems to be a generic interpretation, at the very least their interpretations are not a straightforward indefinite interpretation (i.e., an existential interpretation as in *a tiger is on my front lawn*). Thus, any account must be compatible with both essential and non-essential properties occurring in IS sentences which do not receive a standard indefinite interpretation. Such an account, however, should still predict that there is at least some tendency towards interpreting IS sentences non-existentially when the sentence in question predicates an essential property, and it should, of course, explain any difference between the distribution of BP and IS generics.

One might take (43)-(47) as evidence that no robust or interesting difference in distribution between IS and BP generics exists, however the robustness

of the distribution facts is meant to be illustrated by several further linguistic observations (cf. Cohen (2001a), Greenberg (2003) and Krifka (2009)) outlined below:

(Questions) It is harder to form felicitous questions with a generic IS than with a generic BP:

(48) a. Is a madrigal polyphonic? *(?)* or *existential interpretation*
b. Are madrigals polyphonic?

(Confidence) With a generic BP, it is possible to use expressions describing the confidence the speaker has in the truth of the sentence, but this is odd with a generic IS:

(49) a. Certainly, madrigals are polyphonic.
b. Certainly, a madrigal is polyphonic. *(?)* or *existential interpretation*

(Modified) With a BP, modified noun phrases get generic readings more easily:

(50) a. Tall, left-handed Norwegians wear thick socks.
b. A tall, left-handed Norwegian wears thick socks.*(?)* or *existential interpretation*

(Conjunction) With IS, conjunctions are difficult:

(51) a. Beavers and otters build dams.
b. A beaver and an otter build dams.*(?)* or *existential interpretation*

Each of (Questions), (Confidence), (Modified) and (Conjunction) present data that is meant to confirm a robust difference between the distribution of IS generics and BP generics.

The data on IS generics leaves theorists with the following two main issues to explain:

1. Why does predicating properties that are essential (in the relevant sense) of an IS subject result in a tendency to interpret the IS sentence in question non-existentially?

2. Why do the IS sentences in question and the corresponding BP sentences have different distributions — why is it possible that IS sentences and BP sentences with seemingly similar structure receive different interpretations in the same context , or in some cases, that the interpretation the BP sentence is available while the corresponding IS sentence is not?

Theorists differ on the extent to which they see the two issues as related as we will see (though, of course, issue 1 will have some bearing on issue 2 no matter what account offers). Theorists also, relatedly, differ to what extent these issues are related to generic phenomena (i.e., whether the distinct readings are "generic" or not).

5.7 Existing Theories of IS Generics

There are three main concrete proposals which attempt to explain these issues. I briefly outline each of these below:

1. **Greenberg (2003):** Greenberg thinks that IS generics express definitional "in-virtue-of" generalisations and that BP generics express both descriptive and definitional generalisations. Greenberg describes the cases as follows:

   > In-virtue-of (definitional) generalisations can only be asserted [with respect to] some relatively specific property associated with the property denoted by the [...] subject, in virtue of which every member of the corresponding set has the predicated property. For example, *A boy doesn't cry* will be true [...] if there is some property we associate with the set of boys: a genetic property, or a social norm property [...] in virtue of which every member of the set of boys will not cry. (2003, p.44)

2. **Cohen (2001):** Cohen thinks that IS generics express that a particular rule is "in effect" and that BP generics are ambiguous between the descriptive and this rules-and-regulations interpretation. He describes his view as follows:

   > The respective distributions of generic ISs and BPs are different because their respective interpretations are different. BPs may denote kinds, hence they may be topics and (after the accommodation of the generic quantifier and type-shifting) may yield characterising generics.
   >
   > ISs, on the other hand, cannot be topics, hence the option of characterising generics is not open to them. However, there is another form of generic, one which denotes a rule or a regulation, and this interpretation of ISs is licensed.

> In many cases, the rules expressed by generic ISs are linguistic rules, i.e. definitions. Whenever an IS sentence may, in terms of its meaning or form, be construed as a definition, the sentence may get a generic reading.
>
> The respective logical forms of generic ISs and (one of the readings of) generic BPs are different, and only the latter express quantification. Construed in this way, it is thus not puzzling, but, in fact, hardly surprising, that the two forms have different distributions and different semantic properties. (2001a, p.207)

3. **Krifka (2009):** Krifka contends that IS generics are (for the most part) definitional second-order predications and that BP generics are ambiguous between descriptive generalisations and this definitional interpretation. In a presentation handout on the topic, he says for instance:

> Descriptive generics and definitional generics have a fundamentally different representation.
>
> Descriptive generics are either based on modal quantification or on probability judgments – not discussed here.
>
> Definitional generics make statements about the meanings of expressions and how they should be used. This makes them second-order predications. They are not quantificational or based on probability judgments.
>
> This explains why bare plurals tend to be used for descriptive generics, and why indefinite singular generics tend to be used for definitional generics. But these are mere tendencies. (2009)

Let's consider the below example to illustrate the three theories:

(52) a. A bachelor is an unmarried man.
b. A bachelor throws wild parties. *(?)* or *existential interpretation*

What Greenberg would say about this example is that (52a.) is acceptable since *unmarried man* is (potentially) definitional of *bachelors*. It says that the generalisation *every bachelor is an unmarried men* is non-accidentally true since there is some property we associate with bachelors (linguistic property in this case) in virtue of

which all bachelors are unmarried men. Whereas (52b.) is unacceptable because ISs only express definitional generalisations, and the relevant property cannot be construed as of the appropriate sort.

Cohen, on the other hand, would say that (52a.) is acceptable since it can be construed as a (potential) rule, this licenses the rules-and-regulations interpretation. It says that the rule represented by *a bachelor is an unmarried man* is in effect. (52a.) is true since there is such a rule in effect. Whereas, (52b.) is unacceptable since it cannot be construed as a rule.

Krifka, by contrast, would say that (52a.) is acceptable since it can be construed as making a statement about the meaning of the term *bachelor*. It says that the term bachelor should be used in such a way that it applies to married men. It is a second-order predication applied to the intension of *bachelor*. (52b.) can be interpreted as a descriptive generic, or as characterising linguistic definition (in so far as they are interpretable). The characterising definitional interpretation says that *bachelor* should be used in such a way that in all normal worlds, if x is a bachelor, then x throws wild parties. In so far as (52b.) is unacceptable, this is merely because it cannot be interpreted in any of these ways.

5.8 Diagnosis of Existing Theories and IS Generics

In this section, I address how I think the two issues outlined above should be resolved. In doing so, I offer my own take on the data and present criticisms of existing theories.

Issue 1: As we saw in the previous section, in addressing issue 1, other theorists posit either a definitional or a rule-and-regulations reading of the relevant IS sentences. I would like to claim two things about such readings:

1. While I agree with Cohen (2001a) and Krifka (2009) that the representation (or logical form) of the definitional reading does not involve *Gen* and hence, is not a generic interpretation, I disagree with what the underlying representation (or logical form) they propose: Definitional readings are about the meanings of expressions, and hence, are metalinguistic — they need a metalinguistic device as part of their representational structure.

2. There is no distinctive or systematic rules-and-regulations reading of the relevant sentences: These sentence are either interpreted as descriptive generalisations which merely accord with some (perceived) rule; or they are uttered with a tacit *According to the rules* operator uttered as part of their

logical form; or they are generalisations which have an appropriate intensional lexical restriction provided by context (cf. Section 4.4.3, Chapter 4). There is no systematic reading of the relevant IS sentences which is a rules-and-regulations reading.

I consider each in turn.

1. Evidence for a Metalinguistic Device — Mixed Quotation: Here I will propose an alternative logical form to account for the definitional readings (or at least some of them). I propose that they are cases of *mixed quotation*, and thus, are expressed by sentences like the below, which are not generics:

(53) A'bachelor' is an unmarried man. (*definitional*)

(53) says something like 'bachelor' denotes an unmarried man. There are many different views about the content of mixed quotations like (53) and I will not chose one of those here (cf. for example Geurts and Maier (2005), Recanati (2001) or Cappelen and Lepore (1997)). I merely wish to point out that such a representation is more plausible to account for the relevant interpretations than the representations proposed by Cohen and Krifka, on the grounds that if the relevant readings do indeed talk about meanings or rules, then a metalinguistic device is needed. It is worth noting that such a reading might explain the apparent prevalence of generics in maternal speech: Some of these may in fact be mixed quoted sentences that talk about the meanings of expressions and how they should be used.

2. No Systematic Rules-and-Regulations Reading: On thing to note from the get-go is that whether or not the relevant sentences receive a generic interpretation or rules-and-regulations reading is a matter of context. Consider:

(36) Bishops move diagonally.

(54) Cars drive 30 mph.

Both (36) and (54), I contend, can express standard generic generalisations, but they can also express a rule-and-regulation, though they needn't (semantically) express a rule-or-regulation. One way to describe what is going on is to say that when describing the movements of different chess pieces in such a way that the description accords with the rules, the speaker makes a true descriptive generalisation, but needn't (semantically) express any rule or that what was said is a rule. Another is to say that there is a tacit *According to the rules* operator in addition to

Gen and the reading is a result of the interaction of this operator with *Gen*. A still further way of describing these cases is to say that the generalisations expressed simply have the appropriate intensional lexical restriction and that the lexical restriction which is part of the meaning of *Gen* is allowed to vary (cf. Chapter 4). I think either of these could be going on to yield the relevant readings and that it is a matter of context which one it in fact is.

To see this, it is useful to think about cases where the rules-and-regulations reading and the standard description reading diverge. Consider a country where cars uniformly drive above the speed limit (30 mph), and suppose the speaker is telling someone about driving conditions in this country.

(55) Cars drive 30 mph, but in fact they drive 60 mph.

(56) Cars drive 60 mph, though the rules say not to exceed 30 mph.

(55) is intuitively false. But if the rule was expressed in (55), then it should be true. The speaker should have said something like (56) which is intuitively true and expresses the descriptive (generic) generalisation.

Contradictory conjunctions make the same point:

(57) Cars drive 30 mph, but cars drive 60 mph.

(57) sounds contradictory or at least bad, but if the first conjunct has the rule-and-regulations reading and the second a standard generic interpretation, then it should be fine.

My goal here is not to deny that there are contexts in which one can get a rules-and-regulations reading, but rather to say that there is nothing distinctively generic about the reading and that the reading is not systematically available like many theorists would like to claim. There is not really anything systematic or semantically significant to say about the connection between such readings and the sentences at issue. To see this, observe that one can, in the right contexts, make non-generic, non-IS sentence of the relevant sort have a rules-and-regulations reading. Imagine we are in a context taking about the rules, then (58) and (59) can sound like rules-and-regulations:

(58) All children cross the street right there.

(59) Every two years Nora changes her passport.

These have something like a tacit *According to the rules* operator. Generic ISs with such an operator will have a rules-and-regulations reading, but it is the operator that contributes the rules-and-regulations reading, not anything distinctively

generic or anything distinctively to do with the IS sentences in question.

Issue 2: As we saw, most theorists stipulate a systematic ambiguity to address the distribution facts. Here I will present my own diagnosis of the distribution data. My hypothesis is that there are many contributing factors which play a role in what causes a difference in distribution between the IS and BP sentences at issue, and that there is nothing very systematic to say about it — at least not anything that will tell us something interesting about generics. There are several things that make me think this:

1. First, if an interpreter must choose to interpret the relevant sentence as an existential, a generic or a mixed-quoted sentence, among other possibilities, then we would expect that an existential reading is preferred for IS subjects as opposed to BP subjects since there are relevant semantic differences, for example of singularity and plurality respectively.

2. Second, the availability data seem to be mere tendencies and we should not adopt a solution which doesn't permit the interpretation of ISs as generic generalisations — descriptive readings are often available where Greenberg and Cohen predict they are impossible. Consider:

   (60) A tiger has stripes.

   (61) A roller coaster is fun.

   (62) A raven is black.

   (60)-(62) obviously have an IS subject and I hear the standard generic descriptive generalisation reading as possible, even preferred. Even better are contexts where we've eliminated the existential and definitional/rules-and-regulations readings.

   (63) A: Tell me something interesting about boys / tigers / roller coasters / ravens.
   B: A tiger has stripes. / A roller coaster is fun. / A raven is black.

3. The choice of interpretation and felicity of explicitly A-quantified sentences vary in distribution depending on whether or not they have an IS subject or BP subject. Consider, for example:

   (64) a. A madrigal, in general, is popular. (*existential interpretation of IS subject*)
   b. Madrigals, in general, are popular. (*generic interpretation of BP subject*)

(65) a. A king is generally generous. (*existential interpretation of IS subject*)
b. Kings are generally generous. (*generic interpretation of BP subject*)

The explicit *in general* in (64a.) is comparatively hard to hear as quantifying over instances of madrigals, rather it appears to be (merely) quantifying over events or situations containing an instance of a madrigal. (64b.) by contrast easily receives a generic interpretation. A similar contrast arises in the case of (65a.) and (65b.) as well.

4. Fourth, the supporting robustness data is made out to be more significant than it actually is: Intuitions here are not so robust. Contrast the data in (Questions), (Confidence), (Modified) and (Conjunction) above with the below:

(66) Is a bachelor an unmarried man? (Question)

(67) Certainly a bachelor is an unmarried man. (Confidence)

(68) a. A good child cleans her room. (Modified)
b. A frozen dessert is best served immediately.

(69) A fish and a crab are best served fresh, not frozen. (Conjunction)[5]

In each of these I hear a generic reading or at least a non-existential reading.

5. Finally, considerations of topicality and information structure as discussed in Cohen (2001a) will play a role. These are non-generic features of logical form.

5.9 IS Generics and Indexical Quantification

As mentioned above, the indexical approach could be used to explain at least some of the relevant interpretations of IS and BP sentences (i.e., those which are said to have a definitional/rules-and-regulations reading, and are not mixed-quotations). The approach outlined in Chapter 4 is extremely flexible, and in particular, if the lexical restrictor is chosen in context to reflect the relevant definitional or rules-and-regulations — that is, intensional — situations, then we have an explanation of the sense that the generalisation expressed is somehow

[5]The conjunction data is obviously the most difficult here. It is worthwhile noting that on the indexical approach advocated in Chapter 4 we might expect this: The context shifts required to get the appropriate readings might make the relevant interpretations difficult.

definitional or rules-and-regulations based. For example, if the lexical restriction is to situations in which the rules of chess hold, then (36) would express a generalisation which sounds like a rule-and-regulation. The fact that the lexical restrictor of *Gen* is allowed to vary with context makes these readings possible, and it allows advocates of the indexical approach to preserve the idea that these are indeed paradigm generic generalisations, and a separate interpretation is not always needed. It also helps explain why the line between genuine IS generics and mixed-quoted sentences with a definitional reading, or an *According to the rules* reading, is so fuzzy (i.e., why it is so hard to say in a given case which one is indeed the correct interpretation given a particular essential property of the relevant sort).

Not only this, but the indexical approach might offer avenues to investigate another related issue involving the IS sentences at issue — notably, *arbitrary reference* in natural language. I outline this avenue below before concluding this Appendix.

Another Interpretation of IS Generics? A still different reading of the IS-sentences in question is one on which the IS is interpreted as arbitrarily referring — that is, the IS is interpreted as *an arbitrary x*. Consider, for example the plausibility of interpreting the generics in (70a.) and (71a.) as in (70b.) and (71b.), respectively:

(70) a. A tiger has stripes.
b. An arbitrary tiger has stripes.

(71) a. A smoker smokes one pack per day.
b. An arbitrary smoker smokes one pack per day.

There are various semantics proposed for arbitrary reference in logical language — cf. for example, Fine (1985), King (1991) or Magidor and Breckenridge (2012). None of these, however, deals with arbitrary reference in natural language. A notable exception is Pelletier and Asher (1997), who discuss the possibility of treating generics more generally along these lines, only to dismiss the proposal. One intriguing thing about the existing literature, however, arises when considering the account of King (1991). According to King, in formal language, *an arbitrary x* has a semantics on which it is treated as a tacit and contextually determined quantifier expression. This results in the intriguing hypothesis that IS generics are perhaps the/a natural language counterpart of arbitrary reference in natural language. And, further that if King's account is on the right track, then treating *Gen* as an indexical quantifier can not only account for the IS data, but

also provide an account of arbitrary reference in natural language. Of course, this is only a bold conjecture and requires further investigation, but it is worth noting and saving as future work.

Even if it turns out that the indexical approach is not an explanation of the IS data or of arbitrary reference, I hope to have demonstrated the versatility of such an approach and provided some fruitful avenues of future research and discussion.

References

Abelson, R., & Kanouse, D. (1966). Subjective acceptance of verbal generalizations. In S. Feldman (Ed.), *Cognitive consistency: Motivational antecedents and behavioral consequents.* New York: Academic Press.

Anderson, C., Lepper, M., & Ross, L. (1995). Perseverance of social theories: The role of explanation in the persistence of discredited information. *Journal of Personality and Social Psychology, 39,* 1037-49.

Asher, N. (2006). Leslie on generics. *Workshop at Rutgers University.*

Asher, N., & Morreau, M. (1995). What some generic sentences mean. In G. Carlson & F. Pelletier (Eds.), *The generic book.* Chicago: University of Chicago Press.

Asher, N., & Pelletier, F. (2012). More truths about generic truth. In A. Mari, C. Beyssade, & F. Del Prete (Eds.), *Genericity.* Oxford: Oxford University Press.

Bach, K. (2003). Quantification, qualification and context: A reply to Stanley and Szabo. *Mind and Language, 15(2-3),* 262-83.

Barwise, J. (1993). Constraints, channels, and the flow of information. *Situation Theory and its Applications, 3*(37).

Beaver, D., & Clark, B. (2008). *Sense and sensitivity: How focus determines meaning.* Oxford: Wiley-Blackwell.

Bhatt, R. (1999). Ability modals and their actuality entailments. In S. B. K. Shahin & E.-S. Kim (Eds.), *WCCFL 17 proceedings.* Stanford: CSLI Publications.

Brem, S., & Rips, L. (2000). Explanation as evidence in informal argument. *Cognitive Science, 24(4),* 573-604.

Buchanan, R. (2010). A puzzle about meaning and communication. *Nous, 44(2),* 340-71.

Buring, D. (2011). Topic and comment. In P. Hogan (Ed.), *The Cambridge encyclopaedia of the language sciences.* Cambridge: Cambridge University Press.

Burton-Roberts, N. (1977). Generic sentences and analyticity. *Studies in Language, 1,* 155-96.

Cappelen, H., & Hawthorne, J. (2009). *Relativism and monadic truth.* Oxford University Press.

Cappelen, H., & Lepore, E. (1997). Varieties of quotation. *Mind, 106(423),* 429-50.

Carlson, G. (n.d.). *Reference to kinds in english.* Doctoral Dissertation, Department of Linguistics, University of Massachusetts Amherst.

Carlson, G. (1989). The semantic composition of English generic sentences. In B. P. G. Chierchia & R. Turner (Eds.), *Properties, types and meanings (vol. 2): Semantic issues.* Dordrecht: Kluwer.

Carlson, G. (2009). Generics and concepts. In J. Pelletier (Ed.), *Kinds, things and stuff: Mass terms and generics.* Oxford: Oxford University Press.

Carston, R. (2000). Explicature and semantics. *UCL Working Papers in Linguistics*, *12*, 1-44.

Carston, R. (2002). *Thoughts and utterances: The pragmatics of explicit communication* (Vol. 33). Blackwell Oxford.

Cavedon, L., & Glasbey, S. (1994). Outline of an information-flow model of generics. *Acta Linguistica Hungarica*, *42*(3/4), 2-6.

Chierchia, G. (1998). Reference to kinds across languages. *Natural Language Semantics*, *6*, 339-405.

Cimpian, A., Brandone, A. C., & Gelman, S. A. (2010). Generic statements require little evidence for acceptance but have powerful implications. *Cognitive Science*, *34*, 1452-82.

Cohen, A. (n.d.). *Think generic!: The meaning and use of generic sentences.* Doctoral Dissertation, Department of Linguistics, Carnegie Mellon University.

Cohen, A. (2001a). On the generic use of indefinite singular. *Journal of Semantics*, *18(3)*, 183-209.

Cohen, A. (2001b). Relative readings of many, often and generics. *Natural Language Semantics*, *9*, 41-67.

Cohen, A. (2003). Existential generics. *Linguistics and Philosophy*, *27(2)*, 137-68.

Cohen, A. (2004). Generics and mental representations. *Linguistics and Philosophy*, *27(5)*, 529-56.

Cohen, A. (2008). No alternative to alternatives. *Journal of Semantics*, *26(1)*, 1-48.

Cohen, A. (2011). Generics as modals. *Genius Genericity Conference, Ecole Normale Superieure.*

Cohen, A. (2012). No quantification without reinterpretation. In A. Mari, C. Beyssade, & F. Del Prete (Eds.), *Genericity.* Oxford: Oxford University Press.

Declerck, R. (1991). The origins of genericity. *Linguistics*, *29*(1), 79–102.

Fara, M. (2005). Dispositions and habituals. *Nous*, *39(1)*, 43-82.

Fine, K. (1985). *Reasoning with arbitrary objects.* Oxford: Basil Blackwell.

Fintel, K. von. (n.d.). *Restrictions on quantifier domains.* Doctoral Dissertation, Department of Linguistics, University of Massachusetts Amherst.

Fintel, K. von. (1997). Bare plurals, bare conditionals and 'only'. *Journal of Semantics*, *14*, 1-56.

Gelman, S. (2009). Generics as a window onto young children's concepts. In F. Pelletier (Ed.), *Kinds, things and stuff: Mass terms and generics.* Oxford: Oxford University Press.

Geurts, B., & Maier, E. (2005). Quotation in context. *Belgian Journal of Linguistics*,

17(1), 109-28.

Greenberg, Y. (1998). Temporally restricted generics. In D. Strolovitch & A. Lawson (Eds.), *Proceedings of the eighth conference on semantics and linguistic theory.*

Greenberg, Y. (2003). *Manifestations of genericity.* London: Routledge.

Greenberg, Y. (2007). Exceptions to generics: Where vagueness, context dependence and modality interact. *Journal of Semantics*, *24*, 131-67.

Haslanger, S. (2011). Ideology, generics and common ground. In C. Witt (Ed.), *Feminist metaphysics: Explorations in the ontology of sex, gender and the self.* Springer.

Hollander, M., Gelman, S., & Star, J. (2002). Children's interpretation of generic noun phrases. *Developmental Psychology*, *36*, 888-94.

Holton, R. (2011). Modeling legal rules. In A. Marmor & S. Soames (Eds.), *Philosophical foundations of language in the law.* Oxford: Oxford University Press.

Ichikawa, J., & Jarvis, B. (2009). Thought experiment intuitions and truth in fiction. *Philosophical Studies*, *142(2)*, 221-46.

King, J. (1991). Instantial terms, anaphora and arbitrary objects. *Philosophical Studies*, *61*, 239-65.

Koslilcki, K. (1999). Genericity and logical form. *Mind & Language*, *14(4)*, 441-67.

Kratzer, A. (2012). *Modals and conditionals: New and revised perspectives.* Oxford: Oxford University Press.

Krifka, M. (2006). Basic notions of information structure. In G. F. C. Fery & M. Krifka (Eds.), *Interdisciplinary studies on information structure.* Potsdam: University of Potsdam.

Krifka, M. (2009). Definitional generics as second-order predications. *Genius Genericity Conference, Ecole Normale Superieure.*

Krifka, M., Pelletier, F., Carlson, G., Meulen, A. ter, Link, G., & Chierchia, G. (1995). Introduction. In G. Carlson & F. Pelletier (Eds.), *The generic book.* Chicago: University of Chicago Press.

Kripke, S. (1977). Speaker⊠s reference and semantic reference. In P. French, T. Uehling, & H. Wettstein (Eds.), *Contemporary perspectives in the philosophy of language.* University of Minnesota Press.

Lawler, J. (1973). *Studies in english generics.* Doctoral Dissertation, Department of Linguistics, University of Michigan.

Leslie, S. (n.d.). *Generics, cognition and comprehension.* Doctoral Dissertation, Department of Philosophy, Princeton University.

Leslie, S. (2007). Generics and the structure of the mind. *Philosophical Perspectives*, *21(1)*, 375-403.

Leslie, S. (2008). Generics: cognition and acquisition. *Philosophical Review*, *117(1)*, 1-47.

Leslie, S. (2009). Generics, semantics and semantic cognition. *Presentation, Department of Philosophy, University of St Andrews*, -, -.

Leslie, S. (2012). Carving up the social world with generics. In *Oxford studies in*

experimental philosophy. Oxford: Oxford University Press.

Leslie, S. (2013). The original sin of cognition: fear, prejudice and generalization. *Journal of Philosophy*.

Leslie, S., & Gelman, S. (2012). Quantified statements are recalled as generics: Evidence from preschool children and adults. *Cognitive Psychology, 64*, 186-214.

Liberman, M. (2009, September). *Language log: Mandatory treatment for generic bare plurals?* Available from `http://languagelog.ldc.upenn.edu/nll/?p=1737`

Liebesman, D. (2011). Simple generics. *Nous, 45(3)*, 409-42.

Lombrozo, T. (2006). The structure and function of explanations. *Trends in Cognitive Science, 10*, 464-70.

Lombrozo, T. (2009). Explanation and categorization: How 'why?' informs 'what?'. *Cognition, 110*(2), 248-53.

Ludlow, P. (1989). Implicit comparison classes. *Linguistics and Philosophy, 12(4)*, 519-33.

Ludlow, P. (2000). Interpreted logical forms, belief attribution and the dynamic lexicon. In K. Jaszczolt (Ed.), *Pragmatics of propositional attitude reports*. Amsterdam: Elsevier Science.

Magidor, O. (2009). Category mistakes are meaningful. *Linguistics and Philosophy, 32(6)*, 553-81.

Magidor, O., & Breckenridge, W. (2012). Arbitrary reference. *Philosophical Studies, 158(3)*, 377-400.

Meulen, A. ter. (1986). Generic information, conditional contexts and constraints. In E. Traugott, A. ter Meulen, & C. Ferguson (Eds.), *On conditionals* (p. 103-22).

Milsark, G. (1974). *Existential sentences in English*. Doctoral Dissertation, Department of Linguistics and Philosophy, MIT.

Nagel, J. (2008). Knowledge ascriptions and the psychological consequences of changing stakes. *Australasian Journal of Philosophy, 86*, 279-94.

Neale, S. (2007). On location. In M. O'Rourke & C. Washington (Eds.), *Situating semantics: Essays on the philosophy of john perry*. Cambridge: MIT Press.

Nickel, B. (2008). Generics and the ways of normality. *Linguistics and Philosophy, 31(6)*, 629-48.

Nickel, B. (2010a). Ceteris paribus laws: genericity and natural kinds. *Philosophers' Imprint, 10(6)*, 1-25.

Nickel, B. (2010b). Generic comparisons. *Journal of Semantics, 27(2)*, 207-42.

Nickel, B. (2010c). Mosquitoes and sharks, habit and capacities. *Draft, Department of Philosophy, Harvard University*.

Nickel, B. (2012a). *Between logic and the world: An integrated theory of generics*. Draft Manuscript, Department of Philosophy, Harvard University.

Nickel, B. (2012b). 'dutchmen are good sailors': Generics and gradability. In A. Mari, C. Beyssade, & F. Del Prete (Eds.), *Genericity*. Oxford: Oxford University Press.

Nunberg, G. (1995). Transfers of meaning. *Journal of Semantics*, *12(2)*, 109-32.
Partee, B. (1989). Binding implicit variables in quantificational contexts. *Proceedings of the Chicago Linguistics Society*, *25*, 342-65.
Partee, B. (1991). Topic, focus and quantification. *Proceedings of the First Conference on Semantics and Linguistic Theory*, 159-87.
Pelletier, F., & Asher, N. (1997). Generics and defaults. In J. van Benthem & A. ter Meulen (Eds.), *Handbook of logic and language*. Cambridge: MIT Press.
Ramney, M., Schank, P., Hoadley, C., & Neff, J. (1994). I know one when I see one: Expert and novice classification of evidence. In *Proceedings of the fifth annual american society for information science*. Silver Spring.
Recanati, F. (2001). Open quotation. *Mind*, *110(439)*, 637-87.
Recanati, F. (2002). Unarticulated constituents. *Linguistics and Philosophy*, *25(3)*, 299-345.
Recanati, F. (2003). Embedded implicatures. *Philosophical Perspectives*, *17(1)*, 299-332.
Recanati, F. (2004). *Literal meaning*. Cambridge: Cambridge University Press.
Recanati, F. (2010). *Truth-conditional pragmatics*. Clarendon Press.
Roberts, C. (1998). Information structure: Towards an integrated formal theory of pragmatics. In J. Yoon & A. Kathol (Eds.), *osuwpl volume 49: Papers in semantics*. The Ohio State University Department of Linguistics.
Roeper, T., & Matthei, E. (1974). On the acquisition of *some* and *all*. In *Sixth child language research forum, stanford university* (pp. 63–74).
Roeper, T., Strauss, U., & Zurer Pearson, B. (2006). The acquisition path of quantifiers: Two kinds of spreading. *Current Issues in First Language Acquisition: University of Massachusetts Occasional Papers*, *34*, 97-129.
Roeper, T., & Villiers, J. de. (1993). The emergence of bound variable structures. In E. Reuland & W. Abraham (Eds.), *Knowledge and language: From orwell's problem to plato's problem*. Dordrecht, Kluwer.
Rooth, M. (n.d.). *Association with focus*. Doctoral Dissertation, Department of Linguistics, University of Massachusetts Amherst.
Ross, L., Lepper, M., & Hubbard, M. (1975). Perseverance in self-perception and social perception: Biased attributional processes in debriefing paradigm. *Journal of Personality and Social Psychology*, *32*, 880-92.
Rothbart, M., Fulero, S., Jenson, C., Howard, J., & Birrell, P. (1978). From individual to group impressions: Availability heuristics in stereotype formation. *Journal of Experimental Social Psychology*, *14*, 237-55.
Rothschild, D., & Segal, G. (2009). Indexical predicates. *Mind & language*, *24*(4), 467–493.
Schaffer, J., & Szabo, Z. (n.d.). Epistemic comparativism: A contextualist semantics for knowledge ascriptions. *Philosophical Studies*.
Schubert, L., & Pelletier, J. (1987). Problems in representing the logical forms of generics, bare plurals and mass terms. In E. Lepore (Ed.), *New directions in semantics*. Academic Press.
Sperber, D., & Wilson, D. (1995). *Relevance theory: Communication and cognition*.

Oxford: Blackwells.

Stanley, J. (2002a). Making it articulated. *Mind & Language*, *17*(1-2), 149–168.

Stanley, J. (2002b). Nominal restriction. In G. Preyer & G. Peter (Eds.), *Logical form and language*. Oxford: Oxford University Press.

Stanley, J. (2005). Semantics in context. In G. Preyer (Ed.), *Contextualism in philosophy: Knowledge, meaning and truth*. Oxford: Oxford University Press.

Stanley, J., & Szabo, Z. (2000). On quantifier domain restriction. *Mind & Language*, *15*, 219-61.

Tomasello, M. (2001). Perceiving intentions and learning words in the second year of life. *Language Acquisition and Conceptual Development*, 132–158.

Tomasello, M. (2006). The social bases of language acquisition. *Social Development*, *1*(1), 67–87.

Tversky, A., & Kahneman, D. (1983). Extension versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, *90 (4)*, 293-315.

Varynen, P. (2004). Particularism and default reasons. *Ethical Theory and Moral Practice*, *7 (1)*, 53-79.

Wilson, D., & Sperber, D. (2002). Truthfulness and relevance. *Mind*, *111*(443), 583–632.