FROM KINDS TO PLURAL REFERENCE

ESSAYS ON THE MEANING OF GENERICS

TOM RALSTON

A thesis submitted for the degree of Doctor of Philosophy

at the University of Oxford

May 2025

Tom Ralston, *From Kinds to Plural Reference: Essays on the Meaning of Generics*

Supervisors

Professor Ian Rumfitt

Professor Paul Elbourne

Examiners

Professor Ofra Magidor

Professor Daniel Rothschild

This thesis is typeset using André Miede and Ivo Pletikosić's `classicthesis` package.

Memories and thoughts age, just as people do.
But certain thoughts can never age, and certain memories can never fade.

— Haruki Murakami, *The Wind-Up Bird Chronicle*, *trans.* Jay Rubin

To Pat.

ACKNOWLEDGMENTS

I am very grateful to the Hanfling Fund for their financial support throughout my DPhil, as well as the additional financial support provided by the Faculty of Philosophy and Kellogg College, which enabled me to attend and present at a number of conferences and workshops, and my late grandparents, Gavin and Violet Ralston.

Most of all, I would like to thank my two research supervisors, Ian Rumfitt and Paul Elbourne. Ian was extremely patient as I meandered through the first year of my DPhil and was supportive even as I deviated from the subjects that I had originally joined Oxford to research. I owe a great intellectual debt to Ian, who was the person to impress upon me the importance and irreducibility of plurality. This has turned out to be the central theme of this thesis and I believe it to be a deeply important insight that has been missing from much of the literature on this subject. I would also like to thank Ian for his regular reminders that I am a philosopher, not a linguist.

Paul has always been extremely generous with his time and commentary on my work. His role in this project has been in helping me fine-tune the details and making sure that I have always understood the consequences of my assumptions. Paul's expertise in formal semantics has been invaluable in ensuring that the finished product meets the required standard of technical rigour and has a chance of convincing not only philosophers but linguists of its veracity (although of course I alone bear responsibility for any errors).

Versions of work that appears in this thesis have been presented at the British and Irish Postgraduate Philosophy Association Annual Conference at University College Dublin in 2022, the European Congress of Analytic Philosophy at the University of Vienna, the Joint Session of the Mind Association and Aristotelian Society at Birkbeck University, the Arché Graduate Conference at the University of St Andrews, the Summer School in Social Ontology at the University of Neuchâtel and the Southwest

Graduate Philosophy Conference at Arizona State University in 2023, the Amsterdam Colloquium at the University of Amsterdam in 2024 and several iterations of the Oxford DPhil Seminar. I am grateful to those members of the audiences who engaged with my ideas, especially Friederike Moltmann and Manfred Krifka whose opinions on this topic I value highly. Also those who have been generous enough to read my work in draft, especially James Kirkpatrick, who helped steer me in the right direction while the ideas in this project were still embryonic, Flavia Naehrlich, who provided some very welcome fresh eyes late on, Tim Williamson and Julian Ratcliffe.

I would also like to mention those who introduced me to the philosophy of language, Anders Schoubye and Brian Rabern at the University of Edinburgh and Mat Simpson and Ethan Nowak at UCL. I took Anders's *Formal Semantics for Philosophers* class on a whim in 2016; the wonder that so much of natural language can be explained in terms of function application and a small set of composition rules has yet to leave me. Brian and Mat were both excellent supervisors for my dissertations and I owe them my gratitude for showing me the delights of a research project. Mat has continued to be a source of support and advice in the years since then, for which I will always be grateful.

Nick Jones and James Studd have been excellent informal mentors since I joined Oxford. I have James to thank for helping me navigate the complexities of Oxford's graduate teaching establishment and Nick for all-round moral support. I have been lucky enough to join Nick at St John's for the final year of my degree and it has been a pleasure to be a part of that community. Thanks also to the staff at the Bodleian and Philosophy and Theology libraries for their assistance with my many queries and requests and to the administrative staff in the faculty for their support over the last four years.

Finally, I would like to thank those friends and family who have supported me throughout the last four years, especially after my (foolish?) decision to leave the civil service to pursue another degree in philosophy. First, Gavin, Nicola, Bea, Alice, Joe and (if only for a short time) Miles. There are too many friends to list, but special

thanks go to Alex James for convincing (peer pressuring?) me to accept the offer from Oxford and to Matt Brady for his insights and war stories about life as a PhD student.

Last, but of course not least, I would like to thank Pat, to whom this thesis is dedicated, for putting up with years of me thinking very hard about these strange things called generics, for moving to Oxford (and back) with me and for always being there for me.

TABLE OF CONTENTS

1 GENERICS AND GENERICITY: AN OPINIONATED INTRODUCTION 1

2 COLLECTIVE AND DISTRIBUTIVE READINGS OF GENERICS 47

3 THE TEMPORAL PLURAL REFERENCE THEORY OF GENERICS 85

4 HOW TO MAKE A GENERIC TRUE 169

5 RECONCEPTUALISING THE PSYCHOLOGICAL THEORY OF GENERICS 209

6 CONCLUSION 235

BIBLIOGRAPHY 239

DETAILED TABLE OF CONTENTS

1 GENERICS AND GENERICITY: AN OPINIONATED INTRODUCTION 1

1.1 Introduction 1

1.2 The linguistic forms of generics 1

1.2.1 A diagnostic for existentially interpreted bare plurals 3

1.2.2 Two kinds of generic sentence 4

1.2.3 Two features of the meaning of characterising sentences 5

1.3 Predicational theories 7

1.3.1 The simple predicational theory 8

1.3.1.1 The property objection 9

1.3.1.2 The Excluded Middle objection 9

1.3.2 The sophisticated predicational theory 10

1.3.3 The reflexivity objection to kind-reference 12

1.3.4 Plural predicational theories 14

1.3.4.1 The semantics of plural predication 15

1.3.4.2 Collective and distributive readings of generics 16

1.4 Relational theories 18

1.4.0.1 Structural ambiguity 19

1.4.1 General problems for relational theories 20

1.4.1.1 Unpronounced in all languages 20

1.4.1.2 Lack of uniformity 20

1.4.2 Probability theories 21

1.4.3 Normality theories 23

1.4.3.1 Normal worlds 24

1.4.3.2 Normal objects 27

1.4.3.3 Ways of being normal 29

1.4.4 Referential relational theories 31

1.4.5 Psychological theories 32
1.4.6 Indexical theories 35
1.4.6.1 Some possible constraints 35
1.4.7 Underspecification theories 37
1.5 The nature of genericity 38
1.5.1 Unboundedness? 39
1.5.2 Stability 41
1.5.2.1 Habituals 42
1.6 Roadmap 43
2 COLLECTIVE AND DISTRIBUTIVE READINGS OF GENERICS 47
2.1 Introduction 47
2.2 The data and existing theories 50
2.2.1 The data 50
2.2.2 The binder argument 51
2.2.3 A predicate modifier 52
2.2.4 A scope phenomenon? 53
2.2.4.1 A missing binder 53
2.2.4.2 Reflexive collective readings 54
2.2.4.3 Genuine kind-predication 57
2.2.5 D-relational theories 58
2.2.6 Summing up 62
2.3 Regimenting plurals 63
2.3.1 A plural metalanguage 63
2.3.2 Group noun phrases 65
2.4 Two types of distributivity 69
2.4.1 An operational definition of distributivity 69
2.4.2 Distributivity and groups 70
2.4.3 The sources of distributivity 71
2.5 The pluralist analysis 74
2.5.1 Predictions of a singularist theory 74

2.5.2 Deriving both readings on a pluralist theory 75
2.5.3 Covert binding 78
2.5.4 Co-predications 79
2.6 Rigid designators 80
2.7 Conclusion 82
3 THE TEMPORAL PLURAL REFERENCE THEORY OF GENERICS 85
3.1 Introduction 85
3.2 The landscape 86
3.2.1 The target 86
3.2.2 The orthodoxy 88
3.2.3 The heresy 88
3.2.4 The synthesis 89
3.2.5 The consensus 90
3.2.6 Co-predications 90
3.2.6.1 Co-predications: evidence in favour of uniform kind-predication? 92
3.3 Motivations for a plural reference theory of generics 93
3.3.1 Cross-linguistic data 94
3.3.2 English definite plural generics 95
3.3.3 Instance-distribution predicates 96
3.3.4 Compatibility with collective predicates 98
3.3.5 Distributive and collective ambiguities 98
3.3.6 Cumulative readings 99
3.3.7 The non-maximality/homogeneity paradigm 100
3.3.7.1 Non-maximality 100
3.3.7.2 Homogeneity 101
3.3.7.3 Homogeneity and non-maximality removal 102
3.3.8 The ontology of plural reference 103
3.3.9 Summary of the motivations for a plural reference theory 104
3.4 Existing plural reference theories 105

3.4.1 Chierchia (1998) 105
3.4.1.1 Bare plurals which do not correspond to kinds 106
3.4.1.2 The semantic type of kind-level predicates 108
3.4.2 Koslicki (1999) 109
3.4.3 Moltmann (2013) 111
3.4.4 Diagnosis 112
3.5 The null determiner 112
3.5.1 Formal preliminaries 113
3.5.2 Plural definite descriptions 114
3.5.3 Differences between definite plurals and bare plurals 117
3.5.3.1 Ruling out bare singulars 117
3.5.3.2 Deictic and anaphoric uses of definite descriptions 118
3.5.3.3 Domain selection 119
3.5.3.4 The Inclusivity Principle 120
3.5.3.5 Why are definite plural generics not felicitous? 122
3.5.3.6 A compositional example 123
3.5.4 Temporary generalisations 124
3.5.5 The functional reading of bare plurals 125
3.5.6 Co-predications 126
3.5.7 'All' and 'most' 128
3.5.8 Interim Summary 131
3.6 Ambiguity and information structure 132
3.6.1 Two notions from information structure 132
3.6.2 Mapping material to the restrictor and matrix 133
3.6.3 An alternative analysis 134
3.6.4 Existential bare plurals 136
3.6.4.1 Episodic predicates 136
3.6.4.2 Individual-level predicates 137
3.7 Non-maximality and homogeneity 139
3.7.1 Interest-relativity 139

3.7.1.1 Non-maximality for plural definite descriptions 139
3.7.1.2 Non-maximality for bare plural generics 140
3.7.1.3 Non-maximal noun phrases? 142
3.7.2 The underspecification of plural predication 143
3.7.3 Homogeneity 145
3.7.4 Non-maximality 148
3.7.5 Applying Križ and Spector (2021) to bare plural generics 149
3.7.5.1 Non-maximality and alternative-sensitivity 150
3.7.5.2 Underdeterminacy in the candidates 152
3.7.6 Homogeneity and failures of Generic Excluded Middle 153
3.8 Modality: a missing element? 154
3.8.1 Habitual predicates 155
3.8.2 Three-legged lions 157
3.9 Kind-level predicates in the temporal plural reference theory 158
3.9.1 The *gather-numerous* opposition 159
3.9.2 Kind-level predicates are N-type collective predicates 161
3.9.3 Explaining homogeneity failure 163
3.9.4 Hard cases 164
3.9.4.1 'Be extinct' 165
3.9.4.2 'Be invented' 166
3.10 Conclusion 167
4 HOW TO MAKE A GENERIC TRUE 169
4.1 Introduction 169
4.2 Preliminaries 170
4.2.1 Nomicity 170
4.2.2 The meaning of generics 171
4.2.2.1 Non-maximality 173
4.2.2.2 Temporary generalisations 173
4.2.2.3 Temporally restricted generics 174
4.2.3 Existing theories of generics 176

4.3 The temporal plural reference theory of bare plural generics 178
4.3.1 Plural reference and temporal domains 180
4.3.1.1 Temporal indices 181
4.3.1.2 An unpronounced determiner 181
4.3.1.3 Ruling out temporary generalisations 184
4.4 How to make a generic true 184
4.4.1 Generic and instance-based explanations 184
4.4.2 Intrinsic truthmaker semantics 187
4.4.2.1 The nomicity of generics 189
4.4.2.2 The underspecification of plural predication 190
4.4.3 Facts about categories 191
4.4.3.1 Determinant facts 192
4.4.3.2 Consequent facts 192
4.4.3.3 Relations between universals? 193
4.4.4 Temporally restricted generics 194
4.5 Two objections 195
4.5.1 What if determinism is true? 195
4.5.2 Ambiguity or underspecification? 197
4.5.3 Interim summary 199
4.6 Generics in psychology 200
4.6.1 Generics and essentialism 200
4.6.1.1 Objections 201
4.6.2 Inferences about explanations 202
4.6.2.1 The Generic Explanation Inference 203
4.6.3 The inherence heuristic 204
4.6.4 Other forms of generalisations 207
4.7 Conclusion 208
5 RECONCEPTUALISING THE PSYCHOLOGICAL THEORY OF GENERICS 209
5.1 Introduction 209
5.1.1 The Psychological Theory of Generics 210

5.1.2 Theory confusion 213
5.2 Semantics by opinion poll 214
5.2.1 Collecting truth value intuitions 215
5.2.1.1 Case study: Prasada et al. (2013) 217
5.2.1.2 Case study: Cimpian et al. (2010) 218
5.2.2 Top down semantics 219
5.2.3 Heuristics and truth value judgments 221
5.2.4 Heuristics or meaning? 224
5.3 Evaluation Heuristics 226
5.3.1 Why heuristics? 226
5.3.2 What is preserved? 228
5.3.2.1 Content-sensitivity 229
5.3.2.2 Inferential asymmetry 230
5.3.2.3 Stereotyping and prejudice 231
5.4 Conclusion 232
6 CONCLUSION 235

BIBLIOGRAPHY 239

ABSTRACT

In this thesis, I present a range of arguments in support of the hypothesis that generic bare plural noun phrases in English – those which appear in sentences such as 'Ravens are black' or 'Dodos are extinct' – should be analysed as covert definite descriptions.

In Chapter 1, I introduce the topic of generics and provide a critical summary of the extant semantic theories. I propose a novel way of characterising the phenomenon of genericity as the belief that a sentence is true in virtue of a state of affairs which is stable over time.

In Chapter 2, I present novel data concerning the availability of collective and distributive readings of generics. These ambiguities are paralleled by the availability of collective and distributive readings for non-generic plurals. I show that extant theories cannot explain the availability of both readings and argue that this evidence weighs heavily in favour of a plural reference theory of generics.

In Chapter 3, I present further linguistic parallels between generics and plural definite descriptions and introduce a novel compositional semantic theory which analyses generic bare plurals as covert definite descriptions. I give a lexical entry for a null determiner and show how this predicts both the differences and similarities in the meanings of generic bare plurals and plural definite descriptions.

Chapter 4 deals with the law-likeness of generics, which seems to present a challenge for my plural reference theory. I present a puzzle concerning the law-like interpretation of generics and show that generics with a temporal restriction wholly located in the past elicit diminished intuitions of law-likeness. I resolve this puzzle by proposing a novel truthmaker semantics for plural predication which explains the law-likeness of generics in terms of the different ways that they can be made true.

In Chapter 5, I address Sarah-Jane Leslie's arguments that generics express a default mode of generalisation. I argue that Leslie's approach does not represent a genuine

alternative to semantic theories of generics, since it ignores the non-semantic factors that affect speakers' judgments of the truth values of generics and, as a result, overfits the data. I present a reconceptualised version of Leslie's theory as a theory of how heuristics affect our judgements of the truth values of generics.

Word Count: 55289

Thesis Supervisor: Professor Ian Rumfitt, All Souls College, Oxford

Thesis Supervisor: Professor Paul Elbourne, Magdalen College, Oxford

1

GENERICS AND GENERICITY: AN OPINIONATED INTRODUCTION

1.1 INTRODUCTION

The practice of generalisation is essential to the way that humans think. Many, if not most, of the objects we encounter on a day-to-day basis – on the street, at the supermarket, at the zoo – are novel to us, so if our knowledge about the world consisted only in knowledge of particular objects, we would have to spend much of our time acquainting ourselves with each new object as we encountered it.

By identifying what category an object falls under, we draw many conclusions about that object. For example, when I encounter an apple in my fruit bowl, I believe that it is not only tasty, but a healthy snack, and so, without further ado, bite into it. I do this because I believe that *apples are tasty* and that *apples are healthy snacks*. I may, of course, be wrong about this particular apple and discover, upon biting into it, that it has a dull floury texture. But this particular apple's being less than tasty does not seem like good reason to revise my belief that apples are tasty.

These sentences are *generics* or, more precisely, *bare plural generics*, and are the subject of this thesis. What these sentences mean is one of the big unanswered questions in formal semantics, with just about every conceivable theory having been defended at some point since Greg Carlson's dissertation (1977) launched the field in earnest.

1.2 THE LINGUISTIC FORMS OF GENERICS

Generics come in many linguistic forms:

(1) a) Apples are tasty. *bare plural*

b) An apple is a healthy snack. *indefinite singular*

c) The apple is a fruit. *definite singular*

Each of these sentences has reading which says something about apples *in general*, rather than about any particular apple or apples. They do so without the help of any overt word or expression, such as a quantifier or adverb of quantification, as in the following:

(2) a) Most apples are tasty.

b) An apple is often a healthy snack.

Less well acknowledged is that generics may be expressed using definite plurals in English, albeit in more limited contexts:

(3) a) The poor are unfortunate.

b) The English are keen sailors and enjoy competitive sailing.[1]

c) The women in this family work hard.

Although there is no precise definition of what makes a sentence a generic, I will rely on the intuitive idea that a sentence has a generic meaning when it is about a class of objects in general and does not contain any overt quantifier or adverb or quantification which contributes to this meaning. My focus in this thesis is on generics which take a bare plural form. This is for two reasons:

(i) Bare plural generics are the canonical form of generics in English and are felicitous in the widest range of contexts.

(ii) Whatever different linguistic forms which can express genericity have in common, they are different constructions and will ultimately require distinct compositional semantic theories. As Collins (2015) has argued, the linguistic phenomena are so diverse that it is not reasonable to expect that these sentences having anything in common, at least as far as the compositional semantics go.[2] It makes sense,

1 https://en.wikipedia.org/wiki/England#Sport

2 See Greenberg (2003, 2007) on the differences between indefinite singular and bare plural generics and Krifka (1987) for discussion of definite singular generics.

methodologically speaking, to develop separate compositional semantic theories for each construction; with luck, this approach will explain both the similarities and the differences between them.

1.2.1 *A diagnostic for existentially interpreted bare plurals*

Not all bare plurals are interpreted generically. Instead, many bare plurals are interpreted existentially, particularly in *episodic sentences*, which say something about particular objects or situations.

(4) a) Ravens are nearby.

b) Beavers are building dams.

c) Cryptocurrencies are sold here.

d) Classical logicians are working.

In each case, the sentence is about some particular object(s) or situation. I will employ the following test to identify existentially interpreted bare plural noun phrases: only existential bare plurals are upward-entailing (Krifka et al., 1995).

(5) a) Ravens are black. $\not\Rightarrow$ Birds are black.

b) Classical logicians are intelligent. $\not\Rightarrow$ Logicians are intelligent.

(6) a) Ravens are nearby. $\Rightarrow$ Birds are nearby.

b) Classical logicians are working. $\Rightarrow$ Logicians are working.

This test can serve to disambiguate sentences with bare plurals which may be interpreted existentially or generically:

(7) Champagne socialists frequent this bar.

a) There are champagne socialists who frequent this bar. *existential*

b) In general, champagne socialists frequent this bar. *generic*

On its existential but not its generic interpretation, (7) entails:

(8) Socialists frequent this bar.

1.2.2 *Two kinds of generic sentence*

Some generics say something about members of a class, while others say something about the class as a class. The former are known as *characterising sentences*, the latter as *kind-predication sentences*.[3]

(9) CHARACTERISING SENTENCES

a) Philosophers are cool.

b) Ravens are black.

c) Conservatives have a favourite politician.

d) Philosophy students consider themselves clever.

e) Americans overestimate social mobility.

f) 7 inch records play at 45 rpm.

g) Cows go *moo*.

h) Dolphins use toxic puffer fish to get high.

i) Black lives matter.

(10) KIND-PREDICATION SENTENCES

a) Dodos are extinct.

b) Chickens are numerous.

c) Classical logicians are abundant.

d) YouTube influencers are ten a penny.

e) Elite Scottish football teams from outside Glasgow are rare.

f) Chillis were introduced into India by the Portuguese.

g) Black rats reached Australia by ship.

3 As will become clear in this thesis, I believe that this latter label is a misnomer, since these sentences do not involve predication of kinds. I keep with the standard nomenclature, but readers should be aware that I considers these to be *labels*, not descriptions.

I will argue in this thesis that the differences between characterising and kind-predication sentences correspond to the differences between distributive and a particular type of collective predication in non-generic plural sentences. Very roughly, distributive predicates are true of an argument iff they are true of each individual among their argument, while collective predicates are true of their argument as a whole (see Winter & Scha, 2015).[4]

(11) DISTRIBUTIVE PREDICATES

a) Those students laughed. ⇒ Each of those students laughed.

b) The attendees clapped. ⇒ Each of the attendees clapped.

(12) COLLECTIVE PREDICATES

a) These soldiers are few in number. ⇏ Each of these soldiers is few in number.

b) The Americans reached the moon first. ⇏ Each of the Americans reached the moon first.

The idea is that bare plurals are referential in both characterising and kind-predication sentences and that the differences between them stem from the lexical semantics of the predicates.

1.2.3 *Two features of the meaning of characterising sentences*

Before I discuss the space of semantic theories for generics I will describe two distinctive features of the meaning of characterising sentences:

NON-MAXIMALITY Characterising sentences exhibit a variable tolerance of exceptions. The number of exceptions that a sentence tolerates can vary from case to case. For example, (at least in certain contexts) 'Prime numbers are odd' is false because there is just one exception; 'Dogs bark' is true despite there being many

4 The caveat 'very roughly' is required because, as will become clear, distributive predicates can be true of a plural argument even if they are not true of each thing among their argument. See Champollion (2020: 6-8) for a discussion of this phenomenon, known as *non-maximality* as it relates to the definition of distributivity.

dogs who never bark; 'Sea turtles are long-lived' is true despite the majority of sea-turtles dying very young.

NOMICITY Characterising sentences are often understood to express nomic generalisations, that is, they do not only say something about how the world is, but also how it would be even if things were different in many respects. This means that sentences which describe a state of affairs which holds only accidentally or temporarily are considered 'bad' generics. For example, 'American Presidents are male' is not true even though every actual American President has been male.

In contrast, kind-predication sentences exhibit neither of these features; (10a) depends only on the history of actual dodos and it makes no sense to ask whether some particular dodo could be an exception. (10f) also depends only on the actual history of chillis and, while it is certainly true that individual chillis could be seen as 'exceptions' to that sentence (by virtue of not having been introduced into India), that would be to miss to collective nature of the sentence, which does not attribute any property to individual chillis. No individual chilli can rightly be said to be an exception to this sentence, understood collectively, just as it would be a mistake to say that Joe Biden is an exception to (12b) by virtue of not having reached the moon.

I will now provide an overview of the various semantic theories that have been proposed for generics and, as I go, note various points of interest for the meaning of generics. The discussion will, for reasons of space, be compressed, but I will offer a reasonably comprehensive overview of the existing space of theories, as well as some of the problems that each theory faces.[5]

5 For other overviews of the linguistic meaning of generics, see the comprehensive Krifka et al. (1995) and Mari et al. (2013); for more general discussions of generics, see Sterken (2017), Nickel (2017), Leslie and Lerner (2022) and Cohen (2022).

1.3 PREDICATIONAL THEORIES

According to what I will call *predicational* theories, all generics have the same logical form as simple subject-predicate sentences like (13a).

(13) a) Willard is cool.

b) cool(w)

(14) a) Philosophers are cool.

b) cool(p)

The meaning of the bare plural 'philosophers' in (14a) must therefore be something which can saturate the argument place of the predicate, 'is cool'. The standard answer among predicational theorists is that the bare plural is referential and refers to a kind. I will call any theory which assigns generically interpreted bare plurals a referential semantics (i.e. they refer to an object or some objects), a *referential* theory. All predicational theories are referential, but, as we will see in §1.4.4, not all referential theories are predicational.

Predicational theories assign the same logical forms to both kind-predications and characterising sentences. This means that they are well-placed to explain the following pattern of inference:

(15) a) Rhinos are rare.

b) rare(r)

(16) a) Rhinos are large.

b) large(r)

(17) a) Therefore, rhinos are rare and large.

b) rare(r) $\wedge$ large(r)

The conclusion (17a) is a mixed sentence or *co-predication*, with both a kind-level predicate and an individual-level predicate applying to a single argument. As Koslicki

(1999) and Liebesman and Magidor (2025) argue, being able to validate this very simple pattern of inference is a strength of referential theories, since any theory which assigns (15a) and (16a) radically different logical forms will need some more sophisticated explanation of the validity of the inference to (17a). But for predicational theories, the inference goes through as a simple instance of conjunction introduction.

1.3.1 *The simple predicational theory*

On the very simplest versions of the predicational theory, the logical form of (16a) really is the whole story, as far as the linguistic semantics of generics go. On Liebesman's simple theory, the truth-conditions of a characterising sentence are satisfied iff the kind referred to by the bare plural noun phrase has the property expressed by the predicate (i.e. if the kind RHINO has the property of being large) (Liebesman, 2011; McConnell-Ginet, 2012; Liebesman & Magidor, 2025). The distinctive features of generics, in particular their non-maximality and nomicity, are explained in terms of what it takes for a kind to have a property. Liebesman observes that for many non-generic sentences, an object can hold a property in virtue of the properties of only some of its parts (Liebesman, 2011: 419). For example:

(18) a) The chair is wooden.

b) The kettle is red.

c) The table is touching the wall.

In each case, it is intuitive to think that the object has a property, even though many of its parts do not. For example, we say that a table is touching a wall provided that some part of the table is touching the wall. Things are like this with kinds, Liebesman argues. A kind, PHILOSOPHER, can have a property, coolness, in virtue of some of its members having the property: the kind has "inherited" the property from its members. Importantly, a kind cannot be reduced to its actual members, so it is possible for a kind to possess a property that none of its actual members do. So, for example, if a philosophy-hating maniac assassinated all the world's cool philosophers, (14a)

could still be true, since its truth does not depend only on what the world's actual philosophers are like.

1.3.1.1 *The property objection*

One objection to the simple predicational theory is that it is not clear what it takes for a kind to have a property of individuals, like the property of being large (Teichman, 2023: 1618). Although we might not be able specify necessary and sufficient conditions for an ordinary object's having that property, we intuitively know it when we see it. But for a kind to have the property of being large cannot be, on this proposal, to have many members, since both 'Rhinos are large' and 'Rhinos are rare' are true. So the property must apply to a kind in some different, more mysterious way.

Things are even worse when we consider examples of properties which intuitively cannot be held by abstract objects:

(19) Rhinos are leathery/heavy/ferment food in their hindgut.

Liebesman and Magidor's simple response is that these worries are misplaced (Liebesman & Magidor, 2025: 268-269); since they endorse a more general view they call "property versatility", there is nothing mysterious about a kind having a property usually ascribed to individuals, or an abstract object having a property usually attributed to concrete objects. What it takes for a kind to have the property of being leathery is quite different from what it takes for an individual to have it, but the relevant relation sits within the realm of metaphysics, not linguistic semantics. This response is coherent, yet it may not satisfy Liebesman and Magidor's opponents, since it departs from our intuitions about what it takes for an object to have a property.

1.3.1.2 *The Excluded Middle objection*

Kirkpatrick (2025) makes a more robust objection to the simple predicational theory. He argues that a semantic theory for generics should fail to validate the following principle:

GENERIC EXCLUDED MIDDLE (GEM): For any bare plural characterising sentence of the schematic form ⌜Fs are G⌝, the sentence ⌜Either Fs are G or Fs are not G⌝ is true.

He motivates this desideratum using the following examples where GEM fails:

(20) a) Books are paperbacks or books are not paperbacks.
b) Fair coins land heads or fair coins do not land heads.
c) Lions are male or lions are not male.

Intuitively, none of these disjunctions is true, so they are counterexamples to GEM. The problem for the simple predicational theory is that, assuming a classical metalanguage, the logical form assigned to generics validates GEM:

(21) a) $\text{paperback}(b) \vee \neg\text{paperback}(b)$
b) $\text{land.heads}(c) \vee \neg\text{land.heads}(c)$
c) $\text{male}(l) \vee \neg\text{male}(l)$

Notice that this argument does not rely on the Law of Excluded Middle holding for natural language in general: it says that the metalanguage translations of (20a) are true, but the English sentences are false. One might, of course, choose to reject the Law of Excluded Middle for the metalanguage, but, as Kirkpatrick argues, that seems like a radical response to maintain the simple predicational theory.

1.3.2 *The sophisticated predicational theory*

Sophisticated predicational theories retain the subject-predicate logical form for characterising sentences but add a unary predicate modifier which turns predicates of individuals into predicates of kinds (Carlson, 1977; Teichman, 2023). Using PM for the modifier, (14a) then becomes:

(22) $PM([\lambda x_e.\text{cool}(x)])(p)$

The precise function that PM represents is not of primary importance here (Carlson (1977: 268) suggests that it might not be analysable further; for a more detailed analysis, see Teichman (2023: 1643)). What matters, first, is that the question of property versatility is sidestepped, since kinds can only hold *modified* properties of individuals. Second, Kirkpatrick's Excluded Middle objection is avoided, since PM can scope over negation, so that the metalanguage translation of the counterexample to GEM is not an instance of the Law of Excluded Middle, hence both disjuncts may be false.

(23) $PM([\lambda x_e.\text{paperback}(x)])(b) \vee PM([\lambda x_e.\neg\text{paperback}(x)])(b)$

Third, the introduction of a modifier allows the sophisticated predicational theory to resolve another problem for the simple predicational theory. This objection to the simple theory is that it generates the wrong truth-conditions for generics with bound pronouns (Carlson, 1977: 262-265; Leslie, 2015: 34-35):

(24) a) Cats clean themselves.

b) Dogs obey their masters.

Translating the predicate 'clean themselves' as $[\lambda x.\text{clean}(x)(x)]$, according to the simple predicational theory, the only possible reading should be:

(25) a) $\text{clean}(c)(c)$

b) The kind CAT cleans the kind CAT.

But this fails to capture the meaning where each cat cleans itself. Even granting that cleaning oneself is a versatile property, (25b) is compatible with no cat ever cleaning itself, provided that cats typically clean other cats, so it fails to capture the individual-directed reading of (24a). For the sophisticated theory, however, things are straightforward:

(26) $PM([\lambda x.\text{clean}(x)(x)])(c)$

Now, regardless of the precise meaning of PM, the property attributed to the kind CAT is that of having (enough of the right sort of) instances which clean themselves. This

avoids the problem with the truth-conditions assigned by the simple theory, since this is not true in a scenario in which cats typically clean other cats but never themselves.[6]

1.3.3 *The reflexivity objection to kind-reference*

There is another, more serious problem which generalises to all kind-reference theories, in the threat of paradox caused by a 'Russell-kind' (Kleinschmidt & Ross, 2013: 141-142; Liebesman & Sterken, 2021: 5). The reasoning is as follows:

(i) For a kind-reference theory to be a plausible theory of the meaning of generics, all generic bare plurals in true generics must refer to kinds (whether or not they are 'well-established' or natural kinds). So for any true generic of the form ⌜Fs are G⌝, there is a kind which is referred to by ⌜Fs⌝.

(ii) If kinds exist, most kinds are not instances of themselves: the kind DOG is not an instance of a dog, the kind FOOTBALLER is not an instance of a footballer, the kind MERLOT is not an instance of merlot etc. (The kind KIND may be a rare example, since that kind is an instance of a kind.) So the following sentence should be true:

> Kinds which are not instances of themselves are numerous.

(iii) According kind-reference theories, the generic bare plural 'Kinds which are not instances of themselves' refers to this kind:

> KIND.WHICH.IS.NOT.AN.INSTANCE.OF.ITSELF
>
> By standard Russellian reasoning, that kind is an instance of itself iff it is not an instance of itself, for if it is an instance of itself, it is a kind which is not an instance of itself and if it is not an instance of itself, it is a kind which is an instance of itself. Therefore there is no such kind.

(iv) The English sentence is true, so by (i), that kind exists. Contradiction.

6 For a response by Liebesman and Magidor to this objection, see §1.3.4.2.

(C) No kind-reference theory is a plausible theory of the meaning of bare plural generics.[7]

Call this the *reflexivity objection* to kind-reference. This argument is an instance of the Russellian reasoning that Boolos levied against set-theoretic interpretations of plurals; as he argued, "[t]here are some sets of which every set that is not a member of itself is one, but there is no set of which every set that is not a member of itself is a member" (Boolos, 1984: 442). Analogously, there may be kinds which are not instances of themselves, but there is no kind whose instances are all and only the kinds which are not instances of themselves. We should therefore follow Boolos's advice: "Abandon, if one ever had it, the idea that use of plural forms must always be understood to commit one to the existence of sets (or 'classes', 'collections', or 'totalities')" (Boolos, 1984: 442) – to this list we may add *kinds*.

I should note that this argument against kind-reference theories does not constitute a reductio of the hypothesis that in many (perhaps even most) cases there is a kind which corresponds to the things referred to by a generic bare plural. Just as Boolos's argument does not establish (nor was it intended to establish) that in many cases there is no set all of whose members are to the things referred to by some plural noun phrase, the reflexivity objection does not establish that there may not be a kind all of whose instances are the things referred to by some generic bare plural.

Instead, what this argument shows is that generic bare plurals cannot *in general* refer to kinds, so a semantic theory for generics cannot appeal to kinds as objects of reference. We might, if we choose, go on to invoke an ontology of kinds in other theories, but the reductio shows that we must do without kinds when giving a semantic theory for generics.

7 Liebesman and Sterken make the observation that the kind-reference theory is threatened by this paradox, but do not seem to take it seriously as they suggest that to avoid the paradox some restriction will need to be placed on the condition that, for a kind-predicational generic to be true, there must exist some kind which is the referent of the bare plural, i.e. (i). But the problem is not one that can be evaded: the theory predicts that the Russell sentence is not true (because the argument of the predicate fails to refer). So it makes incorrect predictions about the truth value of the Russell sentence, which is a perfectly meaningful and well-formed sentence of English.
Another obvious response would be to rule out object-language kind-talk, which would make the problematic sentence of English ill-formed. However, it is natural to say things like 'That kind of animal barks' while pointing at a dog (Liebesman, 2011: 433), which suggests that explicit kind-reference in English is perfectly acceptable.

1.3.4 *Plural predicational theories*

According to a less popular version of the predicational theory, generically interpreted bare plural noun phrases are referential, but do not refer to a kind (Declerck, 1991; Koslicki, 1999; Moltmann, 2013).[8] On this view, they do not refer to any *one* object, but to many. This theory retains the simple predicational logical form of characterising sentences, but replaces the constant referring to a kind with a plural constant (notated *aa*, *bb*, *cc* etc.).

(27) cool(*pp*)

The notion of plural reference is now well-established in philosophy.[9] The idea, as it applies to generics, is that a bare plural refers to the things which satisfy its description, rather than to some object, like a kind or a mereological sum, which stands in as a proxy for those things (I will sometimes use the dummy singular term 'plurality' to stand for some things). So in (14a), 'philosophers' refers plurally to individual philosophers, rather than singularly to some composite or abstract object.

One immediate advantage of this theory is that it avoids the reflexivity objection, since the paradoxical sentences only commits the plural predicational theory to the existence of kinds which are not instances of themselves, rather than a kind whose instances are all and only the kinds which are not instances of themselves.

A second advantage of this approach is that it avoids the property objection to the simple predicational theory. This is because while what it takes for a kind to have a property of individuals may seem mysterious, it is straightforward to see how a plurality could have a property of individuals in virtue of the properties of its members.

An important question for any plural reference theory of generics to answer concerns which objects are referred to by bare plurals. For example, which philosophers does the

8 Alternatively, following Chierchia (1998), kinds are sometimes understood as plural objects (see also Dayal, 2006; Guerrini, 2024); according to Chierchia, kinds are functions from worlds to the mereological sums of objects in those worlds. This view is vulnerable to the reflexivity objection above, although on closer examination the stipulation that the objects which bare plurals refer to are kinds does very little work in the analysis and so can be dropped without losing the explanatory power of the analysis.

9 This is thanks to the work of Boolos (1984), Rayo (2002, 2007), Yi (2005), Oliver and Smiley (2016), Moltmann (2016) and Florio and Linnebo (2021).

plural term in (27) refer to? The simplest answer is the plurality of all the philosophers in any possible world (Moltmann, 2013: 37).[10] This suggestion will be good enough for us to consider the merits of the plural theory.

1.3.4.1 *The semantics of plural predication*

The truth-conditions of a plural predication like (27) will depend on what it takes for a plurality to satisfy a plural predicate. On a naïve conception of plural predication, a plurality satisfies a predicate iff everything among the plurality satisfies the predicate. This naïve conception is complicated by two features of plural predication. First, there are collective predicates, which can be satisfied by a plurality without being satisfied by each of its members, such as 'gather', 'scatter' or 'be numerous'. Second, even distributive predicates – those which intuitively are satisfied by a plurality iff they are satisfied by each of things among the plurality, such as 'smile', 'laugh' or 'be asleep' – exhibit non-maximal readings, where at least some exceptions are tolerated (Lasersohn, 1999; Brisson, 2003; Malamud, 2012; Križ, 2016; Križ & Spector, 2021). As Lasersohn argues, in the following scenario, the plural predication (28a) seems true but the universal generalisation (28b) seems false:

(28) *Context:* The vast majority of the townspeople are asleep, although a few are lying in their beds awake.

a) The townspeople are asleep.

b) All the townspeople are asleep.

This complicates the notion of a distributive predicate, but it provides a useful analogy between other plural sentences and generics. Given that non-maximal readings are available for non-generic plural predications, if the plural predicational theory of generics is correct, we would expect generics to exhibit non-maximality as well. This suggests that, at least in principle, the non-maximality of generics could be put down to the same source as it is for other cases of plural predication.[11]

10 For a similar but more measured claim, see Declerck (1991). I discuss the question of the referents of generic bare plurals in Chapter 3, where I will also discuss existing plural predicational theories in more detail.

11 I make good on this suggestion in Chapter 3.

1.3.4.2 *Collective and distributive readings of generics*

Another advantage of the plural predicational theory is its ability to account for a range of similarities between bare plural generics and other sentences which involve plural predication.

Consider again the examples with bound pronouns in (24). Arguably the individual-level readings of those sentences, where each cat cleans itself and each dog obeys its own master, are cases of distributive predication. A good reason to think that the availability of this reading has nothing to do with the fact that these sentences are generics is that the same reading is available for non-generic plural sentences:

(29) a) Cats clean themselves. = (24a)

b) These cats clean themselves.

(30) a) Dogs obey their master. = (24b)

b) The dogs obey their master.

These sentences all have a reading which the predicate covaries with members of the argument; i.e. which says that *each* cat cleans itself or that *each* dog obeys its master. Moreover, as Liebesman and Magidor (2025: 278-288) have argued, all of these sentences also have a collective reading which does not exhibit covariation. Although this reading is perhaps pragmatically odd for (29a) and (30a), it clearly is available for the following cases, where the collective reading is paraphrased in (a) and the distributive reading is paraphrased in (b):

(31) British politicians look out for themselves.

a) In general, British politicians look out for British politicians.

b) In general, each British politician looks out for himself/herself.

(32) These British politicians look out for themselves.

a) These British politicians look out for these British politicians.

b) Each of these British politicians looks out for themselves.

A similar ambiguity is available for other generics:

(33) Conservatives have a favourite politician.

a) There is a politician who is the favourite politician of conservatives in general (exactly one politician).

b) In general, for each conservative there is a politician who is their favourite politician (one politician per conservative).

(34) Computers are PCs or Macs.

a) Either all computers are PCs or all computers are Macs (mixed case excluded).

b) In general, each computer is either a PC or a Mac (mixed case possible).

The important point for the plural predicational theory is that there is an extremely natural explanation of these facts which has nothing to do generics. Following work by Link (1987), Roberts (1987) and Winter (2001), the availability of the distributive readings of sentences like (29b), (30b) and (32) is explained by the optional presence of a quantificational distributivity operator, $\mathcal{D}$, which can be defined using plural variables and the plural inclusion relation:

(35) $[\![\mathcal{D}]\!] = \lambda P_{\langle e,t\rangle}.\lambda x x_e.\forall y(y \preceq xx \rightarrow P(y))$[12]

As I argue in detail in Chapter 2, the plural predicational theory can then use this operator to explain the availability of covarying distributive readings in precisely the same way as the covarying distributive readings of non-generic plurals are explained. This is an important insight, because the availability of these covarying readings has been argued to be evidence in favour of a special relational generic operator (Sterken, 2016). That these readings can be explained on the plural predicational theory on the basis only of assumptions which are motivated independently of the semantics of generics is a significant advantage of the theory.

12 Assume for now that there is some explanation for the availability of non-maximal distributive readings (e.g. by introducing some restriction on the domain of the universal quantifier).

Plural predicational theories therefore represent a promising alternative to kind-reference theories, since they avoid the reflexivity objection and the property objection. Moreover, they make it possible to explain the availability of non-maximal readings and collective-distributive ambiguities in the same terms as other plural sentences, without needing to propose special explanations just for generics.

1.4 RELATIONAL THEORIES

The dominant view of the meaning of characterising sentences is that they contain a covert constituent, 'GEN', which is responsible for their generic meaning. 'GEN' is a binary operator, much like an adverb of quantification, such as 'usually' or 'always'. I will call any theory which introduces a covert binary operator into the logical form of characterising sentence a *relational theory* (Carlson, 1989), since, on this view, characterising sentences express relations between properties or sets of objects/worlds/situations. Calling a theory *relational* is therefore to say something about the logical form it assigns to characterising sentences.

The logical form of a characterising sentence like (14a) will be analysed as follows, where *gen* is the metalanguage translation of 'GEN':

(36) a) Philosophers are cool. = (14a)

b) $gen(x) : \underbrace{\text{philosopher}(x)}_{\textit{restrictor}} \underbrace{[\text{cool}(x)]}_{\textit{scope}}$

This is the same kind of tripartite structure proposed by Heim (1982) for adverbs of quantification (see also Lewis, 1975a).

(37) a) Philosophers are usually cool.

b) $\textit{usually}(x) : \text{philosopher}(x)[\text{cool}(x)]$

On this view, bare plural noun phrases are typically analysed as *indefinites* and come with a free variable which can be bound either by *gen* or an overt adverb of quantification, such as *usually*. Different relational theories will assign different

semantics to 'GEN'. The meaning assigned to 'GEN' will then generate different truth-conditions.

1.4.0.1 *Structural ambiguity*

Probably the most important data point cited as evidence that a non-relational theory cannot succeed concerns the availability of more than one reading of certain characterising sentences (Carlson, 1989; Krifka et al., 1995). On the relational theory, this can be straightforwardly explained by the partitioning of material into the restrictor and scope.[13]

(38) Hurricanes arise in this part of the Pacific.

(39) a) For hurricanes it holds: they arise in this part of the Pacific.

b) $gen(x)$: hurricane(x)[arise.in(x)(this.part.of.the.Pacific)]

(40) a) For this part of the Pacific it holds: hurricanes arise in it.

b) $gen(x)$: this.part.of.the.Pacific(x)[$\exists y$(hurricane$(y) \wedge$ arise.in$(y)(x)$]

In contrast, on both the simple and sophisticated predicational theory, it seems as though only one structure will be possible:

(41) a) arise.in(h) *simple predicational*

b) $PM([\lambda x_e$.arise.in(x)(this.part.of.the.pacific)$])(h)$ *sophisticated predicational*

This evidence, Carlson argued, supports the hypothesis that all characterising sentences contain an unpronounced relational operator. This ambiguity certainly places a burden on any non-relational theory to explain the availability of both readings.

13 Following Diesing (1992), it is assumed that any variables associated with any noun phrase introduced in the scope are bound by an Existential Closure operation.

1.4.1 *General problems for relational theories*

Before considering proposals for the semantics of 'GEN', I will briefly present two very general objections to relational theories.

1.4.1.1 *Unpronounced in all languages*

First, 'GEN' is not pronounced in any known language (Leslie, 2008; Sterken, 2017). Liebesman (2011) combines this fact with the following principle to argue that the burden of proof is on those in favour of relational theories to show that the meaning of generics cannot be explained without 'GEN'.

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

However, we don't need something as strong as PRONOUNCED to place the dialectical burden of proof on the relational theorist:

PRONOUNCED-II If a semantic theory for English posits an unpronounced, semantically significant operator that is not pronounced in any other natural language, then a semantic theory that does not posit that operator is a better theory, provided that it can explain the same range of linguistic data.

A relational theorist would, however, claim that the final condition is not met: non-relational theories cannot explain the same range of linguistic data as relational theories.

1.4.1.2 *Lack of uniformity*

Although characterising sentences are assigned a relational logical form by the relational theory, kind-predication sentences are assigned the same subject-predicate form which the simple predicational theory assigns to all generics.

(42) a) Rhinos are large.

b) $gen(x)$: rhino(x)[large(x)]

(43) a) Rhinos are rare.

b) ✗ $gen(x)$: rhino(x)[rare(x)]

c) rare(r)

This is because (43b) says that individual rhinos are generally rare, which is not what (43a) means. Instead, (43a) says something about rhinos collectively and so is analysed as a kind-predication sentence (e.g. Krifka, 1987; Wilkinson, 1991).

A serious problem for an analysis of this kind is the existence of mixed sentences, such as (17a). Recall that (42a) and (43a) entail:

(44) Rhinos are rare and large. = (17a)

The surface structure of this sentence involves a single argument and a conjunctive predicate, but this conjunctive predicate cannot constitute the scope of *gen*, since that structure asserts that individual rhinos are generally rare and large.

(45) ✗ $gen(x)$: rhino(x)[rare$(x) \wedge$ large(x)]

Since the predicate 'rare' applies to the whole class, not just individuals, this gets the truth-conditions wrong. Any relational theory will therefore need to tell some more complicated story about how sentences like (44) come to have the truth-conditions that they have (Koslicki, 1999).

I leave this structural difficulty aside for now; I discuss a referential variation of the relational theory which can account for at least some mixed sentences below.[14] I now turn to the details of the semantic analyses that have been proposed for 'GEN'.

1.4.2 *Probability theories*

I begin by discussing Ariel Cohen's relational theory, according to which generics express probability judgments. Cohen (1999b, 1999a) treats *gen* as a relation that holds between properties. (36a) then has the following logical form.

14 I discuss mixed sentences in more detail in Chapters 2 and 3; for some further discussion, see the debate in Liebesman (2011) and Leslie (2015).

(46) $gen(\text{philosopher}, \text{cool})$

A generic $gen(\psi, \phi)$ is true (very roughly) iff $P(\phi \mid \psi)$, the probability that something is ϕ, conditional its being ψ, is greater than 0.5.[15] This analysis meets with immediate counterexamples from generics which are true of less than half of the things that have a property, for example:

(47) a) Lions have manes.
b) Swallows lay eggs.

Cohen's solution to this problem is to relativise the probability judgment to instances of the restrictor property that are instances of at least one member of the set of alternatives to the scope property. For example, in (47a), the alternatives to the scope property, having manes, are alternative modes of sexual ornamentation, such as having beautiful plumage, large antlers or, trivially, a mane.[16] Cohen argues that (47a) says that the probability that something has a mane, given that it is a lion that has at least one of the alternatives to having a mane, is greater than 0.5. That explains why (47a) does not entail (48); (48) is evaluated relative to a different set of alternatives: different sexes that a lion might be.

(48) Lions are male.

(48) is true iff the probability that something is male, given that it is a lion of some sex, is greater than 0.5. Since this probability is not greater than 0.5, (48) is false.

One problem for this probabilistic approach is posed by Nickel (2008), who observed that the theory makes the wrong predictions for generics with conjunctive predicates:[17]

(49) a) Bears live in North America, South America, Europe, and Asia.
b) Elephants live in Africa and Asia.

15 Tessler and Goodman (2019) have suggested amending this to an underspecified threshold, which is resolved pragmatically in a context.

16 The set of alternatives may also be affected by the prosodic focus in the utterance (cf. Rooth, 1992). Variation in the alternatives induced will lead to different interpretations of the generic (Cohen, 2004; Kirkpatrick, 2024a).

17 For detailed discussion of generics with conjunctive predicates and some possible strategies for explaining their meaning, see Kirkpatrick (2023).

Consider (49b). This could be interpreted as either of the following, depending on the scope of *gen*:

(50) a) $gen(\text{elephant}, \text{lives in Africa}) \wedge gen(\text{elephant}, \text{lives in Asia})$

b) $gen(\text{elephant}, \text{lives in Africa and Asia})$

On the (a) interpretation, Cohen predicts that the sentence says that the probability of something living in Africa given that it is an elephant is greater than 0.5 and the probability of something living in Asia given that it is an elephant is greater than 0.5. Absent a significant number of transcontinental elephants, that conjunction must be false. Yet (49b) is true, so this cannot be the correct analysis. There presumably will be no difference in the set of alternatives induced by each conjunct (i.e. locations in which an elephant might live), so it is not likely that any attempt to vary the restrictor property between the conjuncts to deliver a different truth-value will be successful.

On the (b) interpretation, the sentence says that the probability that something lives in both Africa and Asia given that it is an elephant is greater than 0.5. Again, assuming that very few (if any) elephants live in both, the sentence is predicted to be false. Either way, the probability theory makes incorrect predictions about the truth value of (49b).

1.4.3 *Normality theories*

According to normality theories, it is part of the meaning of 'GEN' that it quantifies over entities that are normal in some respect. The intuition behind normality theories is that for (36a) to be true is for it to be normal for philosophers to be cool. The fact that there are abnormal philosophers who are not cool is therefore no barrier to (36a) being true. The underlying idea is that for the purposes of the truth of characterising sentences, we can ignore unusual or atypical objects or circumstances.

1.4.3.1 *Normal worlds*

On one version of the normality theory, the semantics of 'GEN' can be stated in terms of a relation between two sets of possible worlds. I take the theory of Pelletier and Asher (1997) as a representative example (see also Delgrande, 1987; Asher & Pelletier, 2013). They argue that (36b) is true at a world w in the following circumstance:

> For every object d the set of worlds that are normal with respect both to w and to the proposition of d's being a philosopher is a subset of the set of worlds in which d is cool (adapted from Pelletier & Asher, 1997: 1160).

Since these truth-conditions depend not only what happens at w, but at other 'normal' worlds too, accidents – like one generation of philosophers being particularly cool – are not on their own enough to make a generic true. Instead, the truth of (36a) depends on how things go in counterfactual but very similar worlds. So if it is only by chance that many philosophers are cool, (36a) may still be false, provided that there are many uncool philosophers in other worlds which are normal with respect to those things being philosophers.

WHAT IS A NORMAL WORLD? Pelletier and Asher define a world's being normal with respect to some proposition as follows:

NORMAL WORLDS w is a normal world with respect to a proposition p and a world w' iff p is true in w and all the typical consequences of p in w' hold in w (Pelletier & Asher, 1997: 1160).

So if we want to know whether a world is normal with respect to the proposition that some person a is a philosopher, we need to look at all the worlds in which a is a philosopher and in which all the typical consequences of a's being a philosopher hold. This approach does not require that there are any worlds where every philosopher is normal; instead, the truth-conditions of a generic require us to look, for every philosopher a, at each world in which a is a philosopher and the typical consequences of a being a philosopher hold, and check whether a is cool in that world.

The first problem with this proposal is that the worlds that are under consideration are relativised to each individual. Consider a very abnormal philosopher, Martin. Now consider a world in which the typical consequences of Martin's being a philosopher hold. Is that a world in which Martin has the properties of a normal or typcial philosopher? Presumably not: the typical consequences of *Martin's* being a philosopher might still be quite unusual for philosophers (imagine that it is in Martin's nature to be a contrarian). So even if it is normal for philosophers to be cool, the generic 'Philosophers are cool' would be falsified by the existence of an abnormal and uncool philosopher.

The second problem is that this definition of a normal world invokes the 'typical' consequences of a proposition holding at a world. What sense should be attached to 'typical' here? It seems that this cannot be a purely statistical matter, since Asher and Pelletier later suggest that what counts as a world which is normal in some respect might not depend only on what are the most frequent consequences of a proposition being true.

(50) Turtles live to be 100.

Regarding (50), they say:

> Suppose we want to evaluate [(50)]. So consider [a turtle] Tim and any normal Tim-turtle world. There are several ways in which we could think of such worlds. One is from the Aristotelian perspective of what is the natural 'telos' of a turtle: if everything goes right or normally for a given turtle, then he lives to be 100. This perspective would make [(50)] true. On the other hand, if we consider the way life proceeds for a turtle based on statistics for turtles in our world, another way of construing normality, then matters are rather more grim for Tim. Most turtles actually die within a few hours of being hatched. This way of construing normality would make [(50)] false at the actual world. A variety of contextual factors, including preceding discourse, typically fixes or at least precisifies the sense of normality at issue (Asher & Pelletier, 2013: 314).

To accommodate the truth of this example, Asher and Pelletier cannot rely on an account of normality solely in terms of prevalence or probability. They claim that the possibility of this 'give' in the notion of normality is an advantage of their view, since it allows them to account for variation in our intuitions about the truth-value of different utterances of the same generic.

The appeal to what happens under ideal conditions is quite unhelpful here, however, since there are many generics which should be true if they can be understood in terms of things going 'right':

(51) a) Professional footballers don't get injured.

b) iPhones last for decades.

c) Spacecrafts survive launch.

It seems as though if everything goes just right for a professional footballer, they would never get injured. Yet it is highly abnormal for a professional footballer never to get injured and (51a) is intuitively false.

This suggests that the 'typical consequences' mentioned in NORMAL WORLDS should not be understood in terms what is the case if everything 'goes right'. Perhaps it would be better to analyse normality in purely statistical terms; that is, q is a typical consequence of p in w iff usually, if p, q in w. But then the truth of (50) goes unaccounted for. Moreover, many generics which say something about what is typically the case are intuitively false (Cohen, 1999a):

(52) a) Humans are over three years old.

b) People have black hair.

c) Bees are sexually sterile.

Consider a world w which is normal with respect to my being a human and the actual world $w_@$. The typical consequences of my being a human hold in w, so I am over three years old in w. Given that the same will be true for all humans and all worlds which are normal with respect to their being human, (52a) is predicted to be true by Pelletier and Asher. Yet it is false.

I conclude that the notion of normality invoked by Pelletier and Asher's normal worlds theory is not adequate for the semantics of generics.

1.4.3.2 *Normal objects*

An alternative to the normal worlds approach is to quantify directly over normal objects. To capture the non-accidental nature of generics, these analyses of 'GEN' typically involve universal quantification over worlds as well as objects, but the restriction in the domain of worlds quantified over is not cashed out in terms of normality (e.g. Eckardt, 1999: 241-243). Greenberg (2003, 2007) states her semantics for generics as follows, where R is an accessibility relation on worlds and $P^{\text{cont.norm}}$ is the property of being a contextually relevant and normal P.[18] 'Ps are Q' is true in a world w iff:

> $\forall w'(wRw' \rightarrow \forall x(P^{\text{cont.norm}}(x, w') \rightarrow Q(x, w'))$
> in all appropriately accessible worlds, every contextually relevant and normal P has the property Q.

According to both Greenberg and Eckardt (1999), the appropriate accessibility relation is one which relates worlds based on their similarity; on Eckardt's version, wRw' iff w' is like w with respect to "causal and statistical regularities", although w' may differ from w with respect to "isolated accidental facts" (Eckardt, 1999: 243).

Proponents of the normal objects view typically have little to say about what makes something a normal P, beyond the triviality that the normal Ps in a world are a subset of the Ps in that world. Even granting a satisfactory analysis of normality, examples like the following pose a challenge to this approach, since they seem, on at least one very plausible reading, to require that *all* prime numbers are odd and that *all* lottery tickets are losers (Nguyen, 2020).

(53) a) Prime numbers are odd.
b) Lottery tickets are losers.

18 Greenberg's additional restriction to contextually relevant individuals is introduced to account for the fact that generics can be felicitously contextually restricted (see Greenberg, 2007).

But on the normal objects approach, there should be no reading of these sentences according to which they are false, since, for example, 2 is abnormal with respect to prime numbers and so should be outside of the domain of quantification.

NORMAL NORMAL OBJECTS Another problem for this approach is that generics with noun phrases modified by 'normal' are perfectly felicitous (Carlson, 1977: 64):

(54) a) Ravens are black.

b) Boys play football.

c) Ticks carry Lyme disease.

(55) a) Normal ravens are black.

b) Normal boys play football.

c) Normal ticks carry Lyme disease.

Yet if the normal objects view is right, then the sentences in (55) mean something like:

(56) a) All normal normal ravens are black.

b) All normal normal boys play football.

c) All normal normal Ticks carry Lyme disease.

One response here would be bite the bullet and say that the normal normal ravens are just the normal ravens, so, while the (56) cases sound odd, they are equivalent to (55). That is a substantive commitment: presumably some normally coloured ravens are more normal than others, and the least normal of these will not be among the normal normal ravens. But even if we grant this response, we would also need to concede that the unmodified sentences in (54) have the same truth-conditions as both the (55) and (56) sentences, which seems intuitively wrong. This is particularly obvious for the pair of (54c), which arguably has a true reading (see Leslie, 2008), and (55c) which does not have a true reading.

The only plausible repair strategy here would be to concede that the restriction to normal objects is not 'hard-wired' into the meaning of 'GEN'. Kirkpatrick (2023) takes

this approach, claiming that the restriction to normal objects arises as the result of the accommodation of a question-under-discussion which concerns only normal objects. The additional implicit restriction could then be absent when there is an explicit restriction to normal objects. Although Kirkpatrick does not discuss the possibility, it is possible, on this theory, that in response to different questions, a different restriction on the domain of quantification will arise. This concession abandons the normality theory, in favour of a more flexible theory in which the restriction on the domain of quantification is always determined by context.

1.4.3.3 *Ways of being normal*

Nickel (2008, 2016), swayed in large part by the difficulties posed by conjunctive predicates (see his examples in (49)) proposed a radical departure from the orthodox normality theories, suggesting that generics involve not universal quantification over normal objects or worlds, but *existential* quantification over *ways of being normal*. He observed that the same challenge he posed to Cohen's probability theory applies to normality theories, since the theories make incorrect predictions when there is more than one way that an object can be normal in some respect (Nickel, 2016: 63).

(57) Elephants live in Africa and Asia. = (49b)

(58) a) $gen(x)$: elephant(x)[lives.in.Africa(x)] $\wedge$ $gen(x)$: elephant(x)[lives.in.Asia(x)]
b) $gen(x)$: elephant(x)[lives.in.Africa.and.Asia(x)]

As he argues, the (a) analysis requires that all normal elephants live in Africa and and all normal elephants live in Asia, which cannot both be true, while the (b) analysis requires that all normal elephants live in both, which is false. Since (57) is true, the orthodox normality analysis fails.

Nickel seeks to repair this by claiming that for a generic ⌜Fs are G⌝ to be true, there must be at least one way of a being a G-normal F such that all Fs that are normal in that way are G. To be a G-normal F is to be an F which is normal with respect to a determinable of G (Nickel, 2016: 59); so for an elephant to be an Africa-normal elephant is for it to be an elephant which is normal with respect to its habitat. Or: for

a raven to be black-normal raven is for it to be a raven which is normal with respect to its colour.

Nickel's solution to the predicate conjunction problem lies in the fact that his semantics for *gen* allow for there to be more than one way of being a G-normal F. So, for example, there is a way of being an Africa-normal elephant such that all elephants that are normal in that way live in Africa, and there is a way of being an Asia-normal elephant such that all elephants that are normal in that way live in Asia. So both conjuncts of (58a) are true (but (58b) is false, since living in both Africa and Asia is not a way of being a normal elephant).

A serious problem with this solution is that it makes sentences like the following true (Hoeltje, 2017: 107):

(59) a) Swallows are female.

b) Peafowl are male.

The reason is simple: there is a way of being a female-normal swallow such that all swallows that are normal in this way are female. That is because there is nothing abnormal about female swallows with respect to their sex; both being male and being female are ways of being normal swallows, according to Nickel (cf. Nickel, 2016: 56 n.7). Although Nickel is happy to bite that bullet, his intuition that these sentences are true is not widely shared (see e.g. Leslie, 2008: 36; Asher and Pelletier, 2013: 327; Hoeltje, 2017).

Moreover, as Kirkpatrick (2023: 390, 420-423) argues, Nickel's insistence that the correct analysis of generics with conjunctive predicates is as sentential coordinations (as in (58a)) is not syntactically motivated, since the conjunctive predicate in (57) passes constituency tests:

(60) FRAGMENT ANSWERS

Q: Where do elephants live?

A: In Africa and Asia

(61) TOPICALISATION

In Africa and Asia is where elephants live.

As the conjunctive predicate is a syntactic constituent of the sentence, it cannot be split in the way that Nickel suggests. So although Nickel's theory is motivated by the need to account for generics with conjunctive predicates, the details of his analysis are not plausible.

1.4.4 *Referential relational theories*

Partly in recognition of the difficulties posed by the compatibility of generic bare plurals with conjunctive and mixed predicates, some attempts have been made to reconcile the relational theory with the plural and kind-predicational theories discussed above (Chierchia, 1998; Nickel, 2016; Cohen, 2020; Kirkpatrick, 2023; Guerrini, 2024). According to Chierchia's influential proposal, generic bare plurals are referential, but individual-level and habitual predicates introduce the generic operator, 'GEN', giving a structure like the following for (14a) (Chierchia, 1995, 1998; see also Guerrini, 2024):

(62) a) [philosophers [GEN [are cool]]]

b) $gen(x, s) : x \leqslant \iota y.\text{philosopher}(y)(s) \wedge C(x, s)[\text{cool}(x)(s)]$

c) It is generally true of objects x and situations s such that x is a mereological part of the sum of all philosophers in s and s is a contextually relevant situation involving x, that x is cool in s.

This structure arises as a result of a type-shifting mechanism, $\cap$ (*nom*), which takes a property-denoting bare plural, $[\![\text{philosophers}]\!] = [\lambda s_s.\lambda x_e.\text{philosophers}(x)(s)]$, into a function from a situation to a mereological sum (see also Dayal, 2006):

(63) $$[\![\cap\text{philosophers}]\!] = \begin{cases} [\lambda s.\iota x\text{philosophers}(x)(s)] \\ \text{if } [\lambda s.\iota x\text{philosophers}(x)(s)] \text{ is in } K \\ \text{undefined otherwise} \\ \text{where } K \text{ is the domain of all kinds} \end{cases}$$

One major advantage of this approach is its ability to explain the felicity of generics with mixed predicates like (17a) (see Nickel, 2016: 255):[19]

(64) a) [rhinos [[are rare] and [GEN [large]]]]

b) $[\lambda K_{\langle s,e\rangle}.gen(x,s) : x \leqslant K(s) \wedge C(x,s)[\text{large}(x)(s)] \wedge$
$\text{rare}(K)]([\lambda s_s.\iota y\text{rhinos}(y)(s)]) =$

c) $gen(x,s) : x \leqslant \iota y.\text{rhinos}(y)(s) \wedge C(x,s)[\text{large}(x)(s)] \wedge \text{rare}([\lambda s_s.\iota y\text{rhinos}(y)(s)])$

Because they resolve this issue, referential relational theories, represent the most promising form of the relational theory; however, the move to a referential theory only helps to resolve the structural problems above. The question of the semantic analysis of 'GEN' remains unresolved.

1.4.5 *Psychological theories*

Another version of the relational theory takes a radically different approach to the meaning of generics, essentially conceding that the task of giving an informative semantic analysis of 'GEN' is unachievable. According to Sarah-Jane Leslie's (2007, 2008) influential theory, generics are the expression of the cognitive system's "default" mode of generalisation. This mode of generalisation is not linguistic but predates the acquisition of language in cognitive development (Leslie, 2007: 381). This explains the "near-paradoxical nature of the acquisition of generics" (Leslie, 2008: 16): children develop the ability to use and understand generics from a very early age, before they learn to use explicit quantifiers (Leslie, 2008: 19).[20] Leslie's hypothesis neatly explains why this should be the case, since children have an innate ability to generalise in this 'default' manner, the linguistic expression of which is a generic.

Leslie gives a disquotational semantics for 'GEN'; that is, she thinks there is nothing more to say, from the perspective of linguistic semantics, about the meaning of a generic than the structure in (36b) which uses the metalanguage operator *gen* to state the meaning of a generic, just as we can say what it takes for a sentence like 'The ball

19 See Kirkpatrick (2023) and Guerrini (2024) for discussion of cumulative predication.
20 See Kirkpatrick (2024b) for some critical discussion of the acquisition data that Leslie relies on.

is red' to be true without giving a substantive analysis of what it takes for an object to be red.

(65) 'Philosophers are cool' is true iff $gen(x)$: philosopher(x)[cool(x)]

However, Leslie claims that we can give an analysis of "how the world must be for [a generic] to be true" (2008: 43):

> We can describe the circumstances under which a generic of the form 'Ks are F' is true as follows:
>
> (i) The counterinstances are negative, and:
>
> (ii) If F lies along a characteristic dimension for the Ks, then some Ks are F, unless K is an artifact or social kind, in which case F is the function or purpose of the kind K;
>
> (iii) If F is striking, then some Ks are F and the others are disposed to be F;
>
> (iv) Otherwise, almost all Ks are F (Leslie, 2008: 43).

What it takes for a counterinstance to be 'negative' is psychological; a counterinstance having some psychologically salient alternative property would count as positive, whereas merely lacking the relevant property would count as a negative counterinstance (Leslie, 2008: 34).

Strictly speaking, this is not a semantic analysis of 'GEN', since Leslie is not claiming that these conditions are part of the linguistic meaning of a generic. But the theory still does make some predictions about the truth values of generics. Sterken (2015c) offers some criticism of the claim that these conditions adequately capture the conditions under which generics are true. For instance, in the following example, the counterinstances are not all negative: there are mammals which lay eggs (platypuses) and laying eggs is a positive property. So (66) should be false, since the first condition is not satisfied.

(66) Mammals give birth to live young.

But (66) has a true reading, so Leslie's theory makes the wrong prediction here. Sterken also identifies counterexamples to the third condition, where Leslie predicts the sentence to be true, but it is in fact false (Sterken, 2015c: 2500):

(67) Homeowners start fires.

By Leslie's standards, 'start fires' expresses a striking or dangerous property, so, given that some homeowners do start fires, (67) should be true. But it is false, so once again Leslie's conditions are not adequate.

Given these counterexamples and Leslie's claim not be giving a semantic analysis of 'GEN', it is reasonable to ask what exactly her theory is a theory of. Leslie's methodology is to reflect on speaker judgments about the truth values of generics and then try to make generalisations about the conditions under which they arrive at judgments of truth. Yet, as I argue in detail in Chapter 5, all that this methodology can tell us is how ordinary speakers make judgments about the truth of generics. The problem is that these truth value judgments are not made relative to clearly specified conditions. For example, a speaker might judge (36a) to be false without knowing much about the actual distribution of cool philosophers or whether there are certain characteristics of philosophers that lead them to be cool or not. So how could we expect these judgments to accurately reflect the truth values of generics? As Jönsson and Hampton (2006) suggest is the case for universal generalisations, it seems likely that speakers employ some kind of heuristic to make these judgments, especially in the absence of clearly specified conditions against which to make the judgments. If that is the case, these judgments are not themselves a good guide to the conditions under which generics are true (as opposed to the conditions under which they are believed to be true).

Leslie's approach may therefore tell us about the factors that cause us to judge generics true or false, but this is not the same as revealing their truth-conditions. It is perfectly consistent with the data that Leslie relies on for her theory that our judgments about the truth values of generics are systematically mistaken in ways that are modulated by cognitive biases. We should, therefore, not view Leslie's theory as a competitor to compositional semantic theories of generics.

1.4.6 *Indexical theories*

The difficulty of giving a uniform semantics for 'GEN' has led some authors to suggest that the meaning of generics is context-sensitive, which would explain why in some contexts generics seem to express very strong generalisations, but only weak ones in others. This could explain, for example, why the examples in (53) have a reading which is false: in some contexts they express a universal generalisation which is falsified by a single exception.

Sterken (2015a) suggests that 'GEN' is a phonologically null indexical expression, whose possible values are adverbs of quantification of different strengths. Sterken goes further and claims that the value that 'GEN' takes in a given utterance can be underdetermined; that is, there can be no fact of the matter which generalisation is expressed by an utterance of a generic.

Given the range of data considered so far, it is extremely plausible that the meaning of a generic can differ from context to context and that it may not be fully determinate which generalisation is expressed by a given utterance of a generic. But granting this is not the same as granting that 'GEN' itself is an indexical expression. Sterken's argument in favour of this conclusion is that the status of 'GEN' as an indexical is the best explanation of these facts (Sterken, 2015a: 22). This claim is hard to assess, since it rests on a comparison between a contextualist theory and a dizzying array of alternatives. Things are made even worse by the fact that Sterken gives little more than a placeholder theory of the semantics of 'GEN', focusing instead on its metasemantics in an attempt to show that this can explain our intuition that the meaning of generics can be underspecified in a context.

1.4.6.1 *Some possible constraints*

A more detailed alternative is proposed by Lee and Nguyen (2022), who attempt to show that there are certain constraints on the quantifiers that an indexical 'GEN' can take as its value. They give a lexical entry for this indexical, where g is a function which assigns semantic values to variables:

(68) $[\![\text{GEN}_i]\!]^g = \begin{cases} \text{defined iff } g(i) \text{ is upward monotone and non-symmetric;} \\ \text{if defined, } g(i) \end{cases}$

The notions of monotonicity and symmetry are from Generalised Quantifier Theory (Westerståhl, 2015):

UPWARD MONOTONE Q is upward monotone iff $Q(A,B) \wedge B \subseteq C \Leftrightarrow Q(A,C)$.

NON-SYMMETRIC Q is symmetric iff $Q(A,B) \Leftrightarrow Q(B,A)$. Q is non-symmetric iff Q is not symmetric.

The restriction on the possible values for 'GEN' is intended to rule out possible meanings like the following:

(69) a) [#Few] mosquitoes transmit malaria. *downward monotone*
b) [#Exactly five] barns are red. *non-monotone*
c) [#At least five] girls cry. *symmetric*

This kind of constrained approach is welcome, but it does not rule out enough. To see why, notice that the possible semantic values for 'GEN' that Lee and Nguyen propose are quite liberal. For example:

(70) a) [Under normal circumstances, almost all] ravens are black.
b) [Ideally, all] boys don't cry.

If the possible meanings for 'GEN' are as unconstrained as this, what prevents the following possible meanings for generics?

(71) a) [Under abnormal circumstances, almost all] ravens are black.
b) [In non-ideal conditions, all] boys don't cry.

These values do not violate the semantic constraints that Lee and Nguyen propose, yet these are intuitively not possible meanings for the bare plural sentences. The difficulty for the indexical account is to specify constraints on possible values for 'GEN' which generate enough of the possible meanings for generics, without predicting

unattested meanings as well. Sterken (2015a: 23) claims that this worry is not really an objection, but a "call for more theorising"; that may well be correct, but until such theorising has taken place, the indexical account of generics constitutes little more than a promissory note for a theory.

1.4.7 *Underspecification theories*

According to a close relative of the indexical theory, the apparent multiplicity of meanings that generics can express is explained by the fact that the meaning of generics is underspecified in some sense. Nguyen (2020), for example, argues that generics express *propositional radicals*, that is, representations which are not fully propositional, but can be transformed into full propositions by a pragmatic process of completion (see also Bach, 1994). The incomplete representation of (36a) will then be as follows:

(72) $__ : \text{philosopher}(x)[\text{cool}(x)]$

Since this structure is missing a determiner, it only expresses a proposition when the structure has been completed. However, as Lee and Nguyen (2022) have observed, the problem with this proposal is that it is insufficiently constrained: there are many completions of this structure which are not possible meanings of (36a), such as:

(73) a) $no(x) : \text{philosopher}(x)[\text{cool}(x)]$

b) $few(x) : \text{philosopher}(x)[\text{cool}(x)]$

That is, (36a) cannot mean that no philosophers are cool or that few philosophers are cool. Since Nguyen's theory fails to rule out these completions, it is not adequate as a semantic theory of generics.

This objection demonstrates that, like indexical theories, any underspecification theory of generics will need to specify some explicit constraints which predict only the attested meanings of generics. Another underspecification theory that fails to meet this criterion is that of Bosse (2021). Bosse claims that generics are semantically complete,

but that they express existential quantification over a domain of other generalisations. That is, (36a) says that some relation holds between philosophers and being cool, without saying which. The challenge for this account is to specify the nature of the domain; why can't the existential generalisation be made true by the relation expressed by the determiner 'no' or 'few'?

Bosse makes no attempt to state explicit constraints and responds to these overgeneration worries by claiming that in different contexts, the domain may be narrowed (Bosse, 2021: 14396). That would explain why (74a) is not true, even though the relation expressed by 'most' is often good enough to make a generic true.

(74) a) Books are paperbacks.

b) Most books are paperbacks.

However, in the absence of explicit constraints, the theory is not falsifiable, since this reply is available in response to any counterexample: for any generic that is predicted to be true by Bosse's theory, but is intuitively not true, say that the domain does not include any relation in the context of utterance. This is in itself a good enough reason not to adopt the theory until constraints on the domain of quantification have been proposed and tested.

1.5 THE NATURE OF GENERICITY

I now turn from semantic analyses of the generic operator to the question of what makes a sentence generic. I frame this as the question: *what is genericity?* Following Nickel (2017), we may distinguish this question from both the semantic questions about the meaning of 'GEN' and the structural questions about the logical form of generics. This question is restricted to characterising sentences, since what makes a kind-predication sentence 'generic' seems to be fundamentally different in character (cf. Nickel, 2017: 440; Mari et al., 2013: 42). Unlike characterising sentences, kind-predication sentences can report particular events, for example:

(75) a) Brown rats arrived in Europe in the 1300s.

b) Tasmanian tigers were hunted to extinction in the 20th Century.

Moreover, as I remarked above, these sentences do not exhibit non-maximality, in the sense that asking whether a brown rat or a Tasmanian tiger is an exception to the sentences seems like a misunderstanding. They don't exhibit nomicity either, since neither sentence suggests that there was anything non-accidental about the arrival of brown rats in Europe or the extinction of Tasmanian tigers. So in what follows, I will restrict my attention to characterising sentences.

The answer to the question *what is genericity?*, I will argue, is that what makes us judge a sentence to be a generic is the belief that the sentence is true in virtue of some stable or non-accidental state of affairs. The view that genericity is not a linguistic phenomenon has already been proposed by Collins (2015, 2018), who argued that the linguistic phenomena that have been claimed to fall under the category of genericity are too diverse to receive a single linguistic explanation. Instead, we should understand genericity in terms of our beliefs about the states of affairs that make sentences true; this means that there will be many distinct linguistic routes to genericity.

I will show that by manipulating the context of an utterance of a sentence so that it can be understood as being true in virtue of a stable state of affairs, we can modulate whether a generic interpretation of the sentence is available. This, I argue, is evidence that genericity consists in the belief that a sentence is true in virtue of a stable state of affairs – I refer to this feature as *stability*.

1.5.1 *Unboundedness?*

Before I present my argument that genericity consists in stability, I will briefly consider an alternative suggestion. Mari et al. (2013) raise the possibility that genericity consists in *temporal unboundedness*. They say:

> [T]he real alternative to the prevailing analysis based on some sort of covert generic operator is not simply to emphasize the presence of an unbounded interpretation, but rather to claim that the unbounded interpretation is not

> the effect of the semantic functioning of an ad hoc generic operator, but an independent ingredient of the semantics of generics which, by itself or in interaction with some other ingredient, produces generic meaning (Mari et al., 2013: 52).

As Mari et al. are aware, the idea that genericity arises from unboundedness cannot be correct, since there are perfectly good generic sentences which are temporally bounded:

(76) a) Before 1870, a Pope was sovereign of the Papal States.
b) In the Middle Ages, poor people went hungry.
c) In the 1990s, baggy jeans were cool.
d) Since 1957, the Coke bottle has been contoured.

There is nonetheless an important relationship between genericity and time: what explains our intuition that a sentence is generic is the belief that it is, if true, true in virtue of a state of affairs which is stable over time.[21] It easy to see how this could be the case for sentences which are unrestricted with respect to time, but even sentences with explicit temporal restrictions, like those in (76), may be true in virtue of stable states of affairs (or at least a state of affairs that is believed to be stable).

This criterion for genericity divorces genericity from the linguistic properties of sentences. A relational theorist might, for example, claim that genericity just consists in a sentence containing 'GEN' as a constituent. A kind-predication theorist, on the other hand, could claim that genericity consists in some constituent in a sentence referring to a kind (although this would not explain the genericity of indefinite singulars, which cannot refer to kinds). But, on the view I propose here, these linguistic features do not on their own determine whether a sentence can be interpreted as a generic.[22]

21 By 'true in virtue of', I intend a relation of truthmaking or grounding.
22 The view I propose is, of course, compatible with those semantic theories. For example, one way that a sentence could be true in virtue of a stable state of affairs might be for it to predicate a property of a kind; if kinds usually have their properties stably over time, this would meet my criterion. But kind-predication would be just one way of a sentence being generic.

1.5.2 *Stability*

I proceed by showing that by providing context which suggests that a sentence is true in virtue of a stable state of affairs enables sentences which would otherwise have no acceptable generic reading to be interpreted as generics. I begin by considering a pair of examples from Carlson (1977: 318).

(77) a) ?Alligators in the next room are intelligent.
b) ?Alligators in the New York sewer system are intelligent.

It is very hard to get a generic reading of either sentence out of the blue. Carlson argues that if we imagine that the alligators in the New York sewer system constitute a race or a type of alligator, then (77b) gets a perfectly acceptable generic reading. But what seems to matter in generating the generic reading of (77b) is that the sentence does not just say something about the alligators which are currently in the New York sewer system, but about the alligators which have been present in the sewers for some sufficiently long period of time. All we really need to get a generic reading is the assumption that alligators have been inhabiting the New York sewer system for some time:

(78) Since the nuclear disaster, there have been alligators inhabiting a number of major US cities, although some are more impressive than others. Alligators in the New York sewer system are intelligent.

It is possible to provide a context which supports a stable interpretation even for sentences like (77a):

(79) Welcome to our alligator experimentation unit. We subject all our alligators to intelligence testing. We keep the stupid ones in this room. Alligators in the next room are intelligent.

The context here supports the idea that we are not just talking about the alligators currently located in a room, but about whatever alligators have been, are or will be in

that room. Given that context, it is straightforward to interpret (77a) as being true in virtue of a stable state of affairs and hence as a generic.

Further evidence in favour of the characterisation of genericity in terms of stability is that, by progressively narrowing the temporal interval for which the generic is said to hold, it becomes harder to understand a sentence as generic:

(80) a) Politicians are intelligent.
b) In the 20th century, politicians were intelligent.
c) In the 1980s, politicians were intelligent.
d) In 1987, politicians were intelligent.
e) In November 1987, politicians were intelligent.
f) ?On 10 November 1987, politicians were intelligent.
g) ?On 10 November 1987 at 14:21, politicians were intelligent.

(80a)-(80c) are easily understood as generics with no supporting context, but it is very difficult to hear (80f) and (80g) as generics. This contrast can be explained by the hypothesis that genericity consists in stability. (80g) seems to say something about a very specific moment in time, which is why it is difficult to hear it as a generic.

That modulating context to suggest that a sentence is true in virtue of a stable state of affairs turns bare plural sentences which have no acceptable generic reading into acceptable generic sentences is good evidence that stability is indeed what the generic reading consists in.

1.5.2.1 *Habituals*

The characterisation of genericity as stability can also help to explain the sense in which habitual sentences are generic. Habituals are sentences like the following:

(81) a) Matt runs to work.
b) When she works from home, Pat takes a walk at lunch.
c) John used to play the guitar.

In each case, the sentence describes some regularity which holds across an extended period of time (see Comrie, 1976: 27-28). As (81c) demonstrates, the series need not be

completely unbounded, since it describes a series of events which is wholly located in the past. But, just as we saw was the case for bare plural sentences, habituals are not felicitous when the period of time they describe is very short:

(82) a) ?On 10 November 1987, Matt used to run to work.
b) ?Today, Pat takes a walk at lunch.

The characterisation of genericity as stability therefore allows us to say that habituals exhibit genericity, without making any commitment to there being some constituent of habituals that is shared with all other generic sentences (see Krifka et al., 1995: 32 for a claim like this).[23]

I conclude that genericity consists in a sentence being believed to be made true by a state of affairs which is stable across time. This account of genericity is not tied to the meaning of any particular expression, but explains genericity in terms of our beliefs about the states affairs that make sentences true. What remains is to deliver semantic theories for different linguistic constructions which generate truth-conditions which explain why those constructions are capable of expressing generic meanings. My objective in this thesis is to offer a semantic theory for bare plurals which achieves this goal.

1.6 ROADMAP

The semantic theory that I will propose is a plural predicational theory: generic bare plurals are referential and refer plurally; characterising sentences do not contain an unarticulated generic relational operator. I will argue that what makes generic bare plurals distinctive is that the domain of their reference is extended throughout time, including into the future. I propose that generic bare plurals are referential because they contain a unpronounced determiner which differs in meaning from the definite article 'the' only in its presuppositions. I call this the *temporal plural reference theory* of

23 See Dobrovie-Sorin (2001) and Rimmel (2004) for arguments against deriving the meanings of habituals and generics from the same source, as well as relevant discussion by Ferreira (2005).

generics. I will argue that this theory represents the best way of accounting for a wide range of linguistic data. The thesis will proceed as follows:

CHAPTER 2 I present novel data concerning the availability of collective and distributive readings of generics and show that the range of readings available for generics is the same as for other, non-generic plural noun phrases. I show that non-referential and certain versions of the referential theory fail to explain this data and argue that this evidence weighs heavily in favour of a plural reference theory of generics.

CHAPTER 3 I present the linguistic motivations for treating generic bare plurals as covert definite descriptions and introduce the details of the temporal plural reference theory. I give a lexical entry for a null determiner and show how this predicts both the differences and similarities in the meanings of generic bare plurals and plural definite descriptions. I show that this theory is compatible with a range of linguistic data concerning generics which have previously been thought to motivate either a kind-predication or relational theory.

CHAPTER 4 I present a puzzle concerning the nomic interpretation of generics and show that generics with a temporal restriction wholly located in the past elicit diminished intuitions of nomicity. I argue that this motivates a theory which takes the temporal aspects of the meaning of generics seriously and present a novel truthmaker semantics for plural predication which explains the nomicity of generics in terms of the different ways that they can be made true. Finally, I show that the temporal plural reference theory can help to make sense of some contested empirical data concerning the relationship between hearing generics about social kinds and the formation of psychological essentialist beliefs about those kinds.

CHAPTER 5 I discuss Leslie's Psychological Theory of Generics in detail and argue that it does not represent a genuine alternative to semantic theories of generics, since it takes speakers' judgments at face value and ignores the non-semantic factors that might affect these judgments. I present a reconceptualised version

of the Psychological Theory of Generics as a theory of how heuristics affect our judgements of the truth values of generics and discuss the application of this reconceptualised theory to some of the puzzles posed by generics, including their apparent content-sensitivity, inferential asymmetry and association with stereotyping and prejudice.

2

COLLECTIVE AND DISTRIBUTIVE READINGS OF GENERICS

2.1 INTRODUCTION

One of the central points of contention regarding the meaning of bare plural generics concerns the semantic contribution that bare plural noun phrases make to the meaning of the sentence. Do they refer to just one object, many objects at once, or introduce variables which may be bound by an operator, such as the covert 'GEN'? According to *referential* analyses of generics, bare plural noun phrases in generics refer to some object or objects. Typically, the reference of a bare plural is taken to be a kind. On the simplest version of the kind-predication view, in a generic, ⌜Fs are G⌝, the bare plural ⌜Fs⌝ refers to the kind F and the generic is assigned the logical form: ⌜G(F)⌝ (Liebesman, 2011). For example:[1]

(1) a) Chickens are widespread.
    b) widespread(CHICKEN)

(2) a) Dodos are extinct.
    b) extinct(DODO)

(3) a) Civil servants work hard.
    b) work.hard(CIVIL.SERVANT)

(4) a) Doctors are intelligent.
    b) intelligent(DOCTOR)

1 Throughout, I use CAPITALS for terms denoting kinds in the metalanguage.

I call referential theories which assign bare plurals a referent which is an individual (e.g. Carlson, 1977, 1982), *singularist*; those which assign bare plurals a referent which is a plurality, I call *pluralist* (e.g. Koslicki, 1999; Moltmann, 2013).[2]

According to the orthodox view of generics, *characterising sentences*, which are sentences which generalise across (normal or typical) members of a kind, such as (3a) and (4a), have a different logical form to *kind-predication sentences*, which predicate a property of a kind as a whole, such as (1a) and (2a). On this view, characterising sentences have a tripartite logical form, which includes a phonologically null operator, 'GEN', which has a meaning similar to 'typically' or 'generally' and is translated into the metalanguage as *gen*, with material divided between a restrictor and scope:[3]

(5) $gen(x) : \underbrace{\text{civil.servant}(x)}_{\text{restrictor}} \underbrace{[\text{work.hard}(x)]}_{\text{scope}}$

(6) $gen(x) : \text{doctor}(x)[\text{intelligent}(x)]$

On this view, the semantic contribution of the bare plural is a predicate with a free variable, which is bound by 'GEN'. Call theories which invoke a relational operator in the logical form of all characterising sentences, *relational* theories (e.g. Krifka et al., 1995; Pelletier & Asher, 1997; Eckardt, 1999; Greenberg, 2003; Leslie, 2008).

The difference between referential and relational theories comes in the analysis of the logical form of characterising sentences, although strictly speaking, the two types of theory are compatible. Some relational theories are referential: they assign bare plurals a referential semantics *and* posit a phonologically null relational operator in all characterising sentences (e.g. Chierchia, 1998; Nickel, 2016; Cohen, 2020; Guerrini, 2024). On (a simplified version of) Chierchia's theory (1998), the logical form of (4a) will be as follows, where $\leqslant$ is mereological parthood relation (I discuss the generation of this structure in more detail below).

(7) a) $gen(x) : x \leqslant \textsc{doctor}[\text{intelligent}(x)]$

2 Chierchia (1998), Dayal (2004) and Guerrini (2024) propose referential theories according to which bare plurals denote individual concepts: functions from possible worlds to mereological sums. I class these theories as pluralist, although strictly speaking the semantic value of a bare plural is not a plurality, but a function whose range contains pluralities.

3 For discussion of this tripartite structure, see Lewis (1975b) and Heim (1982).

b) It is generally true of members of the kind DOCTOR that they are intelligent.

I will refer to relational theories that are referential as *D-relational* (i.e. definite-relational) theories and those which are not referential as *I-relational* (i.e. indefinite-relational) theories.

In this chapter, I will motivate and expound the hypothesis that bare plurals refer plurally. I present data demonstrating that generic bare plurals exhibit both collective and distributive readings in the same way as other non-generic plural expressions. I show that neither singularist referential theories nor I-relational theories predict the full range of collective and distributive readings that generics exhibit. I also show that, while D-relational theories are better placed to explain the data presented here, they make implausible assumptions about the semantic type of collective predicates in generics.

Next, I set out a framework for regimenting plurals using a plural metalanguage and my assumptions about the nature of distributive readings. I show that the hypothesis that bare plurals refer plurally predicts the desired readings, provided that the account is supplemented with a quantificational distributivity operator, which, unlike 'GEN', is motivated independently of the analysis of generics. This distributivity operator can play the variable binding function that relational theorists usually attribute to 'GEN'. Finally, I consider an objection to the analysis based on the observation that bare plurals are not rigid designators, concluding that this constitutes evidence that bare plurals are not plural names, but refer in the same way as plural definite descriptions.

2.2 THE DATA AND EXISTING THEORIES

2.2.1 *The data*

The linguistic data I consider here concerns the availability of collective and distributive readings of bare plural generics.[4] It is well-known that generics exhibit scope ambiguities with quantifiers or connectives (Schubert & Pelletier, 1987):

(8) Conservatives have a favourite politician.

a) There is a politician who is the favourite politician of conservatives in general (exactly one politician).

b) In general, for each conservative there is a politician who is their favourite politician (one politician per conservative).

(9) Computers are PCs or Macs.

a) Either all computers are PCs or all computers are Macs (mixed case excluded).

b) In general, each computer is either a PC or a Mac (mixed case possible).

However, generics also exhibit ambiguities which are very similar in nature to collective-distributive ambiguities for non-generic plural sentences. In the following examples, I offer a paraphrase of (a) the collective reading and (b) the distributive reading.[5]

(10) Scientologists obey their leader.

a) In general, scientologists (collectively) obey the leader of scientologists (maximum one leader).

b) In general, each scientologist$_1$ obeys their$_1$ leader (one leader per scientologist).

4 Although I don't discuss mass nouns here, what I say about bare plurals carries over to them too. For the purposes of this discussion, I will set aside existentially interpreted bare plurals, as in 'Dogs are chasing their tails'.

5 Subscripts are used for co-indexing pronouns. I use 'in general' in the paraphrases to allow for generics to admit exceptions. See §2.4 for discussion of this point in relation to the definition of distributivity.

(11) Philosophy students consider themselves clever.

a) Philosophy students generally consider philosophy students to be clever.

b) In general, each philosophy student$_1$ considers that they$_1$ are clever.

I will argue here that all of the ambiguities exhibited by (8)-(11) are cases of a collective-distributive ambiguity. The distributive readings are characterised by the covariation of the value of a variable with each member of the plurality across which the predicate is distributed. In the collective readings, it is the plurality as a whole which is said to satisfy the predicate. For example, in the distributive reading of (8), the favourite politician may vary with each conservative, whereas in the collective reading there can be at most one politician who is favoured by conservatives in general.

Before moving on to my proposal, I will explore in some detail how existing theories of generics might explain the data in (8)-(11) and show that none of these theories can account for the full range of collective and distributive readings of generics.

2.2.2 *The binder argument*

I begin by briefly recapitulating a criticism of the simple kind-predication theory which Sterken calls the *binder argument* (Carlson, 1977; Leslie, 2015; Sterken, 2016). Consider the logical form of (11) according to the simple kind-predication theory:

(12) $[\lambda x.\text{consider.clever}(x)(x)](\textsc{philosophy.student})$

After β-reduction, this says that the kind PHILOSOPHY.STUDENT considers the kind PHILOSOPHY.STUDENT to be clever. That arguably captures the collective reading of (11). But, as Carlson argued (1977: 262-265), the simple kind-predication theory does not posit enough internal structure to capture the distributive reading, since there is no way to get the philosophy student who is considered clever to vary with each member of the kind PHILOSOPHY.STUDENT. That means that the covariation characteristic of the distributive reading cannot be adequately captured by the simple kind-predication theory.

2.2.3 *A predicate modifier*

In response to this problem, Carlson proposed a 'sophisticated' kind-predication theory. According to this theory, predicates of individuals are modified into predicates of kinds – following Teichman (2023), call this operator 'PM' ('G'' in Carlson, 1977). The two readings of (11) can then be represented as follows:

(13) a) $[\lambda x.\text{consider.clever}(x)(x)](\textsc{philosophy.student})$ *collective*

b) $[\lambda x.PM(\text{consider.clever}(x)(x))](\textsc{philosophy.student})$ *distributive*

The meaning of the modified predicate in (b) is (roughly) a property which holds of a kind iff enough of its members have the unmodified predicate under normal conditions. That gives us the distributive reading. I granted above that a structure like (13a) captures the collective reading according to the simple kind-predication theory. But does it capture this reading according to the sophisticated kind-predication theory? Not on the metaphysical picture painted by Carlson (1977) and not on the semantic one painted by Teichman (2023). First let's consider Carlson's metaphysics.

He argues that predicates which hold of individuals must be modified into predicates which hold of kinds, since predicates which hold of individuals cannot meaningfully be applied to kinds (Carlson, 1977: 266). If that is right, then (13a) should not be interpretable because it attributes an individual-level property to a kind. Therefore, it cannot represent the collective reading of (11).[6]

Turning to Teichman's semantics, he posits an atomic type for kinds, k, drawing a distinction in the semantics between properties of kinds (type $\langle k, \langle s, t\rangle\rangle$) and properties of individuals (type $\langle e, \langle s, t\rangle\rangle$). '$PM$' denotes a function of type $\langle\langle e, \langle s, t\rangle\rangle, \langle k, \langle s, t\rangle\rangle\rangle$, turning properties of individuals into properties of kinds. But if that is right, then either 'consider themselves' denotes a kind-level property, in which case (13b) is not interpretable because of a type-mismatch, or it denotes an individual-level property, in

6 Carlson considers some predicates that could hold of either kinds or individuals, such as 'popular' or 'famous' (Carlson, 1977: 306-7). He considers the stock of such predicates to be quite small, and it does not seem like 'consider themselves clever' could plausibly fall into this category, since it relates to the actions of individuals (considered collectively). Perhaps Carlson could respond by broadening the class of predicates which are ambiguous between a kind-level and an individual reading. Below, in response to Guerrini's D-relational theory, I present some evidence that such a move will not work.

which case (13a) is not. Since the unmodified predicate can express both an individual-level and a kind-level predicate, this analysis fails to capture at least one of the readings of (11).

2.2.4 *A scope phenomenon?*

I now turn to responses on behalf of the I-relational theory. A natural response to the data is to propose that the different readings are a matter of scope (Schubert & Pelletier, 1987). Focusing first on (8) and (9), the analysis on an I-relational theory would be as follows. The distributive readings in (15) arise when *gen* takes scope over the quantifier or disjunction in the predicate. The collective readings in (14), on the other hand, would correspond to the existential quantifier/disjunction taking wide scope over the operator.[7]

(14) *Collective*

a) $\exists x(gen(y) : \text{conservative}(y)[\text{favourite.politician.of}(x)(y) \wedge \text{has}(y)(x)])$

b) $gen(x) : \text{computer}(x)[\text{PC}(x)] \vee gen(x) : \text{computer}(x)[\text{Mac}(x)]$

(15) *Distributive*

a) $gen(x) : \text{conservative}(x)[\exists y(\text{favourite.politician.of}(y)(x) \wedge \text{has}(x)(y))]$

b) $gen(x) : \text{computer}(x)[\text{PC}(x) \vee \text{Mac}(x)]$

This explanation works nicely for (8) and (9) and has been cited as evidence in favour of the presence of 'GEN' in characterising sentences (Sterken, 2016: 510).

2.2.4.1 *A missing binder*

However, although Leslie (2015) and Sterken (2016) cite cases like (10) and (11) as evidence in *favour* of a relational theory, their collective readings pose a problem, at least for I-relational variants.

7 I will ignore the syntactic question of how these structures are generated and simply grant the relational theorist that they are available for the sake of the argument. If they cannot be generated, then so much the worse for the relational theory.

Let's examine the truth-conditions of the two readings in more detail. A situation which would make the collective reading in (10a) false but the distributive reading in (10b) true is one where each scientologist is assigned a personal leader whom they obey, but there is no overall leader whom they collectively obey.[8] The problem for the I-relational theory is that there is no quantifier (or other operator) for *gen* to take scope over. That means we can't put this reading down to a variation in scope. So while the distributive reading of (10) can be analysed as in (16b), there is no other assignment of scope which generates the collective reading.

(16) a) $?$: scientologist(x)[obey.leader.of$(x)(x)$]
b) $gen(x)$: scientologist(x)[obey.leader.of$(x)(x)$]

On this analysis of the distributive reading, the pronoun 'their' must be bound, since its value may vary for each scientologist quantified over. On the collective reading, we want the pronoun to refer back not to each scientologist, but to scientologists in general. One way for the relational theorist to explain the collective reading would be for the pronoun to be left free, with its value determined by the variable assignment.

(17) a) Scientologists obey their$_3$ leader.
b) $gen(x)$: scientologist(x)[obey.leader.of(x)(SCIENTOLOGIST)]

To get the desired truth-conditions, we need a variable assignment of the following form $[3 \mapsto$ SCIENTOLOGIST$]$ (or perhaps one which maps 3 to the plurality of all scientologists).

2.2.4.2 *Reflexive collective readings*

On that analysis, the 'collective' reading is just a special kind of distributive reading, where the leader in question just happens to be the leader of scientologists and so there is at most one leader. We might grant the I-relational theorist this analysis of (10), but it is less plausible when the pronoun is reflexive, as it is in (11), whose

8 Carlson (1977: 262-263) suggests that these collective readings are not available, but he considers only a small range of examples. Liebesman and Magidor (2025: 287-290) clearly establish the availability of the collective reading.

two readings would, after assigning values to free variables, need to be analysed as follows:

(18) a) $gen(x)$: philosophy.student(x)[consider.clever(x)(PHILOSOPHY.STUDENT)]
b) $gen(x)$: philosophy.student(x)[consider.clever$(x)(x)$]

The collective reading in (a) is generated in the same way as (17b). That requires that the reflexive pronoun is interpreted as a free variable at logical form, whose value is given by the variable assignment.

This analysis runs into trouble when consider cases of ellipsis. I assume here Heim and Kratzer's LF IDENTITY CONDITION ON ELLIPSIS (Heim & Kratzer, 1998: 250; see also Sag, 1976; Williams, 1977; Fiengo & May, 1994; Tanaka, 2011):

LF IDENTITY CONDITION ON ELLIPSIS A constituent may be deleted Phonological Form only if it is a copy of another constituent at Logical Form.

Because the analysis under consideration distinguishes the collective reading from the distributive reading in terms of whether the pronoun is bound or free, it predicts that there are only two possible readings for (19), depending on the index on the pronoun which is copied and deleted (see Heim & Kratzer, 1998: 254-257; Keenan, 1971):

(19) Philosophy students consider themselves clever and linguistics students do too.

(20) SLOPPY DISTRIBUTIVE

a) In general, each philosophy student$_1$ considers themselves$_1$ clever and each linguistics student$_1$ considers themselves$_1$ clever.

b) [[philosophy students] [GEN [λ_1 t$_1$ consider themselves$_1$ clever]] and [[linguistics students] [GEN [λ_1 t$_1$ ~~consider themselves$_1$ clever~~]]]]

(21) STRICT COLLECTIVE

a) Philosophy students consider philosophy students to be clever, and linguistics students consider philosophy students to be clever.

b) [[[philosophy students] [GEN [λ_1 t_1 consider themselves$_2$ clever]]] and

[[linguistics students] [GEN [λ_1 t_1 ~~consider themselves$_2$ clever~~]]]]

Where $g(2) =$ PHILOSOPHY.STUDENT, (21b) captures the strict collective reading, i.e. where linguistics students consider philosophy students to be clever (collectively). However, there is at least one more possible reading of (19):

(22) SLOPPY COLLECTIVE

Philosophy students consider philosophy students to be clever, and linguistics students consider linguistics students to be clever.

The problem for the I-relational theory is as follows: if the collective reading of (11) arises only as the result of a referential occurrence of the reflexive pronoun, then the I-relational theory fails to predict the sloppy collective reading paraphrased in (22). The only way of doing so would be to maintain that the elided VP is not syntactically identical to its antecedent (and so violate the LF IDENTITY CONDITION ON ELLIPSIS):

(23) a) [philosophy students] [GEN [λ_1 t_1 consider themselves$_2$ clever]]

b) [linguistics students] [GEN [λ_1 t_1 ~~consider themselves$_3$ clever~~]]

With a suitable variable assignment, $[2 \mapsto$ PHILOSOPHY.STUDENT, $3 \mapsto$ LINGUISTICS.STUDENT$]$, that generates the following truth-conditions:

(24) $gen(x)$: philosophy-student(x)[consider.clever(x)(PHILOSOPHY.STUDENT)] $\wedge$
$gen(x)$: linguistics.student(x)[consider.clever(x)(LINGUISTICS.STUDENT)]

Assuming this adequately captures the truth-conditions of the sloppy collective reading of (19), the question is whether this is method of generating the truth-conditions is plausible. Specifically, is it plausible for the elided VP to contain a free reflexive pronoun whose index is not identical to any antecedent?

There is good evidence that sloppy readings of pronouns can arise even when the pronoun cannot be syntactically (and hence semantically) bound (e.g. because binding would violate island constraints on movement), which suggests that the strict-sloppy

ambiguity cannot straightforwardly be explained by an antecedent pronoun's being free or bound, respectively (Elbourne, 2008). There is also empirical evidence that strict readings are available to speakers even for reflexive pronouns, which is not predicted by the hypothesis that reflexives must be semantically bound (Frazier & Clifton, 2006; Kim & Runner, 2009; Ong & Brasoveanu, 2014; McKillen, 2016). This might provide some motivation for the suggested analysis of (11).

However, on this analysis of the sloppy collective reading of (19), there are two reflexive pronouns, neither of which is semantically bound or is co-referential with any antecedent. That's because the only antecedents are quantificational, so co-reference is not possible. In the evidence just cited suggesting that reflexive pronouns need not be semantically bound, the cases where strict readings were available were ones where the reflexive pronoun nonetheless co-referred with some antecedent. Yet what the I-relational theorist needs to rescue their account is much stronger and less plausible: reflexive pronouns may be referential and need not co-refer with any antecedent (not to mention the violation of the LF IDENTITY CONDITION ON ELLIPSIS). I can see no evidence for that claim, so I conclude that this attempt to explain the collective sloppy reading of (19) fails.

2.2.4.3 *Genuine kind-predication*

Another option open to the I-relational theorist to capture the collective reading of (10) is to say that it is an instance of genuine kind-predication, as in the canonical examples (1a) and (2a). This is consistent with their view of characterising sentences, since they view at least some bare plurals as referring to kinds. (10) is then analysed:

(25) obey.leader.of(SCIENTOLOGIST)(SCIENTOLOGIST)

Even if we grant that this adequately captures the truth-conditions of the collective reading of (10), the analysis can't account for cases of co-predication, since (26) has a reading where the first predicate attributes the property of being diligent to individual scientologists while the second has the collective reading stated in (10a).

(26) Scientologists are diligent and obey their leader.

According to the I-relational theorist, the expression 'scientologists' would need to make two contributions to the sentence, one referring to a kind and one as a predicate contributing variables to be bound by *gen* (this kind of proposal is sketched by Leslie, 2015: 46). But there seems to be no linguistic motivation for the claim that this expression makes two distinct contributions to the interpretation of the sentence, so this solution is at best stipulative (see Collins, 2018: 55-58 for an extended discussion of the problems with this approach to generic co-predication).

This concludes my argument that the I-relational theory does not possess a general strategy for deriving the collective readings of generics.

2.2.5 *D-relational theories*

An obvious fix for the relational theory is to assign all bare plural arguments a referential semantics, which avoids the problems of the missing binder and co-predication. In Guerrini's recent D-relational theory (2024), which follows Chierchia (1998), generic bare plurals begin by denoting properties and come to be referential by means of a type-shifting mechanism, $^{\cap}$ (NOM). NOM is an operation on properties, of type $\langle s, \langle e, t \rangle \rangle$, so that:

$$^{\cap}[\lambda w_s.\lambda x_e.P_w(x)] = \begin{cases} \lambda w_s.\iota P_w \text{ if there are Ps in } w \\ \text{undefined otherwise} \end{cases}$$

$\ulcorner \iota P_w \urcorner$ is a term which refers to the mereological sum of things which are P in w. Following Chierchia, I will refer to theories which assign bare plurals (after type-shifting) a referent of type $\langle s, e \rangle$, as *neo-Carlsonian*. On Chierchia and Guerrini's view, *gen* quantifies over both individuals and situations, resulting in the following structure for (4a):[9]

(27) $gen(x, s) : x \leqslant {}^{\cap}\text{doctors}_s \wedge C(x, s)[\text{intelligent}_s(x)]$

9 $C(x, s)$ means that x is contextually salient in s.

On this view, kinds just are type-shifted properties (Chierchia, 1998: 350) and kind-level predicates take type-shifted properties as their argument; (2a) is analysed as:

(28) $\text{extinct}_s(\cap\text{dodos})$

The distributive and collective readings of (8)-(10) can then be derived as follows:

(29) COLLECTIVE

a) $\exists x(\text{favourite.politician.of}_s(x)(\cap\text{conservatives}) \wedge \text{has}_s(\cap\text{conservatives})(x))$

b) $gen(x,s) : x \leqslant \cap\text{computers}_s \wedge C(x,s)[\text{PC}_s(x)] \vee$
$gen(x,s) : x \leqslant \cap\text{computers}_s \wedge C(x,s)[\text{Mac}_s(x)]$

c) $\text{obey.leader.of}_s(\cap\text{scientologists})(\cap\text{scientologists})$

(30) DISTRIBUTIVE

a) $gen(x,s) : x \leqslant \cap\text{conservatives}_s \wedge C(x,s)[\exists y(\text{favourite.politician.of}_s(y)(x) \wedge \text{has}_s(x)(y))]$

b) $gen(x) : x \leqslant \cap\text{computers}_s \wedge C(x,s)[\text{PC}_s(x) \vee \text{Mac}_s(x)]$

c) $gen(x) : x \leqslant \cap\text{scientologists}_s \wedge C(x,s)[\text{obey.leader.of}_s(x)(x)]$

Finally, consider the two possible structures for (11):

(31) COLLECTIVE

a) [[philosophy students] [λ_1 t_1 consider themselves$_1$ clever]]

b) $\text{consider.clever}(\cap\text{philosophy.students})(\cap\text{philosophy.students})$

(32) DISTRIBUTIVE

a) [[philosophy students] [GEN [λ_1 t_1 consider themselves$_1$ clever]]]

b) $gen(x,s) : x \leqslant \cap\text{philosophy.students}_s \wedge C(x,s)[\text{consider.clever}_s(x)(x)]$

In both structures, the pronoun is bound. What differs is that in the distributive reading, the pronoun is bound by *gen*, so its value covaries with members of the kind. This means that the sloppy collective reading of (19), which posed a problem for I-relational theories, can be accounted for, since the elided VP contains a bound variable:

(33) SLOPPY COLLECTIVE

[[[philosophy students] [λ_1 t_1 consider themselves$_1$ clever]] and

[[linguistics students] [λ_1 t_1 ~~consider themselves$_1$ clever~~]]]

However, the neo-Carlsonian proposal faces a more general problem, relating to the semantic type of the predicate in collective readings. This analysis requires that many predicates are type-ambiguous between a kind-level and an individual-level collective reading. For example, 'consider themselves clever' would need to have a reading where its argument is a kind, type $\langle s, e \rangle$, and a reading where its argument is an individual, type e. This second reading is what gets modified by 'GEN' into a kind-level predicate, which is true of a kind iff it is generally true of its members (cf. the lexical entry for 'GEN' in Guerrini, 2024: 9).

Now consider a variant on (19) and a paraphrase of its collective reading:

(34) a) Philosophy students consider themselves clever and these linguistics students do too.

b) In general, philosophy students consider philosophy students clever and these linguistics students consider these linguistics students clever.

According to the neo-Carlsonian view, the collective reading of the first conjunct of (34a) is derived as a case of kind-level predication (annotated below with subscript K on the verb). But the second conjunct cannot be a case of kind-level predication, since 'these linguistics students' denotes a plurality, not a kind. That should lead to a type-mismatch in the second conjunct in (35b), but the collective reading of the elided VP in (21b) is available.

(35) a) [philosophy students] [λ_1 t_1 consider$_K$ themselves$_1$ clever]

b) [these [linguistics students]] [λ_1 t_1 ~~consider~~$_K$ ~~themselves~~$_T$ ~~clever~~]

This is a general challenge for any theory which assigns kinds and pluralities distinct types, since the data presented here suggests that, if the collective readings of generics are to be accounted for as cases of kind-predication, there will be many predicates which are ambiguous between a kind-level and a collective reading. The acceptability of sentences like (34a) is evidence against this view.

There do not seem to be any linguistic differences between individual, collective and kind-level readings of predicates in terms of the semantic type of their arguments. That is why any of the four answers to this questions are acceptable (cf. Link, 1984):

(36) Q: Who considers themselves clever?
A1: I do.
A2: These linguistics students do.
A3: Each of the students does.
A4: Philosophy students do.

There is no sense of ambiguity in giving any of the four answers to the question and they can be conjoined, which suggests that the question itself is not ambiguous.

Finally, if predicates really were type-ambiguous between an individual and a kind-level reading (the latter being available when an individual-level predicate is modified by 'GEN'), then the we would expect to see type-clashes in the following cases (these tests are suggested by Carlson, 1977: 163):

(37) a) Philosophy students are cleverer than John.
b) Because cats clean themselves, my son thinks that he should too.
c) These men are more devout than Scientologists (are).
d) Jack and Jill run like ducks (do).

On the neo-Carlsonian view, each predicate as it applies to a bare plural argument must take arguments of type $\langle s, e \rangle$, but, as it applies to individuals or pluralities, take arguments of type e. That means that these sentences should be ill-formed because of

a type-clash. But they are acceptable, so we should conclude that the predicates are not type-ambiguous in the way required by the neo-Carlsonian view.

This suggests that any approach which posits a type distinction between the arguments of kind-level predicates and other plural predicates is misguided. I conclude that the neo-Carlsonian theory proposed by Chierchia and Guerrini fails to provide a satisfactory analysis of the data in (8)-(11) because it treats the collectively interpreted predicates as kind-level predicates. There is good linguistic evidence that the predicates are not ambiguous between a kind-level and collective reading.

2.2.6 *Summing up*

To conclude this section, I will briefly summarise the argument so far. Many bare plural generics admit both a collective and a distributive reading. I considered a range of methods of deriving these readings on behalf of existing theories, beginning with kind-predication theories. I found that simple kind-predication theories were unable to derive the distributive reading. Sophisticated kind-predication theories can derive the distributive reading, but are unable to derive the collective readings, since they don't allow kinds to bear the same properties as individuals.

One natural way of representing this ambiguity is by means of a scope ambiguity, which can in some cases be attributed to a scope interaction between *gen* and another operator or connective. However, this analysis failed as a general strategy for deriving the desired collective readings of sentences with bound pronouns. I then considered a D-relational theory, which did not suffer from the same problem, but which required implausible assumptions about the semantic type of the predicates which result in the collective readings.

I can now sketch my own proposal. I will take generic bare plurals to be referring expressions but, unlike kind-predication theorists, I take them to refer plurally. There are two good motivations for this view:

(i) The semantic type of kind-level and plural collective predicates will be the same, avoiding the problems that arose for the kind-predication and D-relational theories described above.

(ii) Bare plurals pattern with other plural referring expressions with respect to the availability of collective and distributive readings. They are in this regard unlike collective or group noun phrases which refer to singular individuals and do not exhibit distributive readings in the same environments.

I replace the predicate modifier of the sophisticated kind-predication theory with a quantificational distributivity operator, '$\mathcal{D}$', used to account for the meaning of plurals (Link, 1987; Roberts, 1987; Lasersohn, 1995; Winter, 2001). This approach explains the availability of all the readings of (8)-(11) and does so without introducing any novel operators (like 'GEN' or '*PM*') or objects (kinds).

2.3 REGIMENTING PLURALS

To properly state my theory, I begin by setting out the details of the metalanguage I use to regiment plural expressions. I will adopt primitive plural resources to regiment natural language plurals. The language I use is a version of PFO+, an extension of first-order logic with plural variables and constants, plural quantifiers and collective predicates (Rayo, 2002; Oliver & Smiley, 2016).[10] To PFO+, I add lambda abstraction over variables of any type.

2.3.1 *A plural metalanguage*

Plural variables and constants are represented typographically with a double-letter, e.g. $xx, yy, zz \ldots$; $aa, bb, cc \ldots$. Unlike singular terms, $x, y, z \ldots$; $a, b, c \ldots$, these terms

10 For discussion of linguistic evidence in favour of plural reference, see Moltmann (2016). I remain agnostic here on the semantics of the language of regimentation. See Florio and Linnebo (2021: Chapter 7) for discussion of the differences between set-based and plurality- or super-plurality based model theories. For my purposes the choice won't make a difference. Friends of a mereological analysis of plurals, following Link (1983), might prefer to recast my discussion of singular objects in terms of mereological atoms.

may take *one or more* object as a value. $\preceq$ is a logical predicate which translates 'is or are among' and relates two terms such that $xx \preceq yy$ iff there are some things among xx and everything which is among xx is among yy. Identity is a special case of inclusion: $xx = yy \leftrightarrow_{df} (xx \preceq yy \wedge yy \preceq xx)$.

I introduce a term connective, $+$, which forms plural terms from two terms. This will translate the English 'and' as it appears between two names. For example, '$a + b$' is a compound term which refers to a and b. A plural quantifier, $\exists$, is interpreted as meaning 'there are some things ... such that ...'. The plural universal quantifier, $\forall$, is defined as the dual of $\exists$. Singular objects, those which can be the value of a singular term, are then defined: $\textsc{sing}(xx) \leftrightarrow_{df} \forall yy(yy \preceq xx \rightarrow xx = yy)$.

Predicates of English are translated into predicates of PFO+. Singular predicates have only singular objects in their extension, e.g. $\lambda x_e.\text{boy}(x)$. Translations of English count nouns can be pluralised by the star operator, *, which denotes the closure operation under $+$ (see Link, 1983). That is, something is in the extension of 'boy' iff it is a boy and some thing(s) is/are in the extension of 'boys' iff it is/they are a plurality all of whose members are members of the extension of 'boy'. For any two objects, a and b, if $P(a)$ and $P(b)$ then $^*P(a + b)$.

Other English predicates, those which have collective readings, are 'born plural', so to speak, and are not derived from singular predicates. For example, 'lifted a box' may have pluralities in its extension without everything among those pluralities being in its extension and is translated: $\lambda xx_e.\exists y(\text{box}(y) \wedge \text{lifted}(xx)(y))$.[11]

I translate definite descriptions as plural terms referring to the maximum plurality which satisfies a predicate (Sharvy, 1980; Link, 1983), which is defined as follows:

(38) $^{\textsc{max}}P =_{df} \lambda xx_e.P(xx) \wedge \forall yy(P(yy) \rightarrow yy \preceq xx)$

Since there can be at most one maximum plurality which satisfies a predicate, definite descriptions presuppose only that such a plurality exists (an additional uniqueness condition is redundant).

11 Talk of 'pluralities' should be viewed only as a way of speaking about objects and need not be seen as entailing an ontological commitment to plural objects. Strictly speaking, singular objects are pluralities (ones with just one member).

(39) $[\![\text{the}]\!] = \lambda P_{\langle e,t\rangle} : \exists xx^{\text{MAX}} P(xx).\ \iota xx^{\text{MAX}} P(xx)$

$\ulcorner \iota xx P(xx) \urcorner$ is a term which, if defined, refers to the things which together exclusively satisfy P.

2.3.2 *Group noun phrases*

Now we turn to the semantics of group noun phrases (NPs), like 'group', 'family', 'committee', 'herd', 'crowd', 'class', 'team' or 'choir'. These are counts nouns which are characterised by their compatibility with *of*-phrases, where they combine with a plural complement, as in 'group of children', 'family of cats', 'committee of judges' (Barker, 1992). Group NPs are relevant to the analysis here because, on a singularist analysis of bare plurals, bare plural NPs should behave like group NPs (Schwarzschild, 1996: 169).[12]

An initially attractive thought is that the extensions of group NPs should consist of the pluralities of the members of groups, since even singular group NPs can appear with collective predicates, such as 'assembled' or 'dispersed', and so seem to be semantically plural:

(40) a) The committee assembled.
b) The crowd dispersed.

But groups cannot be identical to their members. Take Committee C, which has two members, Annie and Bella. If Anna and Bella are both in their 80s, then Anna and Bella have the property of being old. But if Committee C was only formed last week, it does not have the property of being old. Moreover, if Anna and Bella are the only two members of Committee D, which has been in existence for decades, then that committee is old, even though it has all and only the same members as the new committee. So Committee C and Committee D are not identical, even though they have all and only the same members (see Uzquiano, 2004: 141-142). Link, crediting

12 Carlson (1977) is explicitly committed to this view. Liebesman (2011) claims to be neutral on the ontology of kinds (see in particular n. 8), but his discussion of property inheritance (419-420) suggests that this is what he has in mind.

Kaplan, observes that this problem cannot be repaired by assuming that the identity condition on groups requires that they have the same members at all possible points of evaluation, since we can imagine two committees which necessarily have the same members, yet are distinct (e.g. because they have different responsibilities) (Link, 1983: 304).

The differences between group NPs and genuine plural terms also show up in the following contrasts Moltmann (1997: 62):[13]

(41) a) He compared the family members.
b) #He compared the family.

(42) a) The faculty members are distinguishable.
b) #The faculty is distinguishable.

(43) a) Between the chairs, there was a snake.
b) #Between the group of chairs, there was a snake.

Substituting the group NP phrase for what would, on the hypothesis that a group NP phrase denotes the plurality of its members, be a co-referring NP results in unacceptability. Moltmann also observes that the same kind of substitution in other contexts results in meaning change:

(44) a) He classified the people.
b) He classified the group of people.

(45) a) He recognized the family members.
b) He recognized the family.

In each case, the result of the substitution is a sentence which differs in its truth-conditions, suggesting that the two NPs do not refer to the same thing(s).

Group NPs admit quantification over their members with 'part', whereas plural NPs do not (Schwarzschild, 1996):

(46) a) Part of the group arrived early.

13 # indicates a judgment of unacceptability.

b) #Part of the boys arrived early.

(47) a) Part of the herd was hungry.

b) #Part of the members of the herd was angry.

They also behave differently from pluralities when they combine with reciprocal predicates:

(48) a) #The pile is touching each other.

b) The rocks in the pile are touching each other.

(49) a) ?The family lives in different cities.[14]

b) The members of the family live in different cities.

Finally, and most relevant for the purposes of this paper, sentences with plural NPs exhibit a wider range of distributive readings than sentences with group NPs (Moltmann, 1997; de Vries, 2017). Definite plurals are acceptable with both a collective and a distributive reading:

(50) The cows have a favourite kind of grass.

a) There is a kind of grass which is the favourite of the cows (exactly one kind of grass).

b) For each cow there is a kind of grass which is their favourite (one kind of grass per cow).

(51) The children are sleeping or playing.

a) Either the children are sleeping or the children are playing (mixed case excluded).

b) Each child is either sleeping or playing (mixed case possible).

(52) The soldiers obey their leader.

a) The soldiers obey the leader of the soldiers (maximum one leader).

14 ? indicates a judgment of borderline acceptability. This sentence has a true reading in which the family has homes in multiple cities, e.g. they split their time between London and Birmingham.

b) Each soldier$_1$ obeys their$_1$ leader (one leader per soldier).

In contrast, only a collective reading is available for sentences with group NPs, even though the distributive reading, paraphrased in (b), is perfectly plausible and hence not ruled out on pragmatic grounds:

(53) The herd of cows has a favourite kind of grass.

a) There is a kind of grass which is the favourite of the herd of cows (exactly one kind of grass).

b) ✗ For each cow in the herd there is a kind of grass which is their favourite (one kind of grass per cow).

(54) The group of children is sleeping or playing.

a) Either the children are sleeping or the children are playing (mixed case excluded).

b) ✗ Each child is either sleeping or playing (mixed case possible).

(55) The battalion obeys its leader.

a) The members of the battalion obey the leader of the battalion (maximum one leader).

b) ✗ Each member$_1$ of the battalion obeys their$_1$ leader (one leader per member).

This evidence gives us good reason to conclude that group NPs do not refer to pluralities. Instead, I will assume that group NPs refer to singular individuals, i.e. for all NP, if NP is a (singular) group NP, then for all $xx \in [\![NP]\!]$, $\textsc{sing}(xx)$. The relation between groups and their members remains a question of metaphysics, not logic.[15]

15 One caveat is that unlike in American English, in British English, group NPs may take either singular or plural verb agreement. Plural agreement foregrounds a 'member-oriented' reading (Barker, 1992; Sauerland & Elbourne, 2002). Compare:

(i) The committee is old.

(ii) The committee are old.

The group-oriented reading is preferred for (i), while a member-oriented reading is preferred for (ii), which can be paraphrased: the members of the committee are old. Ritchie (2017) suggests that this is

2.4 TWO TYPES OF DISTRIBUTIVITY

I will now explain in more detail what I take distributivity to be and give an account of how distributivity arises which explains the contrast in the availability of collective and distributive readings for plural and group NPs.

2.4.1 *An operational definition of distributivity*

One challenge in defining distributivity is the availability of so-called 'non-maximal' readings of distributive predicates (Gillon, 1987; Brisson, 2003; Križ, 2016). These are cases where an apparently distributive predicate is true of a plurality but is not true of every singular object among the plurality. For example, imagine that there are twelve books on the reading list and John reads eleven of them. It seems that this can be accurately reported by (56a), without entailing (56b).

(56) a) John read the books. $\nRightarrow$
b) John read each of the books.

So we can't simply say that a predicate is distributive with respect to an argument place when it is true of an argument iff it is true of every singular object among the argument, because that would automatically rule out cases like this as distributive. For this reason, I will rely on the following rough-and-ready definition of distributivity:

DISTRIBUTIVITY A predicate P is distributive with respect to an argument xx when
⌜xx is/are P⌝ $\Leftrightarrow$ ⌜Each member of xx is P (modulo exceptions)⌝.

Three comments:

(i) '$\Leftrightarrow$' means 'mutually entails'.

evidence that in British but not American English, group NPs are polysemous, with a group-denoting reading and plurality-denoting reading. One problem with this suggestion is that it does not explain the contrasts in (41)-(47), where the plurality-denoting reading should be available, even in British English. In what follows, I will set aside this phenomenon, on the assumption that the different dialects do not differ in their ontological commitments.

(ii) 'modulo exceptions' is added to deal with non-maximal readings of sentences like (56).

(iii) 'Member' is meant more inclusively than the inclusion relation of PFO+ ($\leq$). For instance, a pupil would be a member of a class, even though the class is a singular object and there are no pupils that are among the class in the strict logical sense.

Predicates which only have distributive readings are called 'distributive'. Predicates which never have distributive readings are called 'collective'. All other predicates are called 'mixed'.

The phenomenon of non-maximality won't play a major role in my analysis here, so I will set it aside here (although see Brisson, 2003; Malamud, 2012; Križ, 2016; Križ & Spector, 2021). However, I will make it clear below where I am assuming that non-maximality comes into the picture with respect to plural predication and distributivity.

2.4.2 *Distributivity and groups*

Using DISTRIBUTIVITY, to conclude that a predicate is not distributive with respect to an argument it is sufficient to show that the entailment fails in either direction. With this in mind, consider the following data from de Vries (2017):

(57) The class is hiding somewhere.

a) $\not\Leftarrow \forall x(\text{pupil}(x) \rightarrow \exists y(\text{hiding.in}(x)(y)))$

b) $\not\Rightarrow \exists x(\forall y(\text{pupil}(y) \rightarrow \text{hiding.in}(y)(x)))$

In (57), where the subject is a group NP, the distributive reading paraphrased in (57a) does not entail the the original sentence, hence the distributive reading of the predicate is not available. The collective reading, paraphrased in (57b), is available. In

contrast, in (58), both the collective and distributive reading are available when the subject is a plural count noun:

(58) The pupils are hiding somewhere.

a) $\Leftrightarrow \forall x(\text{pupil}(x) \rightarrow \exists y(\text{hiding.in}(x)(y)))$

b) $\Leftrightarrow \exists x(\forall y(\text{pupil}(y) \rightarrow \text{hiding.in}(y)(x)))$

We find the same pattern with disjunctive predicates and bound pronouns (cf. examples (53)-(55) and de Vries, 2017: 179-183), as well as reciprocals and other predicates which make reference to the members of a plurality (Moltmann, 1997: 62-65).

2.4.3 *The sources of distributivity*

Following Winter (1997, 2001) and de Vries (2017), the explanation for this contrast is that distributivity may arise from two sources. First: Q(uantificational)-distributivity. This type of distributivity is required to achieve the covarying reading of (58), where there is one hiding place for each pupil. This is achieved by the introduction of an additional quantifier, which can scope over 'somewhere'. This quantifier is introduced by means of a distributivity operator in the object language, which modifies a predicate P into a predicate which is true of a plurality xx iff P is true of each singular object among xx (Link, 1987; Roberts, 1987; Lasersohn, 1995). This is defined:[16]

(59) $[\![\mathcal{D}]\!] = \lambda P_{\langle e,t\rangle}.\lambda xx_e.\forall y(y \preceq xx \rightarrow P(y))$[17]

16 The operator as defined here applies only to one-place predicates. See Lasersohn (1998) for the generalisation to any argument place of a predicate.

17 We might also build in an additional restriction to the antecedent of the conditional which would allow us to account for non-maximal distributive readings (see e.g. Brisson, 2003). The focus here is on the availability of distributive readings of generics, not the nature of non-maximality in generics and other plural sentences, so I will largely ignore this issue.

The availability of the two readings of (58) can be explained by the presence or absence of the $\mathcal{D}$ operator as follows:

(60) COLLECTIVE

a) [[the pupils] are hiding somewhere]]

b) $\exists x(\text{hiding.in}(\iota yy.^{\text{MAX}*}\text{pupil}(yy))(x))$

(61) DISTRIBUTIVE

a) [[the pupils] [$\mathcal{D}$ [λ_1 t_1 are hiding somewhere]]]

b) $\forall y(y \preceq \iota xx.^{\text{MAX}*}\text{pupil}(xx) \rightarrow \exists z(\text{hiding.in}(y)(z)))$

Notice, however, that this does not generate an ambiguity when the definite description contains a singular group NP; on the hypothesis that group NPs denote singular objects, the following sentences are equivalent.

(62) 'COLLECTIVE'

a) [[the class] [is hiding somewhere]]

b) $\exists x(\text{hiding.in}(\iota yy.^{\text{MAX}}\text{class}(yy))(x))$

(63) 'DISTRIBUTIVE'

a) [[the class] [$\mathcal{D}$ [λ_1 t_1 is hiding somewhere]]]

b) $\forall y(y \preceq \iota xx.^{\text{MAX}}\text{class}(xx) \rightarrow \exists z(\text{hiding.in}(y)(z)))$

By the definition of a maximum, if the description refers, there must be just one class. Since groups denote singular objects, there is just one object among the class (i.e. the class itself). So both readings are true iff there is a place in which the class is hiding. This result is welcome: we don't want to predict a distributive reading here (cf. (57a)). But there are examples of predicates which do distribute over the members of the denotations of group NPs. For example:

(64) a) The class laughed.

b) The herd chewed the grass.

If the distributivity operator cannot access the things that are among these pluralities (since they are singular objects), how do the distributive readings arise? This brings us to our second type of distributivity: P(redicate)-distributivity. This is distributivity which stems from the lexical meaning of the predicate. These are special cases, in the sense that "[i]t is as if the properties common to the members of the [group] are extended to the group entity as if by courtesy. *In general this does not happen*... and the effect is enhanced by non-linguistic reasoning about the real world" (Barker, 1992: 90, emphasis added). The reasoning is non-linguistic because it requires knowledge of what sorts of entity can have the property in question. P-distributivity is therefore relatively limited.

Technically, the effects of P-distributivity can be captured by "pseudo-equivalences" (Winter & Scha, 2015), to the effect that the predicate holds of a collection when and only when it holds of its members (modulo exceptions).[18] For example:

(65) xx laughed $\leftrightsquigarrow \forall y(y$ is a member of $xx \rightarrow y$ laughed)

$\ulcorner \phi \leftrightsquigarrow \psi \urcorner$ is read: ϕ is pseudo-equivalent to ψ, i.e. ϕ is lexically reinterpreted to mean that ψ. This reinterpretation has the status of a constraint on the acceptable models for the language. So the pseudo-equivalence in (65) says that all acceptable models of the language are such that if 'laughed' is true of some argument, it is true of all the members of its argument. Since pseudo-equivalences are constraints on admissible models for a language, they do not generate scope interactions and covarying readings in the way that Q-distributivity, arising from a quantificational operator in the object language, does.

Finally, because P-distributivity arises as a result of the lexical meaning of a predicate, we can conclude that any distributive reading of a mixed predicate is a case of Q-distributivity, since, by definition, mixed predicates have meanings that are compatible with a non-distributive construal.

18 As before, 'member' must be interpreted as more inclusive than the logical inclusion relation $\leq$, since groups are singular objects, yet they exhibit P-distributivity.

2.5 THE PLURALIST ANALYSIS

I now return to the analysis of bare plurals and show that unlike singularist referential theories and relational theories, the pluralist account of bare plurals is consistent with the linguistic data introduced above.

2.5.1 *Predictions of a singularist theory*

I begin by ruling out a theory according to which kinds are represented as singular objects in the domain, i.e. objects xx for which $\forall yy(yy \leq xx \rightarrow xx = yy)$. The relevant predictions of this view are that generic bare plurals would pattern with group NPs in admitting P-distributive but failing to admit Q-distributive readings.

The first prediction is borne out, since the following sentences with P-distributive predicates each has a distributive reading which is true iff in general individual birds fly/dogs bark/bottlenose dolphins smile.

(66) a) Birds fly.
b) Dogs bark.
c) Bottlenose dolphins smile.

But the second prediction – that bare plurals do not exhibit Q-distributive readings – is not. Let's revisit (8), repeated below, along with its collective reading (a) and its distributive reading (b).

(67) Conservatives have a favourite politician. = (8)

a) There is a politician who is the favourite politician of conservatives in general (exactly one politician).

b) In general, for each conservative there is a politician who is their favourite politician (one politician per conservative).

The distributive reading of (67) cannot be put down to P-distributivity, since the predicate is clearly acceptable on a collective reading and the distributive reading

involves covariation between individual conservatives and their favourite politician. On the other hand, if 'conservatives' referred to a group, no Q-distributive reading should be available, since, as I have argued, Q-distributive readings are not available when the argument of the predicate denotes a group. We therefore have good evidence that the referent of 'conservatives' is a plurality, not a singular object. The same reasoning applies to the collective readings of sentences (9)-(11), which are all cases of Q-distributivity.

2.5.2 *Deriving both readings on a pluralist theory*

I will now show how a pluralist theory derives both readings. On the pluralist analysis, the two readings are derived using the distributivity operator as follows:[19]

(68) COLLECTIVE

a) [conservatives [λ_1 t_1 have a favourite politician]]

b) $\exists y(\text{favourite.politician.of}(y)(cc) \wedge \text{have}(cc)(y))$

(69) DISTRIBUTIVE

a) [conservatives [$\mathcal{D}$ [λ_1 t_1 have a favourite politician]]]

b) $[\lambda x x_e.\forall y(y \leq xx \rightarrow \exists z(\text{favourite.politician.of}(z)(y) \wedge \text{has}(y)(z))](cc) =$

c) $\forall y(y \leq cc \rightarrow \exists z(\text{favourite.politician.of}(z)(y) \wedge \text{has}(y)(z)))$

Because the distributivity operator can 'access' the members of *cc* through the $\leq$ relation, we get the desired two readings. Next, let's look at a disjunction, (9). The analysis on the pluralist account is as follows:

(70) COLLECTIVE

a) [computers [λ_1 t_1 are PCs or Macs]]

b) $\text{PC}(cc) \vee \text{Mac}(cc)$

19 In what follows, I translate bare plurals using plural constants, e.g. *cc*. As we will see in the final section, this is a simplification, but it serves to demonstrate the relevant point here.

(71) DISTRIBUTIVE

a) [computers [$\mathcal{D}$ [λ_1 t_1 are PCs or Macs]]]

b) $[\lambda x x_e.\forall y(y \leq xx \rightarrow \text{PC}(y) \vee \text{Mac}(y))](cc) =$

c) $\forall y(y \leq cc \rightarrow \text{PC}(y) \vee \text{Mac}(y))$

The final category of example which demonstrates the availability of Q-distributive readings for generics involves bound pronouns. Consider again:

(72) Philosophy students consider themselves clever. = (11)

As in the previous examples, the distributive reading is not predicted by the singularist analysis, yet it clearly is available. But on the pluralist account, the two readings can be analysed as follows:

(73) COLLECTIVE

a) [[philosophy students] [λ_1 t_1 consider themselves$_1$ clever]]

b) consider.clever$(pp)(pp)$

(74) DISTRIBUTIVE

a) [[philosophy students] [$\mathcal{D}$ [λ_1 t_1 consider themselves$_1$ clever]]]

b) $[\lambda x x_e.\ \forall y(y \leq xx \rightarrow \text{consider.clever}(y)(y))](pp)$

c) $\forall y(y \leq pp \rightarrow \text{consider.clever}(y)(y))$

We can derive the strict-sloppy ambiguities of (19) as follows, where in each case the index 2 is not bound and is resolved by the variable assignment to refer to the plurality denoted by 'philosophy students':

(75) SLOPPY DISTRIBUTIVE

a) [[[philosophy students] [$\mathcal{D}$ [λ_1 t_1 consider themselves$_1$ clever]]] and
[[linguistics students] [$\mathcal{D}$ [λ_1 t_1 ~~consider themselves$_1$ clever~~]]]]

b) In general, each philosophy student considers him/herself to be clever and each linguistics student considers him/herself to be clever.

(76) SLOPPY COLLECTIVE

a) [[[philosophy students] [λ_1 t_1 consider themselves$_1$ clever]] and
[[linguistics students] [λ_1 t_1 ~~consider themselves$_1$ clever~~]]]

b) In general, philosophy students consider philosophy students to be clever and linguistics students consider linguistics students to be clever.

(77) STRICT COLLECTIVE

a) [[[philosophy students] [λ_1 t_1 consider themselves$_1$ clever]] and
[[linguistics students] [λ_1 t_1 ~~consider themselves$_2$ clever~~]]]

b) In general, philosophy students consider philosophy students to be clever and linguistics students consider philosophy students to be clever.

These structures are sufficient to generate all the attested readings of (19).[20]

Finally, on this theory, all predicates take arguments of type e, so the data suggesting that predicates are not type ambiguous that represented a problem for the neo-Carlsonian view (see §2.2.5) do not pose any difficulties for the pluralist theory.

My pluralist analysis of bare plurals, coupled with the assumption of a quantificational distributivity operator, is therefore the only of the analyses considered here which predicts the linguistic facts concerning the availability of both collective and distributive readings of generics.

20 This taxonomy raises the question of whether (19) exhibits a strict distributive reading:

(i) STRICT DISTRIBUTIVE

a) [[[philosophy students] [$\mathcal{D}$ [λ_1 t_1 consider themselves$_2$ clever]]] and
[[linguistics students] [$\mathcal{D}$ [λ_1 t_1 ~~consider themselves$_2$ clever~~]]]]

b) In general, each philosophy student considers philosophy students to be clever and each linguistics student considers philosophy students to be clever.

It is hard to say whether such a reading is available, since its truth-conditions are very close to those of the strict collective reading. But here is a reason to think that such a reading is available. When $g(2)$ corefers with 'Democrats', the following has a more straightforwardly comprehensible reading that would correspond to the strict distributive reading of (19):

(ii) Democrats vote for their preferred policies and independents do too.

a) [[democrats [$\mathcal{D}$ [λ_1 t_1 vote for their$_2$ preferred policies]]] and
[independents [$\mathcal{D}$ [λ_1 t_1 ~~vote for their$_2$ preferred policies~~]]]]

b) Each Democrat votes for the preferred policies of Democrats and each independent votes for the preferred policies of Democrats.

This suggests that the strict distributive reading of (19) is genuinely available, although it is difficult to access. It is therefore a virtue that the pluralist analysis that it predicts that this reading is available.

2.5.3 *Covert binding*

As I noted in §2.2.2, Leslie (2015) and Sterken (2016) take examples like (11) and (10) to be evidence of a covert operator which is capable of binding the variable introduced by the pronoun and as evidence against Liebesman's simple kind-predication theory. Sterken (2016: 511) takes that covert operator to be 'GEN'.

Sterken's argument in favour of a covert operator is sound. As we have seen, there is good evidence of a covert operator since there is a reading of (10) where the leader varies from scientologist to scientologist. But this phenomenon is not unique to generics. It is a feature of other plural sentences too, arising for NPs headed by the definite article and demonstrative pronouns too:

(78) a) The soldiers obey their leader.

b) Those soldiers obey their leader.

Both of these sentences have a collective reading (maximum one leader) and a covarying distributive reading (one leader per soldier). So an appeal to a special operator to account for the covarying readings of generics is ad hoc. That's because we can derive both readings of (78a) in the same way as (74a), using the distributivity operator, which takes scope over the pronoun:

(79) COLLECTIVE

a) [[the soldiers] [λ_1 t_1 obey their$_1$ leader]]

b) obey.leader.of$(\iota x x.^{\text{MAX}}{*}\text{soldier}(x x))(\iota x x.^{\text{MAX}}{*}\text{soldier}(x x))$

(80) DISTRIBUTIVE

a) [[the soldiers] [$\mathcal{D}$ [λ_1 t_1 obey their$_1$ leader]]]

b) $[\lambda x x_e.\ \forall y(y \preceq x x \rightarrow \text{obey.leader.of}(y)(y))](\iota x x.^{\text{MAX}}{*}\text{soldier}(x x)) =$

c) $\forall y(y \preceq \iota x x.^{\text{MAX}}{*}\text{soldier}(x x) \rightarrow \text{obey.leader.of}(y)(y))$

We need a way of deriving covarying readings for non-generic plural sentences as well as bare plural generics. I take it that nobody would like to claim that 'GEN'

is responsible for the covarying readings of non-generic plurals. Instead, the quantificational distributivity operator can do the job. But once we have that explanation in place, why not extend the same explanation to generics? This hypothesis has the significant advantage that it does not require any special resources to account for the collective and distributive readings of generics, over and above what we need to explain the availability of those readings for non-generic sentences. While the distributivity operator has a place in a general semantic theory of plurals, 'GEN' does not.[21] The hypothesis that all generic bare plurals refer plurally, in conjunction with the availability of the quantificational distributivity operator, is *all we need* to explain the data. So at least as far as this problem goes, 'GEN' is redundant.

Recently, Liebesman has sought to defend the simple kind-predication theory from Leslie's criticism by suggesting that a distributivity operator could be invoked in the way described here (Liebesman & Magidor, 2025: 303-304). While I agree that this is the best response for a predicational theory, the arguments I have made here show that Liebesman and Magidor are committed to the pluralist analysis of bare plurals (something which Liebesman tentatively rejects elsewhere (2011: n. 8)).

2.5.4 *Co-predications*

I argued in §2.2.4.3 that I-relational theories had a problem with co-predications, where one predicate receives a collective reading and another receives a distributive reading. Here I set out two solutions to this problem consistent with my account of the meaning of bare plurals. First, it is generally assumed that the distributivity operator can operate at the VP-level, so the desired reading of (26) is one where the operator applies only to the first conjunct:

(81) [scientologists [[$\mathcal{D}$ [λ_1 t_1 are diligent]] and [λ_1 t_1 obey their$_1$ leader]]]

21 Nickel frames his analysis of 'GEN' in terms of a distributivity operator (2016: 124-125). However, the semantics he gives for 'GEN' depart radically from those usually assigned to the distributivity operator $\mathcal{D}$, so we should not accept his claim that in his theory 'GEN' "*just is* a distributivity operator".

That means that only the first conjunct distributes over the members of the plurality denoted by 'scientologists'. (81) may be analysed as follows:

(82) $\forall y(y \preceq ss \rightarrow \text{diligent}(y)) \wedge \text{obey.leader.of}(ss)(ss)$

Alternatively, 'are diligent' might be viewed as a P-distributive predicate. In that case, we can dispense with the distributivity operator and get the distributive reading of the first conjunct from a pseudo-equivalence. I believe that these two solutions generalise to all cases of co-predication.[22]

2.6 RIGID DESIGNATORS

In this final section, I will consider an objection to the claim that generic bare plurals refer plurally. The worry is that by taking bare plurals to refer to pluralities rather than groups, a referential theory of generics makes incorrect predictions about the behaviour of bare plurals embedded in modal contexts (Carlson, 1977: 159-160). In particular, sentences like the following should intuitively come out true when the modal 'might' takes wide scope, since there is no reason that there should not have been more dogs than there in fact are:

(83) Dogs might have been more numerous than they are (Liebesman, 2011: n. 8).

If the bare plural 'dogs' is a *name* – either of a kind or a plurality – then we would expect it to be a rigid designator, i.e. it should refer to the same object or objects in every possible world (in which those objects exist) (Kripke, 1980). Consider, for example, the plural name 'Oxbridge', which refers to the Universities of Oxford and Cambridge (Oliver & Smiley, 2016: 79), or 'Universities at Medway', which refers to the University of Greenwich, the University of Kent and Canterbury Christ Church

22 Koslicki (1999) discusses some of the advantages that a plural reference analysis of generic bare plurals has over a relational theory in cases of co-predication. However, her concern is with co-predications with both kind-level predicates (e.g. 'are extinct') and individual-level predicates (e.g. 'are intelligent') and she seems not to have considered the range of data that I have discussed here.

University. Intuitively these names are rigid designators, like singular names, in the sense that the following have no true reading:[23]

(84) a) Oxbridge might not have been Oxford and Cambridge.

b) Universities at Medway might have been more numerous than they are.

So, if 'dogs' is a name of a plurality, (83) says that this plurality might have had more members than it actually does. But a plurality cannot have had more members than it in fact does, since it would be a different plurality if it did (Williamson, 2003: 457), so (83) is predicted to be necessarily false.

Sticking with a referential approach, the singularist analysis offers a way of avoiding this undesirable result. Groups (or kinds), do not have their members necessarily; instead, they may have different members in different possible worlds. That is because groups are something over and above their members, so as we track the identity of the group across possible worlds, we track the group, rather than its members. In contrast, pluralities are nothing over and above their members, so we can only track the members across worlds (Florio & Linnebo, 2021: 206-207). A name for a group (or kind) may be a rigid designator, but that group (or kind) may have different members in different worlds. That explains why (83) may be true.

However, this only constitutes an objection to my pluralist account of generic bare plurals if we have good reason to think that they are rigid designators. That would be plausible if bare plurals are *names*, but not all plural expressions are rigid. Just as definite descriptions may fail to be rigid in the singular case, plural definite descriptions may also fail to be rigid. For example:

(85) The philosophers in Britain are intelligent.

In another possible world, there are some philosophers in Britain who are not in Britain in the actual world, and in that world the plural definite description in (85)

23 More controversial are cases of plural names which have the syntactic form of a definite description, such as 'the Medway Towns', which refers to the towns of Chatham, Gillingham, Rainham, Rochester and Strood, and 'the Channel Islands', which refers to many (but not all) of the islands which lie in the English Channel (Rumfitt, 2005: 120); Rumfitt, however, argues that despite their syntactic form, these terms are indeed rigid (and hence proper names), since the description is used only to fix the reference of the name, rather than to give the meaning of the expression.

would refer to those philosophers, so it is not rigid. Generic bare plurals are like this, in that their reference depends on what objects satisfy some description and this may vary from world to world. Bare plurals pattern with definite descriptions, not names, in this respect. This is not surprising: genericity can be expressed using plural definite descriptions in many languages, including French and Italian (Farkas & de Swart, 2007; Behrens, 2005). If bare plurals refer to pluralities in the same way that definite descriptions do, we would expect their extension to vary from world to world – so there may be more dogs in some other world than there actually are and (83) may be true. This shows that an argument from the observation that bare plurals are not rigid designators does nothing to motivate the singularist analysis. It does, however, strengthen the analogy between bare plurals and plural definite descriptions.

2.7 CONCLUSION

I have argued that neither singularist referential theories nor I-relational theories can account for the full range of readings which are available for generics. D-relational theories fare somewhat better, but are committed to implausible assumptions about the semantic type of the predicates involved in the collective readings of generics. In contrast, a pluralist referential account of bare plurals delivers the full range of attested collective and distributive readings, without introducing any novel operators (like 'GEN' or '*PM*') or objects (kinds), except for the distributivity operator already proposed to explain the same contrast in other plural sentences. Importantly, even those who reject the VP-level distributivity operator for non-generic plurals (e.g. Moltmann, 1997) still need some explanation for the availability of the Q-distributive readings of plural sentences which explains the contrasts between plural and group NPs and it should be possible to extend any such explanation to cover bare plurals too. So the argument that generic bare plurals must denote pluralities does not hang on the details of the analysis proposed here.

The pluralist account of generic bare plurals explains how the collective and distributive readings of generics can be derived in the same way that collective and

distributive readings of other plural sentences are. That means less explanatory work for a semantic theory of generics and more for a general theory of plurality.

Two important questions have been left unanswered here. First, which objects are referred to by bare plurals? There is good reason to think that no extensional account will be adequate; one clue as to the correct answer to this question may lie in Cohen's observation that bare plurals are intensional with respect to their time, but not possible world, index (Cohen, 1999b: 224-225). If that claim is right, then it would constitute prima facie evidence that bare plurals refer to the objects in the extension of a noun phrase across times, but not across possible worlds. Second, there is the compositional question of how bare plurals come to refer to the pluralities that they refer to. I answer both of these questions in the next chapter.

3

THE TEMPORAL PLURAL REFERENCE THEORY OF GENERICS

3.1 INTRODUCTION

In debates over the semantics of bare plural generics, there are two widely held articles of faith: (i) there is a covert adverb of quantification, 'GEN', which is responsible for the genericity of generics and (ii) at least some bare plurals refer to kinds. In this chapter, I will deny both and argue in favour of a semantic theory for bare plural generics which does without either 'GEN' or a domain of kinds.

The error, I argue, is to view generics as special and hence deserving of ad hoc semantic resources to account for their meaning. I will argue that a more fruitful methodology is to recognise the parallels between bare plural generics and other plural expressions. This departure from the orthodoxy is justified on the basis of evidence that bare plural generics are plurally referring expressions, like plural definite descriptions (I will also use 'definite plurals' interchangeably).[1]

My objective in this paper is not to make any novel arguments against existing semantic theories of generics, but to set out in detail a new theory, its motivations and the ways in which it accounts for the linguistic data at least as well as the existing theories. Since this theory proceeds for the most part by adopting resources that are independently motivated within a more general theory of the semantics of plurals, I take it to be superior to the existing theories, which introduce an ad hoc generic operator, an ontology of kinds or both.

The paper is organised as follows. First, I set out the target phenomenon and the existing space of theories. I present evidence from co-predications which weigh heavily

1 Parallels between definite plurals and bare plural generics have been acknowledged in the past, for example by Fodor (1970), von Fintel (1997), Chierchia (1998), Löbner (2000), Gajewski (2005), Dayal (2004, 2006), Heim (2011), Horn (2021) and Guerrini (2024). However, none of these authors fully develop the view that generic bare plurals *just are* definite plurals, as I do here.

in favour of a referential theory of generics. Then I set out in detail the motivations for a plural reference theory of bare plural generics (§3.3), before surveying three existing versions of the theory and finding each wanting (§3.4). In §3.5, I state my own theory in detail and set out some of its virtues. The next sections show how the theory can respond to three challenges relating to the meaning of characterising sentences: structural ambiguity (§3.6), non-maximality and homogeneity (§3.7) and modality (§3.8). In the final section, I explain how pluralities can satisfy predicates which *prima facie* may only be true of kinds and show how several features of these predicates are predicted by the assumption that they are in fact collective predicates.

3.2 THE LANDSCAPE

3.2.1 *The target*

The target phenomenon of this chapter is the generically interpreted English bare plural sentence, such as the following:

(1) a) Ravens are black.
b) Beavers build dams.
c) Cryptocurrencies are digital assets.
d) Classical logicians are intelligent.

These are *characterising sentences*; they say something about what is generally true of the category they describe without specifying proportion or frequency. Linguistically, generics are characterised by the absence of an overt quantifier or adverb of quantification, such as 'most', 'generally', or 'normally'.[2]

2 Characterising sentences come in other linguistic forms too, such as definite and indefinite singulars:

(i) The lion has a mane.
(ii) A beaver builds dams.

My discussion focuses only on bare plural generics, so I will generally use the term 'generics' to refer to bare plural generics unless the context makes it clear that something broader is intended. The reason for the focus on bare plurals is that the linguistic phenomena which fall under the heading of 'generics' are so diverse that it is not reasonable to expect that these sentences having anything in common, as far as the compositional semantics go (Collins, 2015). Though these different linguistic forms are similar in

Although characterising sentences can take a bare plural form, not all bare plural sentences are characterising sentences. Some ascribe a property to a kind which cannot be true of individual members of the kind:

(2) a) Dodos are extinct.
b) Rhinos are rare.
c) Chillis were introduced into India by the Portuguese.
d) Problem gamblers are abundant.

These are known as *kind-predication sentences*. These sentences do not generalise across members of a kind, but say something about a kind as a whole.

Not all bare plurals are interpreted generically. Instead, many bare plurals are interpreted existentially, particularly in *episodic sentences*, which say something about particular objects or situations.

(3) a) Ravens are nearby.
b) Beavers are building dams.
c) Cryptocurrencies are sold here.
d) Classical logicians are working.

In each case, the sentence is about some particular object(s) or situation. I will employ the following test to distinguish generically interpreted bare plural noun phrases (NPs) from existentially interpreted ones: existential bare plurals are upward-entailing, whereas generic bare plurals are not (Krifka et al., 1995).

(4) *Generic bare plurals*
a) Ravens are black. $\not\Rightarrow$ Birds are black.
b) Classical logicians are intelligent. $\not\Rightarrow$ Logicians are intelligent.

(5) *Existential bare plurals*
a) Ravens are nearby. $\Rightarrow$ Birds are nearby.
b) Classical logicians are working. $\Rightarrow$ Logicians are working.

meaning, we should attempt to develop separate compositional semantic theories for each construction; with luck, this approach will explain both the similarities and the differences between them.

3.2.2 *The orthodoxy*

I now turn to the orthodox view of the semantics of generics, which I call *relational theories*.[3] According to relational theories, all characterising sentences contain an unpronounced expression, 'GEN', whose meaning is a relational operator. Characterising sentences are assigned a tripartite structure (Krifka et al., 1995: 26):[4]

(6) $gen(x_1...x_n)$: Restrictor$(x_1...x_n)$[Matrix$(x_1...x_n)$]

For example:

(7) a) Philosophers are cool.
b) $gen(x)$: philosopher(x)[cool(x)]

'GEN' is often viewed as a phonologically null adverb of quantification, with a similar meaning to 'generally' or 'normally'. What goes into the restrictor and what into the matrix is determined at least in part by the focus structure of the sentence (Partee, 1991; Cohen, 1999b; Asher & Pelletier, 2013). 'GEN' is often viewed as a modal quantifier whose domain includes not only actual but possible objects, worlds or situations (Dahl, 1975; Krifka, 1987; Pelletier & Asher, 1997; Thakral, 2018; Kirkpatrick, 2023). The fact that generics admit exceptions is usually taken to be a result of some restriction on the domain of quantification, for example to normal instances of the restrictor (Eckardt, 1999; Greenberg, 2003).

3.2.3 *The heresy*

According to the heretics, or proponents of *predicational* theories of generics view, all generics have a subject-predicate structure. Typically, generic bare plurals are analysed as referring to kinds:

3 Individual authors will of course differ in their views, but I will give a general characterisation of the core aspects of the view, according to which 'GEN' is an adverb of quantification. See Chapter 1 for a more detailed discussion of the different forms of relational theories can take.

4 I use *gen* for the metalanguage translation of the object-language expression 'GEN'. For more on this quantificational structure see Lewis (1975a) and Heim (1982).

(8) cool(PHILOSOPHER)[5]

The simplest versions of the kind-predication theory stop there (Liebesman, 2011; Liebesman & Magidor, 2025), but more sophisticated versions go on to posit some kind of predicate modifier which modifies a predicate which may be true of individuals (call this an 'individual-level' predicate) into a predicate which may be true of kinds (a 'kind-level' predicate) (Carlson, 1977; Teichman, 2023). (7a) is represented by (9), where PM is the predicate modifier:

(9) $PM([\lambda x.\text{cool}(x)])(\textsc{philosopher})$

3.2.4 *The synthesis*

Some theories accept both that all generic bare plurals are referential and that all characterising sentences involve the covert relational operator, 'GEN'. The role of 'GEN' in these theories is to generalise either over ways of being a normal member of a kind (Nickel, 2016) or over the objects which the bare plural refers to, as well as situations involving those objects (Chierchia, 1998; Kirkpatrick, 2023; Guerrini, 2024). Simplifying the analysis somewhat, on Guerrini's proposal, *gen* quantifiers over the parts of a mereological sum, k_s, which is the sum of the members of k at a world s:

(10) $gen(x, s) : x \leqslant \textsc{philosopher}_s \wedge C(x, s)[\text{cool}_s(x)]$[6]

I will refer to referential relational theories as *D-relational* theories, and those which are not referential as *I-relational*, since, on the former, bare plural NPs are treated as definites and, on the latter, as indefinites.

5 Throughout, I use CAPITALS for expressions in the metalanguage which denote a kind.
6 $C(x, s)$ means that x is contextually salient in s.

3.2.5 *The consensus*

If there can be said to be consensus on anything in the semantics of generics, it is that there are some sentences which have the same logical form as (8). These are uncontroversial cases of kind-predication sentences such as:

(11) a) Dodos are extinct.
b) extinct(DODO)

(11a) is said to be an instance of kind-predication in the sense that the property is held by the kind as a whole, not its members. Even proponents of the relational view of characterising sentences acknowledge that sentences like (11a) cannot be assigned the tripartite structure in (6) (Krifka, 1987; Wilkinson, 1991; Gerstner-Link & Krifka, 1993; Leslie, 2015; Asher & Pelletier, 2013), since there is no way of interpreting *gen* which gets the truth-conditions correct for both (11a) and characterising sentences like (7a).

(12) $gen(x)$: dodo(x)[extinct(x)]

This is because (11a) says of the entire kind that it is extinct; (12) says that it is generally true of individual dodos that they are extinct.

3.2.6 *Co-predications*

Before I state my own theory, I will set out what I take to be conclusive evidence that generic bare plurals are referential expressions. This evidence concerns cases of *co-predication*, which are sentences in which two apparently incompatible properties are attributed to a single argument (Liebesman & Magidor, 2025). The problematic cases for generics are ones like the following:

(13) a) Mosquitoes are annoying and widespread.
b) Komodo dragons, which are venomous, are rare.

c) Dodos, which are extinct, were birds.

d) Because pandas only eat bamboo, they are nearly extinct.

What makes these cases of co-predication is that the same subject is said to have two properties, one which intuitively applies only to kinds (e.g. being extinct), the other which intuitively applies only to individuals (e.g. being a bird). There can be no doubt that these are grammatical and meaningful sentences of English, so any theory of generics must explain how they are interpreted.

These sentences pose a challenge to I-relational theories, since those theories posit a systematic ambiguity in the logical forms of kind-predications and characterising sentences. To see this, consider the possible analyses of (13a):

(14) a) $[\lambda k.\text{annoying}(k) \wedge \text{widespread}(k)](\textsc{mosquito})$ *simple kind-predication*

b) $[\lambda k.PM([\lambda x.\text{annoying}(x)])(k) \wedge \text{widespread}(k)](\textsc{mosquito})$ *sophisticated kind-predication*

c) $gen(x) : \text{mosquito}(x)\ [\text{annoying}(x) \wedge \text{widespread}(x)]$ *I-relational*

d) $[\lambda k.gen(x,s) : x \leqslant k_s \wedge C(x,s)[\text{annoying}_s(x)] \wedge \text{widespread}(k)](\textsc{mosquito})$ *D-relational*

Assuming, as relational theorists do, that kinds are not the sort of thing that can be annoying, (14a) is either ill-formed (if there is some syntactic prohibition on applying individual-level predicates to kinds) or just false (if all individual-level predicates are false of kinds). But the I-relational theory's (14c) gets the truth-conditions completely wrong: it says that in general mosquitoes are both annoying and widespread, whereas the English sentence (13a) says that mosquitoes are generally annoying and are, as a class, widespread. The D-relational theory, on the other hand, is much better placed to account for the meaning of (13a), since both the individual-level and kind-level predicate take a kind as their argument in (14d).

One way to repair the difficulty for the I-relational theory is to propose some kind of copying of 'mosquitoes', so that the same word contributes both a kind and a property to the meaning of the sentence (Leslie, 2015: 46). Then (13a) becomes:

(15) $gen(x) : \text{mosquito}(x)[\text{annoying}(x)] \wedge \text{widespread}(\textsc{mosquito})$

Although this ad hoc solution arguably gets the truth-conditions right for (13a), there are some more challenging cases of co-predication which make it clear that this repair strategy won't generalise:

(16) a) Widespread venomous Komodo dragons prevent tourists from enjoying the beach (adapted from Collins, 2018: 58).
b) Mosquitoes are the most widespread annoying insects (Liebesman & Magidor, 2025: 271).

In (16a), it is the collection of widespread venomous Komodo dragons which is said to prevent tourists enjoying the beach, so generalising over individual Komodo dragons will not capture its truth-conditions. As for (16b), Liebesman and Magidor point out that this sentence cannot be analysed conjunctively, since (16b) does not require that individual mosquitoes are the most annoying insects; instead, it says that (among insects) mosquitoes are the most likely to be both widespread and annoying.

They also argue that even a D-relational theory will struggle to account for the meaning of (17), since *gen* would need to take scope over 'annoying' and 'insects' but under 'widespread', since 'widespread' is not of the right semantic type to be modified by *gen*:

(17) Mosquitoes are annoying widespread insects.

This is not to say that a solution to the problem of co-predication cannot be found on behalf of a relational theory. But given the problems which arise from treating individual-level predicates as categorically distinct from kind-level predicates (which is a central commitment of all relational theories), co-predications present a substantial problem for all forms of relational theory.

3.2.6.1 *Co-predications: evidence in favour of uniform kind-predication?*

Authors who present the co-predication data usually do so as evidence in favour of a uniform kind-predication theory of generics (Liebesman, 2011; Liebesman & Magidor, 2025; Teichman, 2023). Yet the move to kind-predication does not necessarily follow,

unless we have good reason to think at least some generics involve reference to kinds. The argument must be as follows:

P1: Some generics involve reference to kinds.

P2: A uniform theory of the logical form of generics is required to explain the co-predication data.

P3: If P1 and P2, then all generics involve reference to kinds.

C: All generics involve reference to kinds.

Liebesman and Magidor (2025: §§10.3-10.6) go to some length to establish P1. Rather than responding to their arguments directly, I will make an alternative positive proposal – generics involve plural reference to ordinary objects – and defend it against some objections. There are two reasons for this dialectic move.

First, as Kleinschmidt and Ross (2013: 141-142) have argued, the semantics of bare plural generics cannot consistently involve reference to kinds on pain of paradox, given the sentence:

(18) Kinds which are not members of themselves are numerous.

There can be no kind whose members are all the kinds which are not members of themselves, since that kind is a member of itself iff it is not a member of itself. The kind-predication view predicts that (18) is either false or undefined (since the argument of the predicate fails to denote), yet it is true (I discuss this argument in more detail in Chapter 1).

Second, I'll show in §3.9 that we do not need to admit kinds into our ontology to act as the bearers of what are usually thought of as kind-level predicates, such as 'are extinct', so the intuitive pull of P1 is diminished.

Together, these two considerations give us good reason to reject P1.

3.3 MOTIVATIONS FOR A PLURAL REFERENCE THEORY OF GENERICS

The theory that I develop here is a referential theory that can account for the co-predication data, but it does so without invoking an ontology of kinds. I will argue

that generic bare plurals are covert definite descriptions and refer plurally to the objects which satisfy the description at some time. I will call this the *temporal plural reference theory* of bare plural generics.

The remainder of this chapter can be viewed as an attempt to motivate this proposal and show that it is a workable semantic theory of generics in the face of the familiar arguments in favour of both relational theories and kind-predication theories. I will show that by treating English generic bare plurals as ordinary plurally referring expressions, we can resolve a number of puzzles about the meaning of bare plural generics. I begin by motivating the analogy between generic bare plurals and definite plurals.

3.3.1 *Cross-linguistic data*

Romance languages, including French, Spanish and Italian, as well as Greek and Hungarian, use definite plurals where English uses bare plurals to express genericity (Farkas and de Swart, 2007; see also Condoravdi, 1992; Chierchia, 1998; Zamparelli, 2002; Gajewski, 2005; Delfitto, 2006; Corblin, 2013; Borik and Espinal, 2015). For example, the English bare plural generic (7a) is translated into French and Italian as follows:

(19) a) Les philosophes sont cools.
b) I filosofi sono fighi.

The hypothesis that English bare plurals contain a covert definite determiner provides a simple explanation of this cross-linguistic variation in the expression of genericity. If the composition of bare plural generics in English includes a covert constituent with the roughly the same meaning as the lexically articulated definite article, we would have an explanation of the fact that the translations of English bare plurals into many other languages require the insertion of an overt definite determiner. This proposal involves a commitment to the existence of a lexical item in English that is not pronounced, but that lexical item is very similar in meaning to one which is

sometimes pronounced in English and always pronounced in many other languages. This is quite distinct from the proposal that English contains a phonologically null generic operator which is not pronounced in any known language (for discussion of this commitment, see Liebesman, 2011: 415).

3.3.2 *English definite plural generics*

There are also cases in which English definite plurals can be interpreted generically, as in the encyclopaedic use (Farkas & de Swart, 2007: 1674):

(20) The coelurasaurs [. . .] were slender-limbed, lightly-build [. . .] and obviously fast-running forms that appear to have preyed upon their smaller reptilian contemporaries (originally from *Encyclopaedia Britannica*).

The definite article is arguably optional in English in this case, as in many others. For example:

(21) a) Italians love to cook.

b) The Italians love to cook.

A generic reading is easily available for both the bare plural and the plural definite description and the insertion of overt 'the' has little effect on the meaning of the sentence, beyond suggesting that the speaker is not among the group described (Acton, 2019).

There are also cases in which an apparently attributive reading of a plural definite description seems almost synonymous with a bare plural. These definite plural sentences are most easily heard as generic when there is an explicit restriction of some kind:

(22) a) Men around here are hard-working.

b) The men around here are hard-working.

(23) a) At this university, professors wear a tie.

b) At this university, the professors wear a tie.

(24) a) Women in this family have a job.

b) The women in this family have a job.

Finally, there are cases of de-adjectival nominalisation where bare plurals are not acceptable in English, but definite plurals behave just like Romance definite plurals in receiving a generic interpretation (Longobardi, 1994; Alexiadou, 2021):

(25) a) The rich are fortunate.

b) #Rich are fortunate.[7]

(26) a) The poor are numerous.

b) #Poor are numerous.

(27) a) The English are keen sailors and enjoy competitive sailing.[8]

b) #English are keen sailors and enjoy competitive sailing.

3.3.3 *Instance-distribution predicates*

Generics with 'instance-distribution predicates' such as 'common', 'rare' or 'widespread' (called "quantificational predicates" by Krifka et al., 1995: 95-98) are particularly amenable to a plural reference analysis:

(28) a) Insects are common.

b) Rhinos are rare.

c) Mosquitoes are widespread.

Krifka et al. suggest that these predicates are unlike kind-level predicates, such as 'are extinct', in the sense that they are about the actual distribution of instances of their argument (e.g. (28a) depends on the distribution of insects in some area). They suggest that these predicates could be treated as second-order predicates, i.e.

7 '#' indicates a judgment of unacceptability.
8 https://en.wikipedia.org/wiki/England#Sport

predicates of predicates, similar to the treatment of quantified NPs in Generalised Quantifier Theory (Barwise & Cooper, 1981). But if these predicates can be treated as monadic second-order predicates, they can also be interpreted, following Boolos (1984), as predicates of pluralities (see also Williamson, 2013: 241-243).[9] In fact this is surely the more natural interpretation: the predicate 'are rare' does not say something about the extension of a predicate, but rather about the instances themselves (collectively) (Izumi, 2012: 210).

The second-order analysis also fails to explain cases of co-predication (see also the examples in (13)):

(29) Insects are annoying and common.

That's because, on this analysis, the same NP is an argument of both an individual-level predicate ('are annoying') and a second-order predicate ('are common'). If one predicate takes individuals or pluralities as its argument and the other takes first-order predicates, how can they both be predicated of a single argument? One option would be to argue that the type mismatch is locally repaired through a type-shifting mechanism, $\Downarrow$, from expressions of type $\langle e, t\rangle$ to expressions of type e, the type of individuals and pluralities.

(30) $\text{annoying}(\Downarrow(\text{insects})) \wedge \text{common}(\text{insects})$

On this view, 'are common' applies to the standard meaning of 'insects' (a function of type $\langle e, t\rangle$ which is true of any plurality all whose members are insects), while 'are annoying' applies to the type-shifted meaning of 'insects' (the plurality of all insects). But what motivation is there for such a mechanism when the same truth-conditions can be arrived at by assuming that both predicates are first-order and 'insects' makes just one contribution to the meaning of the sentence? The plural interpretation is much more natural.

9 Throughout, I will use the singular noun 'plurality' as a stand-in for a plural when the plural is awkward. This should not be read as bearing an ontological commitment to plural objects, but merely as a way of speaking.

3.3.4 *Compatibility with collective predicates*

Unlike indefinite singular generics but like definite plurals, bare plural generics are felicitous with collective predicates (Krifka et al., 1995: 89):

(31) a) Wolves surround their prey.
b) #A wolf surrounds its prey.
c) The wolves are surrounding their prey.

(32) a) Buffaloes form protective circles (Nickel, 2016: 125).
b) #A buffalo forms protective circles.
c) The buffaloes are forming protective circles.

This contrasts suggests that plurality plays a crucial role in the meaning of bare plural generics. The most relevant readings of the bare plural generics (31a) and (32a) are intermediate distributive readings: where it is generally true of *salient groups of* wolves (buffalos) that they surround their prey (form protective circles). As Nickel argues (2016: 124-127), this is analogous to the availability of intermediate distributive readings with other plural NPs, including definite plurals (for further discussion of these readings for non-generic plurals, see Gillon, 1987; Schwarzschild, 1996).

3.3.5 *Distributive and collective ambiguities*

Schubert and Pelletier (1987) observed that bare plural generics exhibit scope ambiguities with certain predicates. For example:

(33) Conservatives have a favourite politician.
a) There is a politician who is the favourite politician of conservatives in general (exactly one politician).
b) In general, for each conservative there is a politician who is their favourite politician (one politician per conservative).

We witness the same pattern with plural definite descriptions and demonstratives (Winter & Scha, 2015: §4.2):

(34) The/these conservatives have a favourite politician.

a) There is a politician who is the favourite politician of the/these conservatives (exactly one politician).

b) For each of the/these conservatives there is a politician who is their favourite politician (one politician per conservative).

The availability of these two readings for non-generic plurals can be explained in terms of a collective-distributive ambiguity (de Vries, 2017). Indefinite singular generics like (35) lack the collective reading, which suggests that the availability of the collective reading of the predicate depends on the argument of the predicate denoting a plurality.

(35) A conservative has a favourite politician.

A very plausible explanation for the availability of both readings of (34) is that a silent quantificational distributivity operator is optionally present in the sentence (Link, 1987; Winter, 2001; de Vries, 2017; Champollion, 2017; also Chapter 2 in this thesis). When the operator is present, we get the distributive (b) reading. If generic bare plurals are plural referring expressions, then the very same operator would explain the ambiguity in (33). We then explain the lack of a collective reading for (35) on the basis that it arises only when a predicate is applied to a plurality; since no plurality is involved in (35), there is no collective reading.

3.3.6 *Cumulative readings*

Kirkpatrick (2023) has observed that generics, like definite plural sentences, exhibit so-called 'cumulative' readings of coordinated predicates (see also Guerrini, 2024).

(36) The ducks are swimming and flying (Winter, 2001).

On its cumulative reading, (36) is true iff each of the ducks is either swimming or flying and at least one duck is swimming and at least one duck is flying. Kirkpatrick discusses the following example of a cumulative generic:

(37) Elephants live in Africa and Asia (originally from Nickel, 2016).

Kirkpatrick argues that the cumulative reading of (37) means: every (normal) elephant lives in at least one of Africa or Asia and each of Africa and Asia has at least one (normal) elephant living in it. But indefinite singular generics do not admit cumulative readings like this:

(38) An elephant lives in Africa and Asia. *existential reading only*

As Kirkpatrick argues (2023: 416), a plausible explanation for this contrast is that the argument in (37) but not (38) denotes a plurality. This suggests that the fact that the semantic plurality of bare plurals makes an important contribution to their meaning and the semantics of (37) and (38) should reflect this. That is easily done if bare plurals, like definite plurals but unlike indefinite singular NPs, denote pluralities.

3.3.7 *The non-maximality/homogeneity paradigm*

3.3.7.1 *Non-maximality*

Discussion of the semantics of generics often begins with the observation that generics admit exceptions. It is also often noted that the number or type of exceptions that they admit varies from sentence to sentence or utterance to utterance (e.g. Leslie & Lerner, 2022).[10] What is not always acknowledged in this literature is that the same observation has been made concerning other sentences with plural NPs, with most discussion focusing on plural definite descriptions. This phenomenon has long been discussed by linguists under the heading of 'non-maximality' (Dowty, 1987; Yoon, 1996; Brisson, 1998, 2003; Lasersohn, 1999; Winter, 2001; Malamud, 2012; Schwarz, 2013a; Križ & Spector, 2021; Bar-Lev, 2021). The general pattern of entailments is this:

10 For detailed discussion of the variable tolerance of exceptions in generics, Sterken (see 2015a) and Nguyen (2020).

(39) a) The boys are sleeping. ⇏ Every boy is sleeping.

b) Dogs bark. ⇏ Every dog barks.

Which and how many exceptions a sentence with a plural definite description tolerates varies from case to case, depending, for example, on the predicate (Yoon, 1996) or on the interests of the speaker (Lasersohn, 1999; Malamud, 2012; Križ, 2016; Križ & Spector, 2021). This is highly reminiscent of the phenomenon of non-maximality as it arises in bare plural generics. A unified explanation of non-maximality for both bare plurals and plural definite descriptions (and other plural NPs that exhibit it) is therefore preferable, absent linguistic evidence that the two phenomena should receive different explanations (I discuss this in greater depth in §3.7). This rules out using 'GEN' as a way of deriving non-maximality for bare plural generics, since there is no reason to think that 'GEN' is involved in non-maximality for non-generic sentences.

3.3.7.2 *Homogeneity*

A closely related phenomenon to non-maximality is that of *homogeneity* (Fodor, 1970; Carlson, 1977; Krifka, 1996; von Fintel, 1997; Löbner, 2000; Breheny, 2005; Gajewski, 2005; Magri, 2012, 2014; Križ, 2015; Bar-Lev, 2021; Amiraz, 2020; Chierchia, 2022; Kirkpatrick, 2025). Homogeneity is the phenomenon of a NP being interpreted as having a (quasi-)universal positive meaning when it is in an upward entailing environment and as having a (quasi-)universal negative meaning when it is in a downward entailing environment, such as within the scope of negation. For example:

(39) a) Matt saw the students. ≈ Matt saw (almost) all of the students.[11]

b) Matt didn't see the students. ≈ Matt saw (almost) none of the students.

(40) a) Peter respects philosophers. ≈ Peter respects (almost) all philosophers.

b) Peter doesn't respect philosophers. ≈ Peter respects (almost) all no philosophers.

This polarity contrast is unexpected if the (b) sentences are the negations of the (a) sentences, since their truth-conditions are not complementary.

11 Throughout, I will use ⌜$\phi \approx \psi$⌝ to mean that ϕ is adequately paraphrased by ψ. The addition of 'almost' accounts for the possibility of non-maximal readings.

Not all predicates give rise to homogeneity. In particular, instance-distribution predicates, kind-level predicates like 'are extinct' and other predicates which hold of a class of objects as a whole do not exhibit this kind of polarity contrast:

(41) a) The students are numerous. $\not\approx$ (Almost) all of the students are numerous.
b) The students are not numerous. $\not\approx$ (Almost) all of the students are numerous.

(42) a) Dodos are extinct. $\not\approx$ (Almost) all dodos are extinct.
b) Dodos are not extinct. $\not\approx$ (Almost) all dodos are not extinct.

(43) a) Pick-up trucks come in four sizes. $\not\approx$ (Almost) all pick-up trucks come in four sizes (Le Bruyn et al., 2017).
b) Pick-up trucks don't come in four sizes. $\not\approx$ (Almost) all pick-up trucks don't come in four sizes.

Instead, each (b) sentence is true iff the (a) sentence is false, so their truth-conditions are complementary. Moreover, for these predicates, the question of non-maximality seems ill-formed: what would it mean for an individual student, dodo or pick-up truck to count as an exception to one of these sentences? Križ (2016) argues that just those predicates which give rise to homogeneity effects are ones which allow non-maximal readings (see also Amiraz, 2020). The failure of homogeneity for these examples is therefore precisely what we would expect given the characterisation of instance-distribution and kind-level predicates as those for which the question of whether an individual could be an exception to the predication is ill-formed.

3.3.7.3 *Homogeneity and non-maximality removal*

For those predicates which do give rise to homogeneity effects, both non-maximality and homogeneity are removed by the addition of 'all'.

(44) a) Matt saw all the students. $\approx$ Every student was seen by Matt.
b) Matt didn't see all the students. $\approx$ Not every student was seen by Matt.

(45) a) Peter respects all philosophers. $\approx$ Every philosopher is respected by Peter.

b) Peter doesn't respect all philosophers. ≈ Not every philosopher is respected by Peter.

The (a) sentences are both clearly false when there is some student Matt didn't see or philosopher Peter does not respect, i.e. non-maximality is removed. And the (b) sentences are much weaker than (39b) and (40b), since just one unseen student or philosopher who is not respected makes them true, i.e. homogeneity is removed.

3.3.8 *The ontology of plural reference*

I now turn from linguistic parallels between generic bare plurals and definite plurals to the question of what generic bare plurals refer to. According to kind-predication theories, they refer to kinds (as already noted, this answer is typically accepted by relational theorists too, at least for a limited set of cases). But this answer can't stand on its own, since it raises the ontological question: what is a kind?

There is of course a vast philosophical literature on the metaphysics of kinds, but that won't help here, since generics with highly modified or gerrymandered bare plural descriptions can be true:

(46) a) Elite Scottish football teams from outside Glasgow are rare.

b) Single-use plastics bags sold by retailers in the UK with more than 250 full-time equivalent employees cost 10p.

Presumably there is no kind corresponding to the bare plural NP 'Elite Scottish football teams from outside Glasgow', yet (46a) is a true and therefore meaningful sentence so the bare plural must refer to something.

The kind-predication theory therefore has a non-trivial question to answer here.[12] In contrast, a plural reference theory has no ontological question to answer: plural reference is ontologically innocent, in the sense that no objects are required to exist beyond the individuals that are among the pluralities referred to by the plural expressions of

12 For some further discussion of the problems associated with reference to kinds, see Moltmann (2013: Chapter 1), Parent (2021) and Kleinschmidt and Ross (2013: 141-142); for a more sympathetic view, see Liebesman and Sterken (2021).

the language (Boolos, 1984). For example, (46a) commits us only to the existence of elite Scottish football teams from outside Glasgow, not a kind all of whose members are those football teams.[13] As such, highly modified bare plurals in felicitous generics pose no problem to a plural reference theory of bare plural generics.

3.3.9 *Summary of the motivations for a plural reference theory*

I pause here to offer a brief summary of the motivations for a plural reference theory of bare plural generics:

(i) Many languages use plural definite descriptions rather than bare plurals to express genericity. English also uses the same construction in more limited contexts.

(ii) Instance-distribution predicates favour a plural reference interpretation of generic bare plurals.

(iii) Bare plural generics and definite plurals are felicitous with collective predicates.

(iv) Cumulative readings of coordinated predicates are available for both bare plural generics and definite plurals.

(v) The non-maximality/homogeneity paradigms are exhibited by both bare plural generics and definite plurals.

(vi) The ontological commitments of kind-predication theories are obscure; plural reference is ontologically innocent.

13 To verify the existence commitment, note the oddness of the following:

(i) ?Lions with purple antlers are rare.

Although it is surely unusual to encounter a lion with purple antlers, this sentence sounds odd on the assumption that there have never and never will be any such lions. See Greenberg (2003: 172) for a similar point.

3.4 EXISTING PLURAL REFERENCE THEORIES

Semantic theories of generics which analyse generic bare plurals in terms of plural reference have been proposed before. Here, I will discuss three existing theories and show why they are inadequate.

3.4.1 *Chierchia (1998)*

In what is the most influential version of a plural theory of generics, Chierchia (1998) argued that bare plurals are not names of kinds, but come to refer to kinds by a type-shifting mechanism, $\cap$ (NOM), which is a function from a property P to a function from a world s to the mereological sum (or plurality) of all instances of P in s (see also Dayal, 2006 and the recent proposal by Guerrini, 2024). Chierchia gives the meaning of the proposed type-shifter as follows (1998: 351):

$$(47)\ [\![\cap P]\!] = \begin{cases} [\lambda s.\iota P_s], \text{ if } [\lambda s.\iota P_s] \text{ is in } K, \\ \text{undefined otherwise} \\ \text{where } K \text{ is the domain of all kinds} \\ \text{and } P_s \text{ is the extension of } P \text{ in } s. \end{cases}$$

This says that the result of applying NOM to a property P is a function from worlds to the mereological sum of instances of P at each world. For Chierchia, kinds just are individual concepts, i.e. functions from worlds to mereological sums, although not all individual concepts are kinds. An inverse operation, $\cup$ (PRED), is also defined, which turns an individual concept d into the property of being among the mereological sum of the members of d at a world s; it is defined for each world s:[14]

14 $\leqslant$ here is the mereological parthood relation. See Link (1983) for discussion.

(48) $[\![^{\cup}d]\!] = \begin{cases} [\lambda x_e.x \leqslant d_s], \text{ if } d_s \text{ is defined,} \\ [\lambda x_e.FALSE] \text{ otherwise} \end{cases}$

For any kind d, $^{\cap\cup}d = d$. The final representation of characterising and kind-predication sentences is as follows (Chierchia, 1995):

(49) a) Rhinos are large. *characterising*
b) $gen(x,s) : {}^{\cup\cap}\text{rhino}(x) \wedge C(x,s)[\text{large}(x)(s)] =$
c) $gen(x,s) : x \leqslant {}^{\cap}\text{rhino}_s \wedge C(x,s)[\text{large}(x)(s)]$
d) Every situation of the appropriate type involving rhinos is a situation in which those rhinos are large.

(50) a) Rhinos are rare. *kind-predication*
b) $\text{rare}(^{\cap}\text{rhino})$
c) The kind RHINO is rare.

The two type-shifting operations allow the bare plural to be interpreted as a property when it is in the restrictor of *gen* and as denoting an individual concept when it is the argument of a kind-level predicate. I will identify two problems for Chierchia's proposal.

3.4.1.1 *Bare plurals which do not correspond to kinds*

Chierchia argues that not all bare plurals should be felicitous in generics, for example 'broken old shoes that Leo left behind' (1998: 351). To rule out such cases, he adds a domain condition to NOM requiring that the resulting individual concept must be a kind. That means that NOM will be undefined for bare plural NPs that do not correspond to a kind and hence a sentence like 'Broken old shoes that Leo left behind are brown' will fail to express a proposition.

However, this restriction is unmotivated because bare plural generics are felicitous even with highly modified or gerrymandered subject terms (see the discussion in §3.3.8). For example:

(51) a) Broken old shoes that Leo left behind are prized relics among the followers of Leo.

b) Norwegian students with names ending with 's' wear thick green socks (Greenberg, 2007: 132).

c) Figs that brown too early are bitter-tasting (Carlson, 1977: 322).

d) Players of *Settlers of Catan* who prioritise accumulating wheat are successful.

We even find felicitous generics with indexicals or demonstratives in the bare plural NP:

(52) a) [Uttered in London Zoo] Animals which have been transported from here to San Francisco Zoo are unhappy.

b) [Indicating a vacuum cleaner] Users of this product are extremely satisfied with its performance.

If Chierchia's constraint is to work, we need an independent specification of what makes an individual concept a kind. One response would be to stipulate that any bare plural that can felicitously be interpreted generically corresponds to a kind, but this suggestion gets things the wrong way around: it explains what a kind is in terms of the felicity of the generic interpretation of a bare plural. We want a theory which makes predications about which bare plurals can be interpreted generically, so this stipulation does not help.

The bare plurals in (51) and (52) can all be interpreted generically, yet intuitively do not correspond to kinds. Absent a substantive theory of kinds which explains why these bare plurals in fact do correspond to kinds, Chierchia's restriction to kinds makes the wrong predictions here. I conclude that a plural reference theory which includes a constraint that only bare plurals that correspond to kinds are felicitous in generics is not adequate.

3.4.1.2 *The semantic type of kind-level predicates*

Another problem for Chierchia's proposal is that it requires that kind-level predicates have arguments of the type of individual concepts, i.e. they denote functions of type $\langle\langle s, e\rangle, \langle s, t\rangle\rangle$ (Chierchia, 1998: 382). For example:

(53) a) $[\![\text{are extinct}]\!] = [\lambda d_{\langle s,e\rangle}.\lambda s_s.\text{extinct}(d)(s)]$
b) $[\![\text{are rare}]\!] = [\lambda d_{\langle s,e\rangle}.\lambda s_s.\text{rare}(d)(s)]$

The oddity here is in attributing properties to *intensional* objects in the actual world, rather than attributing properties (collectively) to pluralities of objects which satisfy some description in the actual world. This oddity is even clearer for non-stative kind-level predicates.

(54) a) Chillis were introduced into India by the Portuguese. = (2c)
b) $[\![\text{were introduced into India by the Portuguese}]\!] =$
$[\lambda d_{\langle s,e\rangle}.\lambda s_s.\text{introduced}(i)(p)(d)(s)]$

Intuitively, whether this predicate is true of $[\![\text{chillis}]\!]$ should depend only on actual chillis. But the predicate here takes an intensional entity as its argument, so whether it is true of its argument may depend not only on what happens in the actual world. Now consider a world in which chillis are the only spicy vegetables that Ralph hates. In other worlds, Ralph hates horseradish. Then $^\cap[\![\text{spicy vegetables which Ralph hates}]\!] \neq {}^\cap[\![\text{chillis}]\!]$. This predicts failures of substitution, but it seems as though (55) should be true iff (54a) is.

(55) Spicy vegetables which Ralph hates were introduced into India by the Portuguese.

That is because the truth of (55) only depends on which vegetables Ralph hates in the actual world.

Finally, these non-stative kind-level predicates can clearly apply not only to kinds, but to ordinary individuals too:

(56) a) Black rats arrived in Australia by ship.

b) My dog arrived in Australia by ship.

c) Therefore, black rats and my dog arrived in Australia by ship.

If 'arrived in Australia by ship' denotes a function of type $\langle\langle s, e\rangle, \langle s, t\rangle\rangle$, then (56b) should be ill-formed. One way of rescuing the account would be claim that the predicate has a flexible type, sometimes denoting a function whose argument is a kind, sometimes an ordinary individual. But that won't explain how (56c) is well-formed, since its argument is a coordination of a kind and an ordinary individual.[15]

I conclude that taking kinds to be individual concepts in the way that Chierchia does is unsatisfactory and that we should aim for a semantic theory which gives the denotation of bare plurals the same semantic type as ordinary individuals.

3.4.2 *Koslicki (1999)*

Koslicki (1999) proposed a 'two-step' analysis of plural predication, according to which generics are cases of plural predication which undergo a reinterpretation procedure. It is the result of this reinterpretation that yields the distinctive generic meaning. Working as she was before the advent of modern plural logic, Koslicki uses a higher-order language in which plural reference is interpreted in terms of second-order concepts (Boolos, 1984; Higginbotham & Schein, 1989). She therefore analyses characterising sentences as follows:

(57) Philosophers are cool.

a) $\text{cool}(\iota X(\forall y(X(y) \leftrightarrow \text{philosopher}(y))))$

b) The ys such that they are in the extension of the concept *philosopher* are also satisfiers of the (second-order) predicate 'cool'.

What does it take for a concept to satisfy a predicate? Koslicki says: "What it ultimately means for the higher-order plural predicate to apply to a concept, then, is generally for the individuals in the extension of the concept to have the characteristic denoted

15 I present a related objection concerning the semantic type of kind-level predicates in Chapter 2.

by the first-order singular predicate" (Koslicki, 1999: 456). She gives the following as a second level of analysis of (57):[16]

(58) $gen(x) : \iota X(\forall y(X(y) \leftrightarrow \text{philosopher}(y))(x)[\text{cool}(x)]$

This says that for (57) to be true, it must be generally true of things which are in the extension of the concept *philosopher* that they are cool. This of course permits exceptions, since it is not necessary for every philosopher to be cool for (58) to be true.

However, this obscures the connection between generic plural predication and non-generic plural predication. The problem is that if the genericity of generics is to be explained in terms of the availability of the reinterpretation, we are due an explanation as to why (58) is a possible for interpretation of (57a), but (59b) is not a possible interpretation of (59a).

(59) The townspeople are asleep (Lasersohn, 1999).

a) $\text{asleep}(\iota X(\forall y(X(y) \leftrightarrow \text{townsperson}(y))))$

b) $gen(x) : \iota X(\forall y(X(y) \leftrightarrow \text{townsperson}(y))(x)[\text{asleep}(x)]$

This reinterpretation is not available; although (59) can admit exceptions – townspeople who are tossing and turning in their beds – it cannot mean: the townspeople are generally or normally asleep. Its analysis presumably does not involve *gen* at all.

The explanation for this contrast cannot be that (57) is a generic and (59) is not, since this is precisely what the theory is trying to explain. It cannot itself be the explanation for the contrast. What we would need, instead, is an explanation for why certain plural predications – those whose argument is a bare plural – admit the reinterpretation in (57), while others – including those whose argument is a definite plural – do not admit reinterpretations like (58). The availability of definite plurals that admit a generic reading (see §3.3.2) makes finding such an explanation an even more difficult task, since these cases show that any explanation for the availability of the reinterpretation cannot be purely down the syntactic form of the argument of the predicate.

16 See also Cohen (2020) for the claim that the analysis of characterising sentences involves the two stages of analysis, in which *gen* is introduced only at the second stage.

3.4.3 *Moltmann (2013)*

Moltmann (2013) claims that generic bare plurals refer to modalized pluralities – pluralities of objects that satisfy a bare plural description in any possible world. Instead of offering a reinterpretation procedure for generic plural predications, Moltmann posits "derived meanings" for predicates which whose argument is a generic bare plural (Moltmann, 2013: 38). For a characterising predicate P, the derived meaning of P, P_{ext}, quantifies over the members of a modalized plurality in 'normal' conditions.

The meaning of P_{ext} is derived from the meaning of P as follows, where pp is a variable ranging over modalized pluralities, R is a relation of normality holding between indices i and i' iff i' is normal with respect to i, C is a relation which holds between a plurality xx and an index i iff xx are contextually relevant in i and $\prec$ is the plural inclusion relation (is/are among):

(60) $[\![P_{ext}]\!] = \lambda i.\lambda pp.\forall i'(iRi' \rightarrow \forall xx(C(xx,i') \wedge xx \prec pp \wedge P_i(xx)))$

Rather than focusing on the details of Moltmann's analysis, I will raise an objection here to the idea that characterising predicates have derived meanings. The problem is that it is possible to conjoin bare plural and non-generic plural NPs as the subject of a single predicate.

(60) a) Doctors and these teachers are intelligent.

b) Civil servants and those ministers wear a tie.

c) Wingers and these centre-backs are quick.

Although these sentences are a little awkward, they are perfectly interpretable. According to Moltmann's interpretations of the derived meanings, they quantify over normal circumstances and say that all relevant sub-pluralities of members of a modalised plurality have the property expressed by the underived meaning of the predicate in those circumstances. That is what gives characterising sentences their distinctive 'generic' meaning.

But in (60a), does the predicate 'are intelligent' have its ordinary or its derived meaning? If it has its derived meaning, then it won't be necessary for the actual teachers referred to by 'these teachers' to be intelligent in order for (60a) to be true. If it has its ordinary meaning, then it will be true only if all the doctors in all possible worlds are intelligent. Neither of these is plausible, so we must conclude that the predicate has neither its ordinary nor its derived meaning in this case. Since Moltmann's suggestion that predicates have derived meanings does not generalise, we may conclude that derived meanings are not involved in the interpretation of generics.

3.4.4 *Diagnosis*

The problems identified with these three theories all stem from their treatment of generics as fundamentally different different from other plural sentences, even though they involve plural reference and plural predication. For example, Moltmann is unable to reconcile her proposal that predicates have derived meanings when they apply to a generic bare plural with the data that suggest that plural predicates have univocal meanings, whether applied to generic bare plurals or any other plural NP. The theory that I propose here does without derived meanings for predicates and locates the differences in the meaning of generics and plural definite descriptions in the size of the domain of the plural NP.

3.5 THE NULL DETERMINER

My proposal is that English bare plural generics are best analysed as including a phonologically null definite determiner with roughly the same meaning as 'the'. A null determiner of this sort has been posited to explain the meaning of proper names (Burge, 1973; Elbourne, 2005; Izumi, 2012; Fara, 2015), in part because, in many languages, the determiner is phonologically realised (Elbourne, 2005: 172; for a typology of languages, see Agolli, 2023: §4.2). This is of course the same as the motivation for treating bare plurals as covert definite descriptions in §3.3.1. In this

section, I will give the semantics for the null determiner and show how the semantics explains the differences between overt plural definite descriptions and generic bare plurals. For reasons that will become clear, this is what I call the *temporal plural reference theory* of generics (TPRT).

3.5.1 *Formal preliminaries*

I begin by setting out the formal system I use to state the theory. Semantic types are defined:

DOMAINS If σ is a type, then D_σ is the domain of objects of type σ.

OBJECTS e is the type of objects and pluralities.

TRUTH VALUES t is the type of truth values. $D_t = \{0, 1\}$.

INDICES s is the type of indices. An index is tuple of a world, a time and a location $\langle w, t, l \rangle$.[17]

FUNCTIONAL TYPES If σ is a type and τ is a type, then $f_{\langle \sigma, \tau \rangle}$ is a function from arguments of type σ to values of type τ.

Throughout, I will adopt the notational conventions of Heim and Kratzer (1998). In $\ulcorner \lambda \phi : \psi .\ \chi \urcorner$, ψ is the domain condition: the function from ϕ to χ is a partial function defined only if ψ is satisfied. Subscripts indicate types of variables and $\ulcorner \lambda \alpha_\tau . \phi \urcorner$ abbreviates $\ulcorner \lambda \alpha : \alpha \in D_\tau . \phi \urcorner$. The interpretation function $[\![\cdot]\!]$ is relativised to a variable assignment g and context c. I will use the following composition rules (Heim & Kratzer, 1998):

FUNCTION APPLICATION If α is a branching node and $\{\beta, \gamma\}$ the set of its daughters, then, for any assignment g, α is in the domain of $[\![\cdot]\!]^{g,c}$ if both β and γ are and $[\![\beta]\!]^{g,c}$ is a function whose domain contains $[\![\gamma]\!]^{g,c}$. In that case, $[\![\alpha]\!]^{g,c} = [\![\beta]\!]^{g,c}([\![\gamma]\!]^{g,c})$.

17 It is plausible that more coordinates should be added to indices but for my purposes, these three will suffice.

PRONOUNS AND TRACES If α is a pronoun or trace, n is an index and g is a variable assignment such that n is in the domain of g, then $[\![\alpha]\!]^{g,c} = g(n)$.

Plural variables and constants are represented typographically with a double-letter, e.g. $xx, yy, zz \ldots$; $aa, bb, cc \ldots$. Unlike singular terms, $x, y, z \ldots$; $a, b, c \ldots$, these terms may take *one or more* object as a value. $\preceq$ is a logical predicate which translates 'is or are among' and relates two terms such that $xx \preceq yy$ iff there are some things among xx and everything which is among xx is among yy. Identity is a special case of inclusion: $xx = yy \leftrightarrow_{df} (xx \preceq yy \wedge yy \preceq xx)$.

I introduce a term connective, $+$, which forms plural terms from two terms. This will translate the English 'and' as it appears between two names. For example, '$a + b$' is a compound term which refers to a and b. A plural quantifier, $\exists$, is interpreted as meaning 'there are some things ... such that ...'. The plural universal quantifier, $\forall$, is defined as the dual of $\exists$. Singular objects, those which can be the value of a singular term, are then defined: $\textsc{sing}(xx) \leftrightarrow_{df} \forall yy(yy \preceq xx \rightarrow xx = yy)$.[18]

3.5.2 *Plural definite descriptions*

I will adopt the following lexical entry for 'the' here, inspired by Sharvy (1980), Schwarz (2009) and Elbourne (2013), which covers both the singular and plural cases:[19]

18 Nothing important in this chapter hangs on the choice of a plural logic; friends of a mereological (set-based) analysis of plurals may substitute out my plural variables in favour of singular variables ranging over sums (sets) (e.g. Link, 1983; Schwarzschild, 1996; Champollion, 2017). As Rieppel (2015) has argued, there may be less at stake between pluralist and mereological analyses of plurals than first appears and for my purposes I don't want to take a strong stand on the issue (although see Moltmann (2016) for linguistic arguments and Oliver and Smiley (2016: Chapter 3) for philosophical arguments against singularist theories of plurals).

19 According to a competing analysis of definite descriptions, these expressions are in fact ambiguous and can denote not only pluralities but super-pluralities, i.e. pluralities of pluralities (Oliver & Smiley, 2016). The linguistic data that supports this hypothesis are cases where a description appears to refer to more than one plurality:

(i) The boys who jointly solved the puzzle did it by a new method.

Oliver and Smiley describe a case in which Pierre and Paul solved the puzzle together by an old method, while Pierre, Paul and Jacques solved the puzzle together by a new method. They point out that, in this case, the description intuitively refers to both the plurality of the three boys and the plurality of the two boys. Then the sentence is false, since only one of those pluralities solved the puzzle by a new method. One way of reconciling this kind of data with my proposal is to introduce *covers*, originally proposed by Gillon (1987) and Schwarzschild (1996) for this purpose, which are maps from pluralities to sub-

(61) a) $[\![\text{the}]\!]^{g,c} = \lambda P_{\langle e,st\rangle}.\lambda i_s : \exists xx^{\text{MAX}} P(xx)(i).\iota xx^{\text{MAX}} P(xx)(i)$

b) ${}^{\text{MAX}}P =_{df} \lambda xx_e.\lambda i_s.P(xx)(i) \wedge \forall yy(P(yy)(i) \rightarrow yy \preceq xx)$

My use of ι follows Elbourne (2013), except that I allow for the plural case. $\ulcorner \iota xx P(xx)(i) \urcorner$ is a term which, if defined, refers to the things which together exclusively satisfy P in i. $\ulcorner {}^{\text{MAX}}P \urcorner$ is a property true of only the maximum of P, i.e. the things which satisfy P which everything that satisfies P is among. If something satisfies ${}^{\text{MAX}}P$, then nothing else does.[20]

This lexical entry includes a variable ranging over indices consisting of a tuple of world, time and location $\langle w, t, l\rangle$, which sets the domain of the definite article.[21] I will refer to the index argument which sets the domain as the *domain index*. Why is this needed? First, this is needed to narrow the domain to account for 'incomplete' descriptions, such as the following:

(62) The table is covered with books.

In cases like this, the definite description cannot have the whole world as its domain, since there is more than one table in the world. This would violate the presupposition. In this case, the domain index restricts the domain of the description to some particular time and location, so that there is just one table in the domain (this is a version of what Neale (1990) calls the *implicit approach* to domain restriction). In addition, definite descriptions, like many determiner phrases (DPs), can be evaluated with respect to a time which is independent of the evaluation of the tense of the main verb of a sentence, suggesting that the domain of the DP may be indexed in some way to a time (Enç, 1986; Musan, 1999):

pluralities (Nicolas & Payton, 2025). Then a predicate may be true of a plurality with respect to a contextually salient cover when it is true of every plurality in the cover (in the example above, when every plurality in the cover, i.e. Pierre+Paul, Pierre+Paul+Jacques, solved the puzzle by a new method). Although I cannot rule out the possibility of a genuine ambiguity in the meaning of the definite article, I believe that the lexical entry here adequately captures at least one reading of the definite article and this will be sufficient for my purposes (for critical discussion of the covers analysis, see Grimau, 2021).

20 *Proof:* consider two pluralities xx and yy which both satisfy ${}^{\text{MAX}}P$. Then, by the definition of a maximum, $xx \preceq yy$ and $yy \preceq xx$. By the definition of identity, $xx = yy$.

21 In Schwarz (2009) and Elbourne (2013), following Barwise and Perry (1983) and Kratzer (2007), this variable ranges over situations; here there won't be any advantage in adopting the situation semantic framework, so I omit it.

(63) a) Some prisoners escaped yesterday. The fugitives have been caught and are now in prison.

b) The director and the vice-director were fired last November. Now the director works as a book salesman (Musan, 1999).

The domain index in the entry above reflects the fact that the domain of a definite description can be determined independently from the domain of the rest of the sentence (see Soames (1986: 357) and the response in Kratzer (2007: §4); also Cooper (1996: §4.1)). Each DP may therefore be interpreted with respect to a distinct index which is not formally constrained by any relationships to other indices. I will assume, following von Fintel and Heim (2011), Schwarz (2012) and Elbourne (2013), that there are (silent) index pronouns in the syntax of the object language. As von Fintel and Heim argue (2011: 89-93), this assumption is required to derive the full range of meanings for certain modal sentences (cf. Fodor, 1970). I represent these pronouns typographically as 'i_n'. When they are not bound by any operator, they are assigned a value by the variable assignment.

I'll give a very simple example to show how the compositional semantics go. Assuming there is exactly one raven at the index $g(1)$, using the following lexical entries and the rules of FUNCTION APPLICATION and PRONOUNS AND TRACES, (64) expresses the proposition in (64c).

(64) [[[the ravens] i_1] [are black]]

a) $[\![\text{ravens}]\!]^{g,c} = \lambda x x_e.\lambda i_s.\text{ravens}(x x)(i)$

b) $[\![\text{are black}]\!]^{g,c} = \lambda x x_e.\lambda i_s.\text{black}(x x)(i)$

c) $\lambda i_s.\text{black}\big(\iota x x^{\text{MAX}}\text{ravens}(x x)(g(1))\big)(i)$

That is, (64) expresses a proposition which is true at an index i iff the ravens at the index the variable assignment assigns 1 to are black at i.

3.5.3 *Differences between definite plurals and bare plurals*

I shall argue that the crucial difference between 'the' and the proposed null determiner, '$\varnothing_{the}$', can be put down to differences in the sizes of their domains and hence can be accounted for by restrictions on the domain index of '$\varnothing_{the}$'.

3.5.3.1 *Ruling out bare singulars*

Before considering the size of the domain, we must first rule out bare singular count nouns, which are not felicitous in English.

(65) a) #Tiger has stripes.
b) #Dog barks.

Without some restriction on the arguments that '$\varnothing_{the}$' can take, we predict that these sentences should be just as grammatical as bare plurals are. I will refine a proposal by Dayal (2013: 58) to explain this; the null determiner is only defined for predicates for which the cardinality of its value varies from index to index. When defined, the cardinality of the value of a singular description is always 1, so this blocks the null determiner from combining with a singular noun. Dayal suggests that the constraint should be formalised as a requirement that the cardinality of the extension of the predicate must vary from one index to another. But that won't block (65a), since the cardinality of the set of tigers may vary from one time or place to another. Instead, we want to take the cardinality of the *iota*-term formed from the predicate, $|\iota x x P(xx)(i)|$, which will always be 1, if defined, when the predicate is a singular count noun.

NULL (TO BE REVISED)

$$[\![\varnothing_{the}]\!]^{g,c} = \lambda P_{\langle e,st\rangle}.\lambda i_s : \exists xx({}^{\text{MAX}}P(xx)(i)) \wedge \exists i' \exists i'' |\iota x x^{\text{MAX}} P(xx)(i')| \neq |\iota x x^{\text{MAX}} P(xx)(i'')|.\iota x x^{\text{MAX}} P(xx)(i)$$

3.5.3.2 *Deictic and anaphoric uses of definite descriptions*

Overt definite descriptions have two uses which are not possible for bare plurals. The first use is *deictic*, where the immediate context of utterance provides a domain for the description. In this case, the value of the domain index can be the immediate location of the speaker at the time of utterance. But bare plurals cannot be used in this way (Krifka, 1987):

(66) *Context:* There are two red books on a table.

a) The books are red.

b) Books are red.

The bare plural in (66b) cannot be used to refer to the two books on the table; the sentence can only mean that books in general are red.

The second use of definite descriptions is *anaphoric*, where the domain of the description is fixed by some preceding discourse. The typical case is an object introduced by an indefinite, which can then be referred to by a subsequent definite description (Hawkins, 1978: 86):

(67) a) Fred was discussing an interesting book in his class.

b) I went to discuss the book with him afterwards.

Here we can assume that the value of the domain index is resolved by the prior discourse (Heim, 2011: 1002). This explanation can also be extended to cover what Hawkins calls "associative anaphoric" (also known as 'bridging') uses of definites. As in the anaphoric case, prior discourse provides a domain for the definite, which can then be used to refer to objects within that domain:

(68) a) I saw a Ferrari driving past.

b) The tyres were smoking.

There seem to be no anaphoric or associative anaphoric uses of bare plurals. Compare the felicitous (associative) anaphoric uses of definite plurals with the bare plurals, which can only have an unrestricted meaning:

(69) Fred was discussing some interesting books in his class.

a) I loved books. *no anaphoric reading*

b) I loved the books. *anaphoric reading preferred*

(70) I saw some Ferraris driving past.

a) I loved tyres. *no anaphoric reading*

b) I loved the tyres. *anaphoric reading preferred*

3.5.3.3 *Domain selection*

I propose that we can locate the relevant differences in the meaning of the null determiner and the overt definite article in English in terms of the values of the domain index for which the two functions are defined. The rationale behind locating the difference in the meaning of the null determiner in its domain is, in part, that it explains the cross-linguistic data. Consider the following description of the linguistic facts in French from Corblin (2013: 356, numbering amended):

(71) a) Les étudiants sont paresseux.

b) The students are lazy.

> The sentence [(71a)] can be used for the students in my class, the students who attend my university, or any possible student. A good theory should capture the fact that these interpretations are similar in nature, and one may be skeptical about the thesis that D-linked [i.e. contextually restricted] [definite plurals] (*les étudiants* = the students of my class) are best analyzed on the basis of a kind-denoting interpretation.

The idea is that the determiner 'les' in French can take a domain of any size, encompassing the meaning of both the English 'the' and '$\varnothing_{the}$'.[22] My proposal is that

22 Robinson (2005) argues that the French definite article does not have the deictic use described above. In other languages, a grammatical distinction is made between *strong* and *weak* definite articles. Schwarz (2009), drawing on Ebert (1971b, 1971a), discusses the case of Fering, a dialect of German, which has two articles (see also Schwarz, 2013b for further cross-linguistic comparisons). In Fering, the weak article generally cannot be used anaphorically or deictically (or 'demonstratively' in Schwarz's terminology.) These contrasts suggest that there are a range of meanings that definite determiners can take cross-linguistically and that the meaning of the English definite article does not line up neatly with the French 'les' or the weak and strong articles of Fering.

we can extend a similar explanation to English. The difference will be that the null determiner has a more restrictive constraint on its domain index than the definite article. How then should the constraint on the domain index of '$\varnothing_{the}$' be formulated?

A first pass attempt to explain these facts would be to say that the domain index of the bare plural must always contain all objects whatsoever; the bare plural in (66b) would then refer to the plurality of all books. But it is not plausible that the domain of a bare plural must always be as inclusive as this, since there are many cases where the reference of a bare plural is either implicitly or explicitly restricted (see (72) below; also Greenberg, 2003: 206).

A second unsuccessful attempt would be to constrain the domain so that generic bare plurals can only refer to pluralities that constitute a kind. That would rule out anaphoric and deictic uses of bare plurals, since those pluralities (e.g. the books Fred was discussing in (69)) don't constitute kinds. We met this constraint already in §3.4.1, where it was ruled out on the basis that many bare plural NPs are felicitous in generics even when they cannot plausibly denote kinds (see (51) and (52)). A constraint in terms of kinds therefore rules out too much.

Let's look in more detail at what we need the constraint to do. On the one hand, it must rule out deictic and anaphoric uses of bare plurals. On the other, it must not exclude all restrictions whatsoever. To account for these two requirements, I propose that the domain of '$\varnothing_{the}$' should, unlike the domain of 'the', be as inclusive as possible, given the context of utterance (see Declerck, 1991 for similar discussion). Call this the Inclusivity Principle. My argument for the Inclusivity Principle will proceed by showing that the desired facts about the possible meanings of bare plurals follow from it.

3.5.3.4 *The Inclusivity Principle*

A domain may be inclusive in (at least) two dimensions: time and space. The most inclusive domain possible is the set of all objects which exist at any time or at any place; that sort of domain is plausible for mathematical generics, such as 'Triangles

have three sides'. But many contexts of utterance will support only less inclusive domains.

(72) a) In the 1980s, politicians were honest.
b) In London, politicians are honest.

The temporal adverbial in (72a) restricts the possible domain of the bare plural; the most inclusive domain (when the sentence is uttered out of the blue) includes all and only the objects in the 1980s. The locative adverbial in (72b) does the same, except with respect to the location parameter of the domain. Although it is sometimes suggested generics never accept implicit contextual restrictions (Condoravdi, 1994: 74-75), the same effect can be observed even when no restriction is made explicit:

(73) a) Software developers are well-paid (but they didn't used to be).
b) Doctors are respected professionals (but only since the start of the 20th Century).
c) Judges wear a wig (but not in the USA).

These sentences presumably do not say something about all software developers, doctors or judges at any time or place, but rather are restricted to some relatively recent period of time, including the present (as well as possibly the future), or some contextually salient location, as evidenced by the felicity of the continuations.

Finally, I can state the lexical entry for '$\varnothing_{the}$', with the INCLUSIVITY PRINCIPLE and the requirement that its argument be a plural noun reflected in the domain condition, where $inc(c,i)$ is a relation between a context c and an index i which is true of the pair iff i determines the most inclusive domain possible in c:

NULL (FINAL VERSION)

$$[\![\varnothing_{the}]\!]^{g,c} = \lambda P_{\langle e,st\rangle}.\ \lambda i : inc(c,i) \wedge \exists xx({}^{\text{MAX}}P(xx)(i)) \wedge \exists i' \exists i'' |\iota xx^{\text{MAX}} P(xx)(i')| \neq |\iota xx^{\text{MAX}} P(xx)(i'')|.\ \iota xx^{\text{MAX}} P(xx)(i)$$

A consequence of this lexical entry for the null determiner is the prediction that bare plurals are never felicitous when used anaphorically or deictically. Anaphoric

and deictic uses of bare plurals are ruled out because they inherently involve the narrowing of a domain. For example, it is the narrowing of the domain in (67) which makes sure that the maximality presupposition (which is equivalent to uniqueness in the singular case) is satisfied. Since the null determiner is only defined for the most inclusive domains, this narrowing of the domain is inconsistent with the use of the determiner.

3.5.3.5 *Why are definite plural generics not felicitous?*

Nothing I have said so far rules out the felicity of definite plurals having a generic meaning. But as the following examples demonstrate, definite plurals typically only have a contextually restricted reading:

(74) a) The birds fly.
b) The lions have fur.
c) The ravens are black.

So why can't these sentences have generic meanings? As we saw in §3.3.2, some English definite plurals can have a generic meaning, so we don't want to encode a lexical restriction on the domain of the overt definite article, since that would rule out those cases.

If definite plurals can refer to objects within a domain of any size, but bare plurals can only refer to objects within a domain which satisfies the Inclusivity Principle, it follows that using a definite plural when a bare plural is felicitous should indicate that a speaker has a more restricted domain in mind. That is because the use of the definite plural is less informative than the use of a bare plural would have been, since the definite plural could have been used to refer to some contextually salient objects, whereas a bare plural could not. By Grice's maxim of Quantity (1975), the use of a definite plural will therefore typically indicate that the speaker intends to refer to objects within a domain that does not satisfy the Inclusivity Principle, since they would otherwise have used the more informative bare plural. If we consider an example of a sentence that can get a plausible generic reading with a definite or a

bare plural, it is intuitive that, when uttered in a context which does not determine a particular restriction on the domain, a speaker of the bare plural will intend to refer to a less bounded plurality of individuals than the speaker of the definite plural (cf. Hawkins, 1978: 217).

(75) a) Italians love to cook. = (21a)

b) The Italians love to cook. = (21b)

This could be explained on the present theory by a difference in the temporal dimension of the domains; plausibly the bare plural refers more inclusively with respect to time than the definite plural, so (75a) can be interpreted as a statement about Italians within some extended period of time, whereas (75b) would be interpreted as a statement about the Italians in the actual world *right now*. The bare plural therefore has a more 'general' meaning, since its domain is more inclusive than the definite plural and so (75b) makes a statement about a larger plurality of Italians.

3.5.3.6 *A compositional example*

I will pause briefly to show how the composition of the semantics of a simple bare plural generic will go. Using the lexical entries from (64) and NULL and the rules of FUNCTIONAL APPLICATION and PRONOUNS AND TRACES, (76a) expresses the proposition in (76e), assuming that $g(1)$ is an index which contains at least one raven and satisfies the other presuppositions of '$\varnothing_{the}$':

(76) a) [[[$\varnothing_{the}$ ravens] i_1] [are black]] =

b) $[\![\text{black}]\!]^{g,c}([\![\text{the}]\!]^{g,c}([\![\text{ravens}]\!]^{g,c})([\![i_1]\!]^{g,c})) =$

c) $[\lambda x x_e.\lambda i_s.\text{black}(xx)(i)]\Big([\lambda P_{\langle e,st\rangle}.\ \lambda i : inc(c,i) \wedge \exists x x(^{\text{MAX}}P(xx)(i) \wedge \exists i' \exists i'' |\iota x x^{\text{MAX}}P(xx)(i')| \neq |\iota x x^{\text{MAX}}P(xx)(i'')|).\iota x x^{\text{MAX}}P(xx)(i)]([\lambda x x_e.\lambda i_s.\text{ravens}(xx)(i)](g(1)))\Big) =$

d) $[\lambda x x_e.\lambda i_s.\text{black}(xx)(i)]\big(\iota x x^{\text{MAX}}\text{ravens}(xx)(g(1))\big) =$

e) $\lambda i_s.\text{black}\big(\iota x x^{\text{MAX}}\text{ravens}(xx)(g(1))\big)(i)$

If we suppose that $g(1) = w$, an index denoting all of time and space in the actual world (I assume worlds always satisfy the INCLUSIVITY PRINCIPLE), then (76a) is true at an index i iff the maximum plurality in $[\![\text{ravens}]\!]$ in w is black in i.

How is the fact that not all ravens must be black in order for (76a) to be true accounted for? I will defer the complete answer to that question until §3.7, but for now, I will give a brief preview. I will argue that the non-maximality of (76e) stems from the meaning of the predicate. Not all members of a plurality xx must be black for xx to satisfy the predicate 'are black'. This is because the meaning of plural predicates is underspecified with respect to the proportion of the argument that must satisfy the predicate in order to the argument as a whole to satisfy the predicate. This is resolved contextually by the purposes of conversation, so, very roughly, the meaning of (76a) will be that enough of the ravens in the world are black for the current purposes of conversation. I explain this theory in much more detail below.

I now turn to some advantages of TPRT.

3.5.4 *Temporary generalisations*

A consequence of this semantics for bare plurals is that in most contexts bare plurals should resist interpretations according to which they hold in virtue of merely temporary states of affairs. Consider Cohen's example (Cohen, 1999a: 225):

(77) a) *Context:* By chance, every current Supreme Court judge has a prime Social Security number.
b) Supreme Court judges have a prime Social Security number.

First, in the context of Cohen's discussion, the intended interpretation of this sentence is restricted spatially to Israel; judges in other jurisdictions don't matter. This means that Israel across some period of time including the present is the largest domain in discourse context. If other jurisdictions are under consideration, then the bare plural cannot be used to refer only to Israeli judges because that would violate the INCLUSIVITY PRINCIPLE.

Second, as Cohen observed, the fact that all the current Supreme Court judges have a prime Social Security number does not suffice for the truth of (77b). But the INCLUSIVITY PRINCIPLE requires that the domain of the bare plural should extend beyond the moment of utterance, so the reference of the bare plural would not include only the current Supreme Court judges, but those in the past and future too (although not necessarily all of time). That would explain why the temporary accident of all the current judges having prime Social Security judges is not sufficient to make (77b) true.

3.5.5 *The functional reading of bare plurals*

TPRT can explain the availability of what Condoravdi (1994) calls the "functional" reading of bare plurals, which is a quasi-universal reading of a bare plural in a non-generic context:

(78) In 1985, there was a ghost on campus.

a) Students were afraid.

b) The students were afraid.

c) . . . but the students at nearby universities found it funny.

As Condoravdi argues, (78a) seems synonymous with (78b) in this context. The functional reading is quite unexpected on theories which treat generic bare plurals as semantically distinct from definite plurals. The bare plural is restricted to the campus in question in 1985, as evidenced by felicity of the continuation in (78c). On the current proposal, the context-setting sentence in (78) sets the domain for the bare plural. Given the INCLUSIVITY PRINCIPLE, the bare plural can only be interpreted with a domain which is at least as inclusive as the one specified by (78), even given other contextual support for a restricted domain. This prediction is borne out (Condoravdi, 1994: 87), as the bare plural in (79b) cannot be restricted to the immediate context which contains students, even when a definite plural could accept the same restriction.

(79) a) *Context:* The speaker and hearer know that there is a ghost haunting the campus and are looking at the library, where several students are visible.

b) Students are afraid to enter the library.

c) The students are afraid to enter the library.

The definite plural in (79c), but not the bare plural in (79b), can be restricted to the visible students.

Finally, it is worth acknowledging that the continuation in (78c) involves a widening of the domain to include nearby universities. Does that mean that the index of the bare plural in (78a) violates the INCLUSIVITY PRINCIPLE? No – the domain can always be widened *after* an utterance of the bare plural, but once a wider domain is made salient, it cannot be ignored. If the discourse is reordered so that the widening takes place *before* the bare plural, the bare plural can no longer be restricted to the original campus:

(80) a) In 1985, there was a ghost on campus.

b) The students at nearby universities found it funny.

c) ?Students were afraid.

This is presumably another instance of Lewis's (1979) observation that once a conversational standard has been raised, it is hard to lower it. Once the domain has been made more inclusive, it is hard to restrict it.

3.5.6 *Co-predications*

Next, I will show that TPRT theory accounts for the co-predication data that posed a challenge to the relational theory. First, take a simple example like (13a).

(81) a) Mosquitoes are annoying and widespread. = (13a)

b) [[[$\varnothing_{the}$ mosquitoes] i_1] [are annoying and widespread]]

c) $[\![\text{annoying and widespread}]\!]^{g,c} = [\lambda x x_e.\lambda i_s.\text{annoying}(xx)(i) \wedge \text{widespread}(xx)(i)]$

d) $[\![\varnothing_{the} \text{ mosquitoes } i_1]\!] = [\![\text{the}]\!]^{g,c}([\![\text{mosquitoes}]\!]^{g,c})([\![i_1]\!]^{g,c}) =$

e) $[\lambda P_{\langle e,st\rangle}.\lambda i : inc(c,i) \wedge \exists x x(^{\text{MAX}}P(xx)(i) \wedge$
$\exists i' \exists i'' |\iota x x^{\text{MAX}}P(xx)(i')| \neq |\iota x x^{\text{MAX}}P(xx)(i'')|).\iota x x^{\text{MAX}}P(xx)(i)]$
$([\lambda x x_e.\lambda i_s.\text{mosquitoes}(xx)(i)](g(1)))) =$

f) $[\lambda x x_e.\lambda i_s.\text{annoying}(xx)(i) \wedge \text{widespread}(xx)(i)]\big(\iota x x^{\text{MAX}}\text{mosquitoes}(g(1))(xx)\big) =$

g) $\iota x x^{\text{MAX}}\text{mosquitoes}(g(1)))$

h) $[\![\text{annoying and widespread}]\!]^{g,c}([\![\varnothing_{the}\ \text{mosquitoes}]\!]^{g,c})([\![\text{i}_1]\!]^{g,c})) =$

i) $\lambda i_s.\text{annoying}\big(\iota x x^{\text{MAX}}\text{mosquitoes}(xx)(g(1))\big)(i) \wedge$
$\text{widespread}\big(\iota x x^{\text{MAX}}\text{mosquitoes}(xx)(g(1))\big)(i)$

Because TPRT does not distinguish between kind-level and individual-level predicates, the case which Liebesman and Magidor identified as being more difficult for I-relational theories, namely (17), poses no special difficulties:

(82) a) Mosquitoes are annoying widespread insects. = (17)

b) [[[$\varnothing_{the}$ mosquitoes] i_2] [are annoying widespread insects]]

c) $[\![\text{are annoying widespread insects}]\!]^{g,c} =$
$[\lambda x x_e.\lambda i_s.\text{annoying}(xx)(i) \wedge \text{widespread}(xx)(i) \wedge \text{insects}(xx)(i)]$[23]

d) $[\![\text{annoying widespread insects}]\!]^{g,c}([\![\varnothing_{the}\ \text{mosquitoes}]\!]^{g,c})([\![\text{i}_2]\!]^{g,c})) =$

e) $\lambda i_s.\text{annoying}\big(\iota x x^{\text{MAX}}\text{mosquitoes}(xx)(g(2))\big)(i) \wedge$
$\text{widespread}\big(\iota x x^{\text{MAX}}\text{mosquitoes}(xx)(g(2))\big)(i) \wedge$
$\text{insects}\big(\iota x x^{\text{MAX}}\text{mosquitoes}(xx)(g(2))\big)(i)$

Finally, we can consider a case with an ambiguity, where one predicate could either hold of pandas in general, or each individual panda.

(83) Pandas take care of themselves but are an endangered species.

I will assume that the two readings can be disambiguated based on the presence or absence of a quantificational distributivity operator which operates at the level of the verb phrase (Link, 1987):[24]

(84) a) $[\![\mathcal{D}]\!]^{g,c} = \lambda P_{\langle e,st\rangle}.\lambda x x_e.\lambda i_s.\forall y(y \leq xx \rightarrow P(y)(i))$

b) $[\![\text{take care of themselves}]\!]^{g,c} = \lambda x x_e.\text{take.care.of}(xx)(xx)$

c) $[\![\text{endangered species}]\!]^{g,c} = \lambda x x_e.\text{endangered.species}(xx)$

23 We can arrive at this conjunctive predicate by using (an intensional version of) Heim and Kratzer's (1998) PREDICATE MODIFICATION rule.

24 The operator as defined here applies only to one-place predicates. See Lasersohn (1998) for the generalisation to any argument place of a predicate.

(85) a) [[[$\varnothing_{the}$ pandas] i_3] [[$\mathcal{D}$ [take care of themselves]] but [are an endangered species]]]

b) $[\lambda x x_e.\lambda i_s.\forall y(y \preceq xx \rightarrow \text{take.care.of}(yy)(yy)(i) \wedge \text{endangered.species}(xx))(i)]\big(\iota x x^{\text{MAX}}\text{pandas}(xx)(g(3))\big) =$

c) $\lambda i_s.\forall y(y \preceq \iota x x^{\text{MAX}}\text{pandas}(xx)(g(3)) \rightarrow \text{take.care.of}(yy)(yy)(i) \wedge \text{endangered.species}\big(\iota x x^{\text{MAX}}\text{pandas}(xx)(g(3))\big)(i)$

(86) a) [[[$\varnothing_{the}$ pandas] i_3] [[take care of themselves] but [are an endangered species]]]

b) $[\lambda x x_e.\lambda i_s.\text{take.care.of}(xx)(xx)(i) \wedge \text{endangered.species}(xx)(i)]\big(\iota x x^{\text{MAX}}\text{pandas}(xx)(g(3))\big) =$

c) $\lambda i_s.\text{take.care.of}\big(\iota x x^{\text{MAX}}\text{pandas}(xx)(g(3))\big)\big(\iota x x^{\text{MAX}}\text{pandas}(xx)(g(3))\big)(i) \wedge \text{endangered.species}\big(\iota x x^{\text{MAX}}\text{pandas}(xx)(g(3))\big)(i)$

The distributive reading in (85a) involves quantification over individual pandas and says that individual pandas look after themselves, while the collective reading in (86a) says that pandas in general take care of pandas.[25]

3.5.7 *'All' and 'most'*

Finally, I will show how TPRT fits into a wider theory of the meaning of determiners. According to Generalized Quantifier Theory, quantifiers express relations between sets and are of type $\langle et, \langle et, t\rangle\rangle$ (Barwise & Cooper, 1981). One challenge for this analysis is the possibility of partitive constructions like the following:

(87) Most/many/some/three/few/all/both of the chiefs spoke (Matthewson, 2001: 160).

25 Two comments are in order here. First, although the semantics of '$\mathcal{D}$' are given in terms of universal quantification, this will need to be modified to allow for non-maximality; but since there are non-maximal readings of plural sentences which are not generics, this presents no special difficulty here (see Križ and Spector, 2021: 160 for discussion of the non-generic case). Second, we could also introduce a version of Schwarzschild's 'Part' operator (1996), which involves quantification over salient sub-pluralities (i.e. members of a *cover*); we could then generate a reading of (83) that says, for example, that families of pandas look after themselves (for further discussion of this point as it relates to generics, see §3.3.4 and Nickel, 2016: 124-127).

Here, the restrictor of the quantifier looks more like an entity-denoting DP than a set-denoting NP. Matthewson (2001) takes this as evidence that these quantifiers take an argument of type e as their restrictor and so are of type $\langle e, \langle et, t \rangle \rangle$. This allows us to take the meaning of the definite description in (87) at face value, i.e. denoting the maximum plurality of the chiefs at the index and it is this plurality which acts as the restrictor of the quantifier. We might therefore posit the following meaning for the determiner 'most':[26]

(88) $[\![\text{most}]\!]^{g,c} = \lambda x x_e . \lambda P_{\langle e,t \rangle} . \mid \{y : y \leq xx\} \cap \{zz : P(zz)\} \mid > \mid \{y : y \leq xx\} - \{zz : P(zz)\} \mid$

This says that 'most' denotes a relation between a plurality and a set, which holds iff the number of the members of the plurality which are members of the set is greater than the number of the members of the plurality which are not.

I will show that the INCLUSIVITY PRINCIPLE plays a role in the availability of contextually restricted readings of sentences with non-partitive 'most' and 'all'. Consider the contrast between partitive and non-partitive 'all'/'most' with respect to their domains (Partee, 1995; Cooper, 1996; Matthewson, 2001; Dobrovie-Sorin & Giurgea, 2021):

(89) I am meeting some linguists and philosophers for lunch today.

a) All (of) the linguists are intelligent. *contextual restriction preferred*

b) All linguists are intelligent. *contextual restriction unavailable*

c) Most of the linguists are semanticists. *contextual restriction preferred*

d) Most linguists are semanticists. *contextual restricted unavailable*

When there is no overt definite article, it is much harder to get the anaphoric reading with a contextual restriction to the linguists that the speaker is meeting for lunch. But, as before, some contextual restrictions are acceptable, provided they adhere to the INCLUSIVITY PRINCIPLE:

(90) I am meeting with some judges from the Supreme Court and High Court of the UK today...

26 The analysis of 'all' is more complex, as it is arguably not a determiner at all, but a modifier (Partee, 1995: 583; Križ & Spector, 2021: 1162; Dobrovie-Sorin & Giurgea, 2021: 186-188).

a) All (of) the Supreme Court judges have a law degree. *contextual restriction preferred*

b) All Supreme Court judges have a law degree. *restriction to UK preferred*

c) Most of the Supreme Court judges are male. *contextual restriction preferred*

d) Most Supreme Court judges are male. *restriction to UK preferred*

If the non-partitive sentences contain the null determiner, then, assuming, as Matthewson does, that 'of' is semantically vacuous, 'all' and 'most' can be given an unambiguous interpretation in each of the sentences. The lack of the contextually restricted meaning in (89b) and (89d) can then be explained by the domain condition on the null determiner. In the context introduced by (90), an index i^* which contains just the judges the speaker is meeting would not satisfy the Inclusivity Principle, so the domain condition of 'the' in (90c) is satisfied, but the domain condition of '$\varnothing_{the}$' in (89d) is not satisfied by that index. This predicts, correctly, that (90c) but not (90d) can be used to say that most of the contextually salient Supreme Court judges are male.

To verify this prediction, consider the compositional calculation.

(91) a) $[\![\text{Supreme Court judges}]\!]^{g,c} = \lambda x x_e.\lambda i_s.\text{scjs}(xx)$

b) $[\![\text{are male}]\!]^{g,c} = \lambda x x_e.\text{male}(xx)$

(92) $[\![$ [Most [[$\varnothing_{the}$ [Supreme Court judges]] i_2]] are male$]\!]^{g,c} =$

a) $[\![\text{most}]\!]^{g,c}([\![\varnothing_{the}]\!]^{g,c}([\![\text{Supreme Court judges}]\!]^{g,c})([\![i_2]\!]^{g,c}))([\![\text{are male}]\!]^{g,c})$

b) $[\![\varnothing_{the}]\!]^{g,c}([\![\text{Supreme Court judges}]\!]^{g,c})([\![i_2]\!]^{g,c}) =$

c) $[\lambda i_s : inc(c,i) \wedge \exists x x(^{\text{MAX}}\text{scjs}(xx)(i) \wedge$
$\exists i' \exists i'' |\iota x x^{\text{MAX}}\text{sjcs}(xx)(i')| \neq |\iota x x^{\text{MAX}}\text{scjs}(xx)(i'')|).\iota y y^{\text{MAX}}\text{scjs}(yy)](g(2))$ ✗

Assuming a variable assignment g where $g(2) = i^*$ – where i^* is the index containing just the judges introduced in (90) – the composition terminates here, because the domain condition of '$\varnothing_{the}$' is not satisfied.

If, however, we assume that $g(2)$ is an index which satisfies the domain condition of '$\varnothing_{the}$', the composition can proceed:

(92) d) $[\lambda x x_e.\lambda P_{\langle e,t\rangle}. \mid \{y : y \leq xx\} \cap \{zz : P(zz)\} \mid > \mid \{y : y \leq xx\} - \{zz : P(zz)\} \mid]$
$(\iota y y^{\text{MAX}}\text{scjs}(yy)(g(2)))([\lambda x x_e.\text{male}(xx)]) =$

e) $\mid \{y : y \leq \iota y y^{\text{MAX}}\text{scjs}(yy)(g(2))\} \cap \{zz : \text{male}(zz)\} \mid >$
$\mid \{y : y \leq \iota y y^{\text{MAX}}\text{s.c.js}(yy)(g(2))\} - \{zz : \text{male}(zz)\} \mid$

That is, (92) is true iff the number of Supreme Court judges at the index $g(2)$ who are male is greater than the number of Supreme Court judges at the index $g(2)$ who are not male.

3.5.8 *Interim Summary*

Let's take stock. I have given a semantics for a null determiner which differs only in its presuppositions from the definite article 'the'. I showed how the Inclusivity Principle derives the generic readings of bare plurals but rules out deictic and anaphoric uses. I showed that the cross-temporal nature of the domain of the null determiner explains why generics cannot be made true by merely temporary states of affairs and that the semantics of the null determiner can explain Condoravdi's functional uses of bare plurals. I showed that TPRT can explain the co-predication data and, finally, that the presence of the null determiner can explain the constraints on the domain of non-partitive 'all' and 'most'.

Now that the details of the proposal are in order, I turn to some of the questions that the proposal raises:

(i) Don't structural ambiguities show that bare plural generics contain a covert relational operator (§3.6)?

(ii) How should non-maximality and homogeneity be explained (§3.7)?

(iii) Aren't generics supposed to be modals (§3.8)?

(iv) How can pluralities be extinct (§3.9)?

3.6 AMBIGUITY AND INFORMATION STRUCTURE

I begin by discussing a well-known objection to predicational theories of generics. An example from Milsark (1974) provided a Damascene moment for Carlson, leading him to abandon his original kind-predication view (Carlson, 1977) in favour of a relational theory (Carlson, 1989). The reason for this shift was that Milsark's example seemed to show that generics admit scope ambiguities which should not be available according a theory which gives generics as little internal structure as Carlson's 1977 account does.

(93) Hurricanes arise in this part of the Pacific.

a) Hurricanes generally have an origin in this part of the Pacific.

b) There arise hurricanes in this part of the Pacific.

The problem is that (93) seems to be ambiguous between two readings, one which is about hurricanes in general and one which is about some part of the Pacific Ocean. But if we have only a unary generic operator – call it *PM*, as above - we only get the first reading:

(94) $PM([\lambda x x_e.\text{arise.in}(xx)(\text{this.part.of.the.pacific})])(\textsc{hurricane})$

3.6.1 *Two notions from information structure*

Before I discuss the orthodox solution to this problem, I will introduce two notions from information structure which will play an important role in the analysis. First, *topic*: the *topic* of a sentence will be the entity or set of entities that the sentence is about. If an entity denoted by an expression X is the topic of a sentence, then it will typically be possible to rephrase the sentence by saying 'As for X,...', or by reporting an utterance of the sentence by saying 'He said about X,...' (Reinhart, 1981). A sentence may have more than one topic, as in 'As for Jack and Jill, they climbed the hill' and in some sentences no constituent will identify the topic, as in 'It is raining', where the topic of the sentence will be some contextually salient situation or event; these sentences are sometimes known as *thetic* sentences (Krifka & Musan, 2012).

Second, *focus*: focus is a prosodic feature of discourse, marked by certain pitch accents and stress (Büring, 2016). In what follows, I will annotate focus on an expression α in a sentence: $[\alpha]_F$. The primary function of focus is to indicate the presence of alternatives; one semantic role that focus plays is to assign material to the restrictor or matrix of an adverb of quantification (Rooth, 1992; Partee, 1991):

(95) a) Mary always takes [John]$_F$ to the movies.
b) *always* : Mary takes someone to the movies [Mary takes John to the movies]
c) Whenever Mary takes anybody to the movies, that person is John.

(96) a) Mary always takes John to the [movies]$_F$.
b) *always* : Mary takes John somewhere [Mary takes John to the movies]
c) Whenever Mary takes John somewhere, she takes him to the movies.

Finally, expressions which are focused are typically understood as providing new information, while expressions that are deaccented are given, in the sense that they do not provide new information (Schwarzschild, 1999). In particular, expressions in focus correspond to the wh-phrase of a (possibly implicit) question (Rooth, 1992; Roberts, 1996):

(97) Q: Who does Mary like?
A: Mary likes [Sue]$_F$.

(98) Q: Who likes Sue?
A: [Mary]$_F$ likes Sue.

3.6.2 *Mapping material to the restrictor and matrix*

Now we can turn to the relational theory's explanation of the ambiguity. With the relational operator 'GEN', we can partition different material into the restrictor and matrix to get both readings (Krifka et al., 1995: 26):

(99) a) $gen(x, y)$: hurricane(x)[this.part.of.the.pacific$(y) \wedge$ arise.in$(x)(y)$]

b) $gen(x)$: this.part.of.the.pacific$(x)[\exists y(\text{hurricane}(y) \wedge \text{arise.in}(y)(x))]$

The two structures correspond to differences in the information structure of the sentence; when 'hurricanes' denotes the topic of the sentence, we get the (a) reading; when it is in focus, we get the (b) reading (Krifka, 1995; Partee, 1991). That means we can disambiguate the (a) and (b) readings based on the accenting of the sentence, where CAPITALS mark the main pitch accent of the sentence and | marks a pause (Krifka et al., 1995: 27):

(100) a) Hurricanes | arise in this part of the pacific.

b) HURRICANES arise in this part of the pacific.

As a broader generalisation of these facts, Diesing (1992) argued that a bare plural in the matrix of an operator is obligatorily bound by Existential Closure (cf. Heim, 1982), which explains the presence of the existential quantifier in the matrix of (99b). Bare plurals outside of the matrix may be bound by another operator, such as 'GEN', explaining the generic interpretation of the bare plural in (99a).

3.6.3 *An alternative analysis*

In general, definite NPs, such as definite descriptions and pronouns, denote familiar or given objects, whereas indefinites must introduce new referents (Krifka & Musan, 2012: 22).

This suggests an explanation for the ambiguity of (93) in line with TPRT. When a bare plural is focused, it can introduce a new referent, so it can be interpreted as an indefinite and introduce a variable which can be bound by Existential Closure. That is why it gets an existential reading in (100b). But when the bare plural is outside of the focus of a sentence or is deaccented, it must be given or familiar, so it cannot be interpreted as an indefinite. That is why the bare plural in (100a), which is the topic and hence given, is interpreted generically. This generates two truth-conditionally distinct readings, which correspond to the different accents in (100) (I suppress the intensional elements of the analysis here):

(101) a) Hurricanes [arise in this part of the Pacific]$_F$.

b) arise.in($\iota x x^{\text{MAX}}$hurricanes($x x$))(this.part.of.the.pacific)

(102) a) [Hurricanes]$_F$ arise in this part of the Pacific.

b) $\exists yy$(hurricanes(yy) $\wedge$ arise.in(yy)(this.part.of.the.pacific))

(101a) is a generic about hurricanes in general, while (102a) is a habitual sentence about a particular location in the Pacific. De Swart (1996: 175) observes that a French translation resolves the ambiguity of (93), because in the reading in which the bare plural denotes the topic, (101), it is translated as a definite plural, *les ouragons*, and in the reading where the bare plural is in focus, (102), it is translated as an indefinite plural, *des ouragons*. This is as we would expect, on the assumption that the generic reading of the bare plural in English arises because of a covert definite determiner and that the determiner is not present when the bare plural is interpreted existentially. The truth-conditions in (101) and (102) also nicely line up with Krifka et al.'s (1995: 25) glosses on the two readings of (93), which make the topics of the sentences explicit:

(103) a) For hurricanes it holds: they arise in this part of the Pacific.

b) For this part of the Pacific it holds: there arise hurricanes.

TPRT also gets the entailments of the two readings of (93) right. (102a) entails (104b), but (101a) does not entail (104a).

(104) a) For heavy storms it holds: they arise in this part of the Pacific.

b) For this part of the Pacific it holds: there arise heavy storms.

In summary, when a bare plural is not focused, it must be given, so it can't be interpreted as an indefinite, which would explain why the covert determiner is felicitous for unfocused bare plurals. Conversely, bare plurals may be interpreted existentially only when they are in the focus of a sentence (see Cohen and Erteschik-Shir, 2002 for discussion of the claim that bare plurals are always interpreted existentially when in focus). This explains the ambiguity of (93) and its relationship to the information structure of the sentence, without the need for a relational operator.

3.6.4 *Existential bare plurals*

3.6.4.1 *Episodic predicates*

This solution suggests an explanation for another puzzle about bare plurals. Some bare plurals admit only existential readings. For example:

(105) a) Dogs are playing on my lawn. *existential reading only*

b) Basketball players are warming up. *existential reading only*

If bare plural NPs can be interpreted as indefinites only when they are in focus, we predict that these sentences should have existential readings when the bare plural is in focus. But if the bare plurals aren't topics in these examples, what are the topics? Erteschik-Shir (1998) argues that all sentences have a topic, but that these may be "stage topics", which are spatiotemporal locations (or perhaps situations or events).[27] For example, the stage topic of (105a) would be the location denoted by 'my lawn' at the time of utterance; for (105b) it would be the spatiotemporal location that the sentence is describing (e.g. a basketball court at the time of utterance). As the second example demonstrates, stage topics need not be overt. So when there is a spatiotemporal location that can act as the stage topic of a sentence, the bare plural may be focused and hence receive an existential rather than a generic interpretation.

This explains the possibility of an existential reading of (105a) and (105b). But what prevents the bare plural from acting as a topic and therefore receiving a generic interpretation in these sentences? This is blocked pragmatically, since a reading according to which dogs in general have the property of playing on my lawn (right now) is not plausibly true. If it is blocked pragmatically, we would expect there to be felicitous sentences with episodic predicates and quasi-universally interpreted bare plural subjects where the bare plural is the topic. This is indeed the case. Consider some

27 The assumption that all sentences have a topic follows from the characterisation of a topic as being the entity or set of entities that a sentence is about. When a sentence has no topic (for example because the expression purporting to denote the topic fails to refer to anything) then that sentence is about nothing.

more examples with bare plurals and episodic predicates, along with paraphrases of their generic and existential readings:[28]

(106) Bears are preparing for winter.

a) It is true of bears in general that they are preparing for winter. *generic*

b) There are some bears which are preparing for winter. *existential*

(107) Students receive their exam results today.

a) It is true of students in general that they receive their exam results today. *generic*

b) There are some students who receive their exam results today. *existential*

(108) Dinosaurs roamed the Pioneer Valley (Diesing, 1988).

a) It is true of dinosaurs in general that roamed the Pioneer Valley. *generic*

b) There are some dinosaurs which roamed the Pioneer Valley. *existential*

In each case, the bare plural can have either an existential or a generic interpretation; like (93), the sentences may be disambiguated by the prosodic focus of an utterance: when focused, the existential reading of the bare plural is preferred. The availability of generic readings of bare plural sentences with episodic predicates suggests that the lack of a generic reading of (105a) and (105b) can indeed be put down to the implausibility of those readings.

3.6.4.2 *Individual-level predicates*

Another generalisation that it would be good predict is that bare plural sentences with individual-level predicates such as 'are intelligent' or 'are tall' should have only a generic reading. I propose to explain this in the following way. It is typically understood that individual-level predicates either do not have an event argument (Kratzer, 1995), or that their event argument is of a sort which has no spatial location

28 The quasi-universal readings of these sentences are arguably more like the functional readings discussed in §3.5.5 than typical generics. Whether these examples count as true generics is not relevant. What matters is that the bare plural NP gets a quasi-universal interpretation, which is the case for both generic and functional bare plurals.

(Maienborn, 2005). If that is right, then there is no entity which can play the role of spatiotemporal stage topic for a sentence with only an individual-level predicate. Assuming, following Erteschik-Shir (1998), that all sentences have a topic, the bare plural must denote the topic, ruling out an existential reading. This prediction is borne out, since existential readings of bare plurals are not available for sentences with individual-level predicates:

(109) a) Dogs are intelligent. *generic only*

b) Basketball players are tall. *generic only*

On the other hand, there are some bare plural sentences with predicates which have both an episodic and an individual-level interpretation. On the episodic reading of the predicate, it is possible for there to be a stage topic, allowing an existential reading of the bare plural only on that interpretation:

(110) Sharks are visible.

a) There are some sharks which are visible. *existential*

b) In general, sharks are visible. *generic*

(111) Ancient figures are carved on the walls of this cave.

a) There are some ancient figures carved on the walls of this cave. *existential*

b) In general, ancient figures are carved on the walls of this cave. *generic*

These sentences, which are originally due to Kiss (1998), can be understood as describing what is the case at some time or location (i.e. as having a stage topic). On this reading, the bare plural can be interpreted as an existential, since it does not denote the topic of the sentence. On the generic reading, the topics are sharks and ancient figures, respectively, and the sentences are not understood as making an assertion about some particular spatiotemporal location, but about their topics.[29] That it correctly predicts the two readings of sentences with predicates which are ambiguous in this way is another advantage of TPRT.

29 See also de Swart (2002) for a similar explanation of the availability of existential readings of bare plurals.

3.7 NON-MAXIMALITY AND HOMOGENEITY

I now turn to the question of non-maximality: why do bare plural generics, analysed as plural predications, admit exceptions? I mentioned in §3.3.7 that the variable tolerance of exceptions was something that unifies both generic bare plurals and definite plurals. In this section, I will set out the theory of Križ and Spector (2021), which attempts to explain both non-maximality and its correlate homogeneity in terms of the underspecification of plural predication. Since I have analysed bare plural generics as cases of plural predication, this explanation naturally extends to the generic case.

3.7.1 *Interest-relativity*

One of the hallmarks of the phenomenon of non-maximality is that it is interest-relative: which exceptions an assertion of a plural predication admits depends on the purposes of conversation. If it is right that the exception-tolerance of generics arises from the same source as the non-maximality of other plural sentences, this would go some way to explaining why giving a semantic theory for generics has proven so difficult. I will begin by showing that non-maximality with plural definite descriptions is interest-relative. I then go on to show how the non-maximality of bare plural generics is interest-relative in a similar way.

3.7.1.1 *Non-maximality for plural definite descriptions*

Although theorists differ significantly in the analysis of the sources of non-maximality and homogeneity, what is widely accepted is that the range of non-maximal readings of a sentence available in a context depends at least in part on the interests of the conversational participants (cf. Bar-Lev, 2021: 1075). To see how, consider the following pair of scenarios from Malamud (2012: 19):

(112) a) *Context:* Mary has a large house with over a dozen windows in different rooms. She locks up and leaves to go on a road trip with her friend Max, forgetting to close just a few of the many windows in various rooms. A few minutes into the ride, Max says, "There is a thunderstorm coming. Is the house going to be OK?" Mary replies:

b) Oh my, we have to go back – the windows are open!

c) *Interpreted as:* At least some of the windows are open.

(113) a) *Context:* Mary has a large house with over a dozen windows in different rooms. She is about to have the window frames painted, and is preparing the house for the arrival of painters. Her friend Max says, "I think you still have a ton of work to do, right?" Mary replies:

b) Actually, I'm done – the windows are open!

c) *Interpreted as:* Every window is open.

Here, we have two utterances of the same sentence-type in (112b) and (113b). The sentence receives a weak existential reading when it matters whether at least one window is open, but a strong universal reading when it matters whether every window is open. It is therefore a desideratum on any theory of non-maximality that it allows for the exceptions that a sentence tolerates to vary depending on the interests of the participants in a conversation.[30]

3.7.1.2 *Non-maximality for bare plural generics*

I will now give some examples which demonstrate that the purposes of conversation can affect truth value judgments for bare plural generics. The formula for generating the examples is to vary the immediate interests of the participants in the conversation, so that different standards are in play.[31]

30 I am aware that some speakers find the sentences in non-maximal scenarios either unacceptable or degraded. The possibility of non-maximal readings is well-attested in the literature (Dowty, 1987; Yoon, 1996; Krifka, 1996; Brisson, 1998; Lasersohn, 1999; Burnett, 2016; Križ, 2016; Križ & Spector, 2021; Horn, 2021) and there is some empirical evidence that sentences with definite plurals are judged true in non-maximal scenarios (Schwarz, 2009; Augurzky et al., 2023).

31 For some more discussion and examples of the variable tolerance of exceptions of generics, see Nickel (2008), McConnell-Ginet (2012), Sterken (2015a), Nguyen (2020) and Bosse (2021).

7-INCH RECORDS

(114) a) *Context:* A website has a page explaining the different sizes and speeds at which vinyl records play. The overwhelming majority of 7-inch records play at 45 rpm, for reasons to do with audio fidelity, although there are many 7-inch records which play at 33 rpm.

b) 7-inch records play at 45 rpm. *judged true*

(115) a) *Context:* A record shop owner is sorting a large stack of vinyl records into two piles, one for records which play at 33 rpm and one for records which play at 45 rpm. She is instructing her assistant, who is not familiar with records, how to sort the records.

b) 7-inch records play at 45 rpm. *judged false*

In both of these cases, the sentence can be understood as a generic about 7-inch records in general. In the first case, it seems sufficient that it is typical of 7-inch records to play at 45 rpm for the (114b) to be true. But in (115), the context of utterance raises the stakes, since sorting the records according to size is likely to lead the assistant to make an error, by placing a 7-inch record which plays at 33 rpm in the 45 rpm pile. So in that context, (115b) sounds false, since it is not sufficient that it is typical of 7-inch records play at 45 rpm for the sentence to be true.

DOLPHINS

(116) a) *Context:* A website has a page stating distinctive and unusual facts about animals. Most dolphins do not use puffer fish to get high, but there is documented evidence that some do.

b) Dolphins use toxic puffer fish to get high.[32] *judged true*

(117) a) *Context:* Uttered during a conversation explaining the differences between typical dolphins and typical whales.

b) Dolphins use toxic puffer fish to get high. *judged false*

32 https://factanimal.com/animal-facts/

In (116), the context seems not to require that the activity be typical of dolphins, since it is presented as an 'unusual fact' about dolphins. In that context, the sentence sounds true, since it is sufficient that at least some dolphins exhibit the behaviour. In (117), however, the sentence is presented as a statement about what is typical for dolphins, in which case it is judged false, as the majority of dolphins do not exhibit this behaviour.[33]

Now that we have established that the non-maximality of both bare plural generics and definite plurals varies depending on the purposes of conversation, I will consider how to account for this in the semantics. I will suggest that plural predication is underspecified (Schwarzschild, 1994; Krifka, 1996; Lasersohn, 1999; Malamud, 2012; Križ, 2016; Champollion et al., 2019; Križ & Spector, 2021) and that this can account both for the homogeneity paradigm and the availability of non-maximal readings of plural sentences.

3.7.1.3 *Non-maximal noun phrases?*

I begin by ruling out one possible source of non-maximality. Non-maximality cannot be a result of restriction on the reference of NPs. If it were, then a sentence with one plural NP and a conjunctive predicate could not tolerate different exceptions relative to each predicate. Consider the following example (adapted from Brisson, 2003: 140):

(118) a) *Context:* There are eleven girls. Ten played a game of basketball, while one, Annie, watched from the sideline. Then nine of the basketball players and Annie went swimming while the remaining basketball player, Bella, watched from the shore.

b) The girls played basketball and went swimming.

(118b) is judged true in this situation even though the girls who played basketball are not the same as the girls who went swimming; but the conjoined verb phrases have just one NP as an argument, so if non-maximality was a result of reference restriction

33 What is the relationship between this example and the striking property generics discussed by Leslie (2008), such as 'Mosquitoes carry the West Nile virus'? Certainly the property is 'striking' in Leslie's sense, but, contra Leslie, I have argued that the sentence is not true *simpliciter*, since whether it is true depends on the purposes of conversation. For a similar point, see Sterken (2015c: 2500).

on NPs, the sentence should be judged false, since the group of girls all of whom played basketball did not all go swimming (and vice versa). Moreover, when 'all' is added to the second conjunct, it is not possible to interpret the second conjunct as holding of only a restricted group of girls:

(119) The girls played basketball and then all went swimming.

That is, (119) is false when ten of the eleven girls played basketball and went swimming (Križ, 2015: 74-75). This provides strong evidence that non-maximality is not a result of a restriction on the reference of NPs.

The same reasoning applies to generics:

(120) a) Dogs have four legs and are good companions.
b) Dogs have four legs and are all good companions.

(120a) is true, even though presumably there are some dogs which have four legs but are not good companions and some which are good companions but do not have four legs. (120b) requires that all dogs are good companions, no matter how many dogs have four legs. It is not enough that all the dogs that have four legs are good companions, if there are other dogs which do not have four legs and are not good companions. We should therefore conclude that, for both definite plurals and bare plural generics, non-maximality is not the result of a restriction on the reference of the NP.

3.7.2 *The underspecification of plural predication*

Instead, I will pursue the idea the meaning of some plural predicates is underspecified. Consider the following case (adapted from Schwarzschild, 1994), in which there are two girls, Annie and Bella:

(121) a) *Situation 1:* Both Annie and Bella have brown hair.
b) *Situation 2:* Neither Annie nor Bella have brown hair.
c) *Situation 3:* Annie has brown hair and Bella does not have brown hair.

(122) a) The girls have brown hair.

b) The girls do not have brown hair.

(122a) is clearly true in Situation 1 and (122b) is clearly true in Situation 2. But what about should we say about these sentences with respect to Situation 3? It seems as though neither sentence has anything to say about this kind of situation.

The idea that I will explore here, following Križ and Spector (2021), is that what causes both homogeneity and non-maximality is the fact that (some) plural predicates are underspecified with respect to the proportion of a plural argument that must satisfy the predicate in order for the argument to satisfy the predicate as a whole. When a homogeneous plural predicate – i.e. one which exhibits the homogeneity paradigm, cf. §3.3.7 – is applied to an argument, there are many possible propositions that the sentence can express, each corresponding to a different way of selecting some of the members of the plural argument and applying the predicate to those members. These propositions are the *Candidate Interpretations* for the sentence. The proposition expressed by a plural predication is true iff all of its Candidate Interpretations are true.

The homogeneity effect arises because the Candidate Interpretations of a plural predication ϕ embedded under negation are the negations of the Candidate Interpretations of ϕ. So $\neg\phi$ is true iff all the Candidate Interpretations of ϕ are false, leaving a gap between the conditions under which ϕ and $\neg\phi$ are true.

Non-maximality arises because not every Candidate Interpretation is relevant for the purposes of conversation; these propositions may be 'filtered out'. The effect of this is that the strongest Candidate Interpretations, i.e. those that say that the plurality as a whole satisfies the predicate, may not enter into the calculation of the truth-conditions of the sentence, leading to non-maximal readings.

Before I present the details, I should make it clear that my proposal does not stand and fall with Križ and Spector's theory. The most prominent competing account assigns definite plurals (and hence, on my proposal, bare plurals) a weak existential semantics and derives their quasi-universal meaning by means of a scalar implicature (Magri, 2012, 2014; Bar-Lev, 2021; Chierchia, 2022). On this account, homogeneity

arises because no relevant implicatures arise when a definite plural is embedded within a downward entailing environment, so it gets interpreted as an existential. One distinctive prediction of this account is that there should be no non-maximal readings of definite plurals in downward entailing environments; this is because, in downward entailing environments, the trigger for both the quasi-universal meaning of the definite plural and the availability of non-maximal readings – the calculation of maximality implicatures – is not present, since the basic semantics of the definite plural give the strong negative existential meaning (cf. Bar-Lev, 2021: 1051). But non-maximal readings do seem to be available for definite plurals in downward entailing environments:

(123) a) *Context:* A class of students has been set a list of books to read. The expectation is that they read all of the books.
b) None of the students read the books.
c) *Predicted interpretation:* No student has read any book.
d) *Can be interpreted as:* No student read every book.

(124) a) *Context:* As part of her film club, Hannah is supposed to watch all the films on a list.
b) Hannah hasn't watched the films.
c) *Predicted interpretation:* Hannah was not watched any film.
d) *Can be interpreted as:* Hannah has not watched every film.

The availability of non-maximal readings of downward entailing environments seems like good evidence against this account. This is not conclusive, however, but I will now proceed to set out Križ and Spector's account.

3.7.3 *Homogeneity*

On the theory of Križ and Spector (2021) (K&S), the source of homogeneity and non-maximality is the predicate; plural terms, such as definite plurals, have a simple, maximal referential meaning and are not ambiguous or underspecified. The underspe-

cification of the predicate results in plural sentences being associated with a range of propositions – *Candidate Interpretations* (CIs) – which are built up compositionally from the members of the plurality denoted by the plural argument. I will use the notation $[\![\alpha]\!]^{g,c}_{CI}$ to denote the set of candidate interpretations for α relative to a variable assignment g. What is communicated by a sentence ends up being the conjunction (in set-theoretic terms, intersection) of all its CIs, i.e. $[\![\alpha]\!]^{g,c} = \bigcap [\![\alpha]\!]^{g,c}_{CI}$ (this will be weakened to take account of non-maximality; some CIs will be filtered out).[34]

CI-TRUTH A sentence is true iff all of its Candidate Interpretations are true.

CI-FALSITY A sentence is false iff all of its Candidate Interpretations are false.

Clearly, this leaves the possibility of a gap between the conditions under which a sentence is true and false, i.e. where some but not all the CIs are true.

What are the CIs of a sentence with an underspecified plural predicate? Begin with a set $\mathscr{P}art(xx)$, which is the set of all pluralities (including individuals) which are among xx. Then, for a plural predication $\ulcorner xx \text{ are } P \urcorner$, there is a CI for each subset S of $\mathscr{P}art(xx)$ which is generated by forming a disjunction each of whose disjuncts is of the form $\ulcorner P(yy) \urcorner$, where yy is either a member of S or a plurality formed by conjoining a member of S with some member of $\mathscr{P}art(xx)$.[35] Formally:

$$[\![P(xx)]\!]^{g,c}_{CI} = \{\bigvee P(yy) \mid yy \in S \wedge S \subseteq \mathscr{P}art(xx) \wedge \forall zz(zz \in \mathscr{P}art(xx) \rightarrow (\exists ss(ss \in S \wedge ss \preceq zz) \rightarrow zz \in S))\}$$

Why are the CIs *disjunctions*? The answer, which will become clearer in the discussion of non-maximality, is that a speaker need not be committed to which objects are the exceptions to their predication; they may be ignorant or indifferent (Malamud, 2012: 19). So CIs should permit a speaker to be indifferent as to which objects satisfy

34 As a historical note, Stephen Read (2009) has discussed the notion of plural signification proposed by the medieval logician Thomas Bradwardine to resolve the Liar paradox. Plural signification is the notion that sentences may express more than one proposition. The resulting theory of truth takes a sentence to be true (false) iff every proposition it expresses is true (false), leaving the possibility of a gap when a sentence expresses at least one true and one false proposition, closely mirroring the suggestion by Križ and Spector (2021). For further discussion of this theory of truth, see Dutilh Noaves (2008).

35 Strictly speaking, the CIs are propositions, but for ease of exposition I just list the sentences here.

the predicate, hence the CIs are disjunctions of propositions each of which applies the predicate to some thing(s) among a plurality.

Let's consider a toy example. Imagine there are two red balls, a and b.

(125) The balls are red.

$\mathscr{P}art([\![\text{the balls}]\!]^{g,c}) = \{a, b, a+b\}$. Then the procedure for generating CIs is as follows, for each subset S of that set:

(i) $S = \{a, a+b\} : R(a) \vee R(a+b)$

(ii) $S = \{b, a+b\} : R(b) \vee R(a+b)$

(iii) $S = \{a, b, a+b\} : R(a) \vee R(b) \vee R(a+b)$

(iv) $S = \{a+b\} : R(a+b)$

This generates the following set of CIs:

$$[\![(125)]\!]^{g,c}_{CI} = \{R(a) \vee R(a+b), R(b) \vee R(a+b), R(a) \vee R(b) \vee R(a+b), R(a+b)\}$$

When R is distributive, the conjunction of these CIs is $R(a+b)$; that is, if all CIs are taken into account, (125) is true iff a and b are both red. The homogeneity paradigm, where the truth-conditions of (126) require that *neither* a nor b is red, is explained by the fact that the CIs for (126) are just the negations of the CIs of (125):

(126) The balls are not red.

(i) $S = \{a, a+b\} : \neg(R(a) \vee R(a+b))$

(ii) $S = \{b, a+b\} : \neg(R(b) \vee R(a+b))$

(iii) $S = \{a, b, a+b\} : \neg(R(a) \vee R(b) \vee R(a+b))$

(iv) $S = \{a+b\} : \neg R(a+b)$

This generates the following set of CIs:

$$[\![(126)]\!]^{g,c}_{CI} = \{\neg(R(a) \vee R(a+b)), \neg(R(b) \vee R(a+b)), \neg(R(a) \vee R(b) \vee R(a+b)), \neg R(a+b)\}$$

The conjunction of this set of CIs is the proposition that neither *a* nor *b* is red, generating the negative existential truth-conditions characteristic of the homogeneity paradigm.

3.7.4 *Non-maximality*

So far we have explained the homogeneity effect for plural predication, but we still predict strong universal truth-conditions in positive environments. To explain non-maximality, K&S propose that some of the CIs get filtered out for reasons of relevance. For example, in Malamud's example (112), in which what matters is whether or not there is at least one window which is open, CIs which entail that more than one window is open are not relevant for the purposes of conversation, so those CIs are filtered out.

K&S define a notion of *strong relevance* to an issue. Taking an issue to be a partition on the common ground (a set of possible worlds), strong relevance is defined as follows (K&S: 1145):

STRONG RELEVANCE A proposition p is *Strongly Relevant* to a partition I iff $\exists X(X \subset I \wedge p = \bigcup X)$.

Or: a proposition is strongly relevant to a partition iff it espouses the contour of one of its cells. The idea is that speakers care about establishing whether the actual world is in one cell of a partition or another; they don't care which world it is among the worlds in that cell. So they treat all the worlds in that cell as equivalent. CIs (i.e. propositions) that do not help to establish which cell the actual world is in are not strongly relevant in this sense.

Propositions which are not strongly relevant are 'filtered out' when calculating the proposition expressed by a sentence. So we revise CI-TRUTH and CI-FALSITY to take into account Strongly Relevant Candidate Interpretations (SRCIs):

SRCI-TRUTH A sentence is true iff all of its Strongly Relevant Candidate Interpretations are true.

SRCI-FALSITY A sentence is false iff all of its Strongly Relevant Candidate Interpretations are false.

Returning to (125), imagine that what matters is whether at least one of the balls are red. Then there will be a partition of worlds I_1 consisting of two cells, C_1, which contains all the worlds in which at least one ball is red and C_2 which contains all the worlds in which no balls are red. The purpose of the conversation is to decide whether the actual world is in C_1 or C_2. Since the proposition that both balls are red is more informative than is required to establish which cell the actual world is in, it is not strongly relevant and so is filtered out. That leaves the following strongly relevant CIs:

$$[\![(125)]\!]^{g,c}_{\text{SRCI}} = \{R(a) \vee R(a+b), R(b) \vee R(a+b), R(a) \vee R(b) \vee R(a+b)\}$$

$\bigcap [\![(125)]\!]^{g,c}_{SRCI} = [\![(125)]\!]^{g,c} = R(a) \vee R(b) \vee R(a+b)$. When R is distributive, this is equivalent to: $R(a) \vee R(b)$. So (125) is true, relative to I_1 iff at least one of a or b is red.

3.7.5 *Applying Križ and Spector (2021) to bare plural generics*

Križ (2015: 185-187) considers the possibility of applying a theory of this kind to bare plural generics, based in part on the homogeneity paradigm exhibited in §3.3.7 by both bare plurals and definite plurals.

One immediate objection to this suggestion is that non-maximality seems to be much more pervasive for generics than it is for definite plurals. It is, in general, much easier to access non-maximal readings of generics than it is for definite plurals. However, this is easily explained by a feature of non-maximal readings in general: they are more easily accessible when the number of objects under consideration is higher (Dalrymple et al., 1994; Brisson, 1998; Križ & Spector, 2021). Since generic bare plurals, unlike definite plurals, typically refer to very large pluralities of objects, it follows that non-maximal readings should be more easily accessible for generics.[36]

36 A corresponding prediction is difficult to verify: non-maximal readings of generics which refer to a very small number of objects should be harder to access. Perhaps 'Astronauts who have set foot on the moon

3.7.5.1 *Non-maximality and alternative-sensitivity*

Another problem is that generics exhibit exception-tolerance in a systematic way which seems to be sensitive to salient alternatives to the predicate (Cohen, 1999b; Kirkpatrick, 2023, 2024a). For example:

(127) a) Swallows lay eggs. ≈ Swallows that produce offspring lay eggs.
b) Lions have manes. ≈ Lions with sexual ornamentation have manes.

In (127a), the generic is not falsified by male swallows which do not lay eggs, because they are not even *candidate* egg-layers, that is, they do not satisfy any member of the set of alternatives to the predicate 'lay eggs', a set which contains alternative methods of producing offspring, e.g. giving birth to live young, undergoing mitosis, laying eggs.[37] Moreover, although necessarily all the swallows that lay eggs are female, the following entailment is not valid:

(128) Swallows lay eggs. ⇏ Swallows are female.

Although on the surface this sensitivity to alternatives seems quite different to the highly variable non-maximality discussed above, I will argue that it can be explained in the same way. First, consider an alternative explanation, according to which the restriction is achieved through domain restriction. The domain of the bare plural in (127a) would then be restricted to candidate egg-layers. One problem with this proposal is that it would not explain why (129) is false.

(129) All swallows lay eggs.

I suggested in §3.5.7 that the restrictor argument of 'all' is a full DP, in this case containing the null determiner, so that generalisations with non-partitive 'all' should be subject to the same constraints on domain restriction as bare plural generics. But (129) can only mean that all swallows of any sort lay eggs. This suggests that the

are American' is one such example, since intuitively this would be falsified by a single non-American astronaut walking on the moon.

37 Laying eggs is trivially an alternative to laying eggs. For discussion of the generation of alternatives for generics, see Cohen (1999b: Chapter 5); for the formal semantics of the generation of alternatives, see Rooth (1992).

tolerance of exceptions to (127a) which are not candidate egg-layers is not a result of domain restriction, since, if it were, we would expect to observe the same effect in (129).[38]

Second, the alternative-sensitivity of generics is mediated by factors relating to discourse and is more flexible than an initial glance at the examples in (127) suggests. Cohen (1999b: 110) gives the minimal pair (130) to demonstrate the effect of prosodic focus on the availability of alternatives (see Rooth, 1992).

(130) a) Veterinarians shoot [livestock]$_F$. ≈ Veterinarians that shoot something shoot livestock.

b) Veterinarians [shoot]$_F$ livestock.≈ Veterinarians that do something to livestock shoot it.

(131) a) Swallows lay [eggs]$_F$. ≈ Swallows that lay something lay eggs.

b) Swallows [lay]$_F$ eggs. ≈ Swallows that do something to eggs lay them.

c) Swallows [lay eggs]$_F$. ≈ Swallows that produce offspring lay eggs.

We can explain the different focus structures in (130) and (131) in terms of the questions that they can felicitously answer. For example:

(132) a) What do swallows lay?

b) Swallows lay [eggs]$_F$.

(133) a) What do swallows do to eggs?

b) Swallows [lay]$_F$ eggs.

(134) a) How do swallows produce offspring?

b) Swallows [lay eggs]$_F$.

What this pattern suggests is that the non-maximal interpretations of the bare plural generics in (127) uttered out of the blue involves the accommodation of a question-

38 One confound here is that, according to an empirical study by Leslie et al. (2011), sentences like (129) are judged true by many people (although this effect was not replicated in Lazaridou-Chatzigoga and Stockall, 2013). It is difficult to say whether informants are interpreting these generalisations as implicitly restricted to candidate egg-layers (see Lazaridou-Chatzigoga et al., 2019 for some discussion), but these judgments are plausibly explained by the employment of heuristics by informants, leading to errors (see Jönsson and Hampton, 2006, 2012; Hampton, 2012 for the empirical picture, as well as Chapter 5 in this thesis for relevant theoretical discussion).

under-discussion. A question-under-discussion is a way representing the purposes of discourse (Roberts, 1996), and on the proposed theory of non-maximality, we predict that exceptions which make no difference for the purposes of conversation can safely be ignored. That prediction is borne out: if we are trying to establish how swallows produce offspring, we are only interested in the swallows that produce offspring in some way, so swallows which are not candidate egg-layers are ignored.

Therefore this kind of systematic non-maximality is predicted by a theory which takes the availability of non-maximal readings to be modulated by the purposes of conversation. The focus-sensitivity of generics comes as a consequence of the role that focus plays in structuring discourse with respect to the question-under-discussion. I conclude that K&S's theory of non-maximality can explain why bare plural generics exhibit alternative-sensitive tolerance of exceptions.

3.7.5.2 *Underdeterminacy in the candidates*

Recent research into the semantics of generics has stressed the possibility of underdetermination in their meaning (Sterken, 2015a; Nguyen, 2020). According to these views, there are many propositions that a generic sentence-type can express and it may not be fully determinate on any given utterance which proposition is in fact expressed. What matters, as far as successful communication goes, is that a speaker and hearer are able to coordinate on a suitable range of propositions. In many contexts, it may not matter *precisely* what a speaker intends to convey by an utterance of a generic, so the fact that their communicative intentions are not made clear in a context is not a barrier to communication.

On a given utterance of a generic, context may fail to determine a unique set of SRCIs for the sentence. In these cases, its truth-conditions will be underspecified, since this domain of the universal quantifier in SCRI-TRUTH is underspecified. This result is welcome: although in many cases it may be clear that a speaker intends to convey a restriction to normal instances, in others they may have no particular restriction in mind. This is probably the case with the expression of many stereotypical beliefs:

(135) a) Snakes are slimy.

b) Lawyers are sneaky.

c) Italians love pasta.

d) Benefit recipients are lazy.

If pressed on their utterances, it seems unlikely that a speaker would say that they had some particular restriction in mind, although they could certainly claim that they did not mean, for example, that *all* Italians love pasta. In these cases there is no single set of SRCIs expressed by the utterances, so their truth-conditions are underdetermined.

The possibility of underdeterminacy offered by this explanation of non-maximality is a virtue, since the ability of speakers to use generics even when they have no particular generalisation in mind is a distinctive feature of the use of generics (Bosse, 2021) and this feature may help to explain why speakers of generics are able to shift between different interpretations of a generic (McKeever & Sterken, 2021).

3.7.6 *Homogeneity and failures of Generic Excluded Middle*

I conclude my discussion of non-maximality and homogeneity by considering a recent argument against simple predicational theories, which turns on the question of whether they predict the homogeneity paradigm. Kirkpatrick (2025) argues that a semantic theory for generics should fail to validate the following principle:

GENERIC EXCLUDED MIDDLE (GEM): For any bare plural characterising sentence of the schematic form ⌜Fs are G⌝, the sentence ⌜Either Fs are G or Fs are not G⌝ is true.

He motivates this desideratum using the following examples where GEM fails:

(136) a) Books are paperbacks or books are not paperbacks.

b) Fair coins land heads or fair coins do not land heads.

c) Lions are male or lions are not male.

Kirkpatrick argues that simple kind-predication theories cannot explain why GEM fails: if a kind term F either satisfies a predicate G or it does not, either ⌜G(F-kind)⌝ or

⌜¬G(F-kind)⌝ should be true. It will be helpful to verify that the present theory does not validate GEM.

Following SRCI-TRUTH, the first disjunct of (136a) is true iff all of its SRCIs are true. The second disjunct is true iff all of the SRCIs of the first are false. Since it is possible that some of the SRCIs of the first disjunct are true and some false, it is possible that neither disjunct is true. This explains why generics fail to validate GEM according to TPRT.

3.8 MODALITY: A MISSING ELEMENT?

One issue which has been absent from my discussion so far is modality. This may strike readers as surprising, since the orthodox view of generics holds that their truth depends on what is the case in non-actual possible worlds, typically reflected in the semantics of 'GEN' as a quantifier over possible worlds, situations or objects. The argument in favour of the view that generics are modals is based on the possibility of true generics with no supporting instances (Pelletier & Asher, 1997: 1133):

(137) a) Kim handles the mail from Antarctica.
b) Members of this club help each other out in an emergency.

The idea is that these sentences can be true even if no mail ever arrives from Antarctica or there never are any emergencies involving members of the club. Yet, the argument goes, both are true, so generics cannot depend only on what is actually the case, otherwise they would be false.

My response to this argument is that it is consistent with the modality being introduced entirely by the predicate. The examples which are cited in favour of a modal relational theory involve habitual or dispositional predicates. Whether something satisfies one of those predicates depends not only on what happens in the actual world, but on what would happen in non-actual circumstances too. For example, for it to be true of some members of a club that they help each other out in an emergency, it must be the case that if an emergency arose, they would help each other out. But all

this shows is that whether some actual objects satisfy that predicate depends on what they would do in non-actual situations.

A proponent of a relational theory might reply that it is the presence of 'GEN' in the sentence which explains this modality, so this would actually constitute evidence in favour of their view (Krifka et al., 1995). However, there is good reason to think that the modality of habituals cannot be reduced to a relational operator, so a commitment to the relational operator 'GEN' is independent of a commitment to the modality of habituals and dispositionals.

3.8.1 *Habitual predicates*

Rimmel (2004) presents some compelling evidence that simple habituals like (137a) do not involve 'GEN'. She contrasts the following pair:

(138) a) Kim drinks beer. *habitual reading available*
b) #Kim drinks a beer. *habitual reading unavailable*

Out of the blue, (138b) has no acceptable habitual interpretation, as it suggests that there is a single beer which Kim habitually drinks. Applying the relational theory to these habituals, the following analyses should be available, where C is a contextually salient predicate of situations (e.g.: 'is a situation suitable for drinking beer'):

(139) a) $gen(s) : C(s)[\text{Kim drinks beer in } s]$
b) $gen(s) : C(s)[\exists x(\text{beer}(x) \wedge \text{Kim drinks } x \text{ in } s]$

Both of these are intelligible, yet (139b) is not a possible meaning for (138b). Rimmel argues that this is because the indefinite must take wide scope in the sentence, so the only way of making sense of (138b) has a habitual is for there to be a single beer which Kim regularly drinks. Evidence for the indefinite taking wide scope is the fact that bare habituals are fine with indefinite objects when the action is repeatable:

(140) Kim drives a Toyota. $\approx$ There is a Toyota that Kim drives (regularly).

In addition, there can be no general prohibition on indefinites taking what would, on the relational theory, be narrow scope with respect to 'GEN':

(141) Conservatives have a favourite politician. ≈ (33)

a) For each conservative, it generally holds that there is a favourite politician who is their favourite.

b) $gen(x)$: conservative$(x)[\exists y(\text{favourite.politician.of}(y)(x) \wedge \text{has}(x)(y))]$

As Rimmel argues, a monadic habitual operator, *hab*, which operates on verb phrases, would explain these linguistic facts better than the relational generic operator, since there is no possibility of a scope interaction between *hab* and an indefinite (see also Dobrovie-Sorin, 2001 for some further linguistic evidence of a monadic habitual operator and Boneh, 2019 for a proposal for a monadic dispositional operator).

Since the focus of this paper is not the meaning of habituals, I won't give a precise semantics for *hab*, although see Boneh and Doron (2013) and Ferreira (2016) for proposals which require that the event argument of *hab* is non-atomic. Informally, we can think of *hab* as transforming a predicate P of individual events into a predicate of states which holds iff P is true of a sufficient number of events across a sufficient period of time. *hab* therefore turns predicates of non-stative eventualities into predicates of states (Smith, 1991; Vlach, 1993) and is related to the meaning of imperfective aspect (Ferreira, 2016). Since habitual meaning has something to do with what an object *would* do under appropriate circumstances, the meaning of *hab* must be modal (see Boneh and Doron, 2013 for more discussion of the modal aspects of habituality).

The addition of the unary predicate modifier *hab* explains why many generics elicit the intuition that they depend not only on how things are in the actual world, but how they would be under certain counterfactual circumstances, without conceding that this intuition must be explained through the presence of the relational operator 'GEN'. The standard argument for the modality of generics therefore fails to motivate the existence of 'GEN'.

3.8.2 *Three-legged lions*

I will briefly address one other argument that is intended to show that generics exhibit modality. Nickel claims, regarding (142):

> Not only is [(142)] true, it would still be true even if every single lion were to have some number of legs other than four, perhaps by some combination of mutation, accident, combat, and the machinations of a madman (Nickel, 2016: 16).

(142) Lions have four legs.

Curiously, Nickel suggests that in such a situation, (142) would be vacuously true if we treated it as an extensional universal generalisation. But (142) would just be false in that situation, since there are counterexamples: lions which do not have four legs. He seems to have in mind the idea that (142) could be true even if the madman executed all the lions in the world, so that there were no more lions.

To explain Nickel's observation that (142) seems true in the situation he describes, it is sufficient to note that on the current theory bare plural generics involve plural predication over a domain of objects not only from the present, but from the past and future too. So when the madman executes the lions there are still some objects in the domain that satisfy the description 'lions', i.e. those in the past (although in this case it would probably be better to say 'Lions *had* four legs') and, presumably, the future. Similarly, if the madman amputated one leg from every actual lion, (142) would still be true when every SRCI is true, i.e. if enough lions among the plurality of lions past, present and future have four legs for the purposes of conversation.

Further evidence in favour of this analysis is that in a world in which there never have been any lions and never will be, (142) seems infelicitous (cf. Greenberg, 2003: 259). What matters is whether there are some things that satisfy the bare plural description *at some time*, rather than at the time of utterance. This is reflected in the existence condition in the domain condition of the null determiner.

Nickel's argument therefore fails to establish that generics have what he calls "modal import". Instead, it shows that generics are not just about what is the case at the time of utterance. That is accounted for in the current proposal by the temporal extension of the domain of the null determiner, so that a bare plural may refer to objects not only at the moment of utterance but to objects in the past and the future too.[39]

3.9 KIND-LEVEL PREDICATES IN THE TEMPORAL PLURAL REFERENCE THEORY

In this final section, I turn to an ontological question which arises from the hypothesis that generic bare plurals denote pluralities. How can pluralities satisfy kind-level predicates, such as 'are extinct'? It is natural to hold that pluralities are not the sort of thing that can be extinct, but kinds are, which seems to militate against my account. That goes along with an ontological commitment: if a sentence of the form '⌜Fs are G⌝', where G is a predicate that can be true only of kinds, is true, '⌜Fs⌝' must refer to a kind (e.g. Zamparelli, 2002: 309; Liebesman & Sterken, 2021: 4).

I will argue here that we need not, as far as the compositional semantics of generics go, think of kinds as a special sort of object, distinct from pluralities of individuals. I will remain neutral on the ontological question of whether kinds just are pluralities of objects.

It is tempting to think that because certain predicates appear only to permit kinds as their argument, it will be necessary to make a linguistic distinction between kinds and ordinary individuals (see, for example Carlson, 1977, 1982). I will show that apparently kind-level predicates can be analysed as collective predicates and so can be true of pluralities.[40] I do so by showing that apparently distinctive features of kind-level predicates are paralleled by 'numerous'-type collective predicates. Finally, I consider two hard cases, 'be extinct' and 'be invented', which pose a problem for the hypothesis that their bare plural arguments refer to pluralities of actual objects, concluding that both are consistent with TPRT.

39 See Chapter 4 for more discussion of this point.

40 I will continue to use the term 'kind-level' predicate as a label for predicates which are usually thought to only to be true of kinds.

Explicit arguments in favour of treating kinds as a special sort of individual are rare, but appeals to intuitions about 'what sorts of objects' can have certain properties are used to motivate the introduction of kinds as objects of reference. For instance, Teichman says of the sentence '*Star Wars* fans are numerous': "It is a flat out category error to think that a *Star Wars* fan could be (or fail to be) numerous" (2023: 1640), using this intuition to motivate the view that predicates like 'be numerous' are defined only for arguments which denote kinds, rather than ordinary individuals (for a similar argument, see Ritchie & Knobe, 2020: 144-145).

This kind of argument is unconvincing unless we think that collective predicates whose arguments are plural terms, such as plural demonstratives, definite descriptions or lists, involve ontological commitments beyond individuals. There is no reason to posit new objects to act as the holders of collective predicates (and there are good reasons not to, cf. Moltmann, 2016), so this argument fails to establish that bare plurals must refer to objects which are ontologically distinct from their members.[41]

3.9.1 *The gather-numerous opposition*

I now turn to the linguistic analysis of kind-level predicates. There is a puzzle as to why certain collective predicates are compatible with 'all', whereas others are not (Kroch, 1974; Dowty, 1987; Winter, 2001). For example:

(143) a) (All) the students gathered in the hallway.
b) (#All) the students are numerous.

Collective predicates which are compatible with 'all' are therefore known as 'gather'-type (G-type) predicates; those which are not are known as 'numerous'-type (N-type) predicates. One immediate caveat here is that N-type predicates may be compatible

41 On Link's (1983) mereological analysis of plurals, plural expressions do refer to 'new' objects, i.e. mereological sums. These objects are ontologically dependent on their atomic parts and, on the extensional theory of mereology endorsed by Link and most others (a notable exception being Moltmann, 1997), mereological sums are individuated by their members (for the view that mereology is ontologically innocent, see Lewis, 1991). Kinds, on the other hand, are not ontologically dependent on their members, since they could have had no members, and are not individuated by their members, since co-extensive but distinct kinds are possible. I have not endorsed Link's theory anyway, as I see no advantage in the additional (although lightweight) ontological commitments of the mereological analysis of plurals.

with 'all' when it is possible to distribute the predicate over its argument. This happens when the argument of the predicate refers to a plurality of pluralities or groups, such as 'armies' or 'sets', which may themselves be numerous.

(144) a) All the enemy armies are numerous. ≈ Every enemy army is numerous (Kroch, 1974: 194).

b) All the sets of axioms are inconsistent. ≈ Every set of axioms is inconsistent.[42]

So 'all' forces a distributive reading, paraphrased with 'every', in N-type predicates if such a reading is possible, otherwise it is unacceptable. But with G-type predicates, adding 'all' is not equivalent to the distributive reading, even when its argument is a plurality of groups. (145a) has a reading which is not equivalent to (145b) (Winter, 2001: 203):

(145) a) All the committees met in Room C. ≉

b) Every committee met in Room C.

(145b) can only report the actions of each individual committee, whereas (145a) has a reading which says that there was one large meeting of all the committees. So for G-type predicates, adding 'all' is not equivalent to the distributive reading, since every committee could meet on its own, without it being true that all the committees met (together).

The *gather-numerous* opposition also corresponds to a difference in *homogeneity*. G-type predicates are homogenous, N-type predicates are not (cf. Križ, 2015: 23).

(146) HOMOGENEITY

a) The students gathered in the hallway. ≈ All the students gathered in the hallway.

b) The students didn't gather in the hallway. ≈ None of the students gathered in the hallway.

42 The observation that 'are inconsistent' (on its external 'with each other' reading) is N-type is due to Kuhn (2020).

(147) HOMOGENEITY FAILURE

The students are not numerous. ≈ It is false that the students are numerous.

3.9.2 *Kind-level predicates are N-type collective predicates*

My contention is that the class of predicates usually called 'kind-level' are in fact just N-type collective predicates, some of which have lexical selection restrictions to species, and so can be true of pluralities. To establish this conclusion I will first show that kind-level predicates do indeed pattern with N-type predicates in the relevant ways. Then, I will go on to discuss why the view that these predicates can only be true of kinds is so widespread, before going into some more detail on some hard cases in order to bolster the case that these predicates can indeed be true of pluralities in place of kinds.

The most obvious example of a putative kind-level predicate which is also an N-type collective predicate is, of course, 'numerous', but since there is some controversy as to whether instance-distribution predicates like 'numerous' are genuine kind-level predicates (see Krifka et al., 1995: 95), let's use a less controversial pair of examples to show that they pattern with N-type predicates. First, kind-level predicates are incompatible with 'all', except on a sub-kind reading, equivalent to 'all types of F are G', analogous to the sub-group distributivity in (144):

(148) a) All dinosaurs are extinct.
b) *Can only mean:* Every species of dinosaur is extinct.

(149) a) All pick-up trucks come in four sizes.
b) *Can only mean:* Every model of pick-up comes in four sizes.

Second, kind-level predicates exhibit homogeneity failure:

(150) a) Dodos are not extinct. ≈ It is false that dodos are extinct.
b) Pick-up trucks do not come in four sizes ≈ It is false that pick-up trucks come in four sizes.

So 'be extinct' and 'come in four sizes' behave exactly as we would expect them to with respect to 'all', homogeneity and distribution over sub-groups if they were N-type collective predicates.

Both are stative predicates, as are many other apparently kind-level predicates, such as 'be widespread', 'be rare', 'be common', 'be abundant' or 'be scarce'. However, other kind-level predicates are collective achievement predicates: 'be invented', 'be introduced into…', 'reach…'.[43] These predicates also pattern with N-type predicates:

(151) a) (?All) black rats reached Australia by ship (acceptable only on a strict distributive or taxonomic reading).[44]

b) Black rats did not reach Australia by ship. ≈ It is false that black rats reached Australia by ship.

(152) a) (?All) chillis were introduced into India by the Portuguese (acceptable only on a strict distributive or taxonomic reading).

b) Chillis were not introduced into India by the Portuguese. ≈ It is false that chillis were introduced into India by the Portuguese.

Given TPRT, how could it be true that black rats reached Australia by ship when the vast majority of black rats did not arrive in Australia by ship? This is an example of 'team credit', which is common for collective achievement predicates, even when most of the members of the argument of the predicate did not participate in the achievement. For example:

(153) a) The soldiers surrendered.

b) The Americans reached the moon first.

c) The Republicans voted the proposal through.

It would be a mistake to aim for a general analysis of the relationship between the parts and the whole which would explain under what conditions a collective

43 The categories of state and achievement relate to the lexical aspect of the predicate and are properties of the eventualities that the predicate describes (Vendler, 1957). The observation that collective achievement predicates are incompatible with 'all' is due to Taub (1989), following the discussion in Dowty (1987). See also Brisson (2003) for further analysis and discussion of counterexamples.

44 The acceptability of the strict distributive reading is due to the fact that these collective achievement predicates may also be true of individuals.

achievement predicate may be true of a plural argument; instead the relationship seems to be driven by pragmatic considerations. For example, in (153c) for the Republicans to have voted for the proposal is for some sufficiently large proportion of Republicans to have voted for the proposal. So, with regards to (151a), we may say that the property of reaching Australia by ship can be held by black rats in general, on behalf of the particular black rats that made a sea voyage.

3.9.3 *Explaining homogeneity failure*

It is reasonable to ask why apparently kind-level predicates and other N-type predicates exhibit homogeneity failure. We can locate the explanation for this failure in the relationship between the pluralities that a predicate is true of and their members. Non-homogenous predicates are true of their arguments as integrated wholes, whereas homogenous predicates are true of their arguments in virtue of the nature of their members (distributive predicates being the homogenous predicates *par excellence*). Looking beyond the domain of plural predication, many singular count nouns, such as 'ball', 'boy', or 'table', are non-homogenous, since they are true of an object as an integrated whole (Löbner, 2000: 237; Moltmann, 1997: 2). And their pluralised counterparts, 'balls', 'boys' and 'tables', are homogenous because they are true of a plurality in virtue of the nature of the members of that plurality. In the same way, non-homogenous collective state predicates like 'are numerous' or 'are widespread' measure a plurality as an integrated whole and are insensitive to the nature of its members.[45]

Collective achievement predicates, like 'voted the proposal through' or 'reached Australia by ship' are true of pluralities in virtue of the existence of some particular event (e.g. a successful vote on the proposal, an arrival in Australia), but there is no requirement that all of the members of the plurality which holds the achievement must be agents of that event (e.g. there may be Republicans who did not vote for the proposal, rats that were not on the ship). However, there do seem to be some

45 For some recent discussion of this issue, see Champollion (2015, 2017), Kuhn (2020) and Amiraz (2020).

constraints on the achievements that pluralities can have in virtue of their members. For a plurality to hold an achievement in this way, it must have a high degree of unity or integrity (Moltmann, 2016), for example in virtue of belonging to a political party or constituting a species.

(154) a) *Context:* All democrats voted against the proposal, but enough Republicans voted for the proposal that it passed.
b) The Democrats and the Republicans voted the proposal through.

(155) a) *Context:* Black rats were the first mammals to reach Australia and they arrived by ship. Reptiles never reached Australia.
b) Mammals and reptiles reached Australia by ship.

In (154b), the attribution of the collective achievement to both the Democrats and the Republicans seems false or at the very least grossly misleading, even though they did together vote the proposal through. Similarly in (155b), although there are representatives of mammals and reptiles who reached Australia by ship (particular black rats), the sentence is false because the plurality of mammals and reptiles does not have the required unity to hold a collective achievement of this kind.

What seems to cause the failure of homogeneity in these cases is that the plurality holds the achievement as a group, not as a collection of individuals. And this is distinctive of plural predicates which exhibit failures of homogeneity: they are true of arguments as integrated wholes, rather than being true of an argument in virtue of the nature of their members.

3.9.4 *Hard cases*

So far, I have made a linguistic argument that kind-level predicates are N-type collective predicates and that this is sufficient to explain the supposedly distinctive behaviour of kind-level predicates. Now I turn to some ontological questions concerning the nature of the objects that satisfy kind-level predicates according to TPRT.

3.9.4.1 *'Be extinct'*

One worry for the claim that 'be extinct' is a predicate which holds collectively of a plurality is that, by definition, anything which it is true of at time t has no members which exist at t. That raises the question: how could it be true of a plurality of objects if, by its very meaning, there are no individuals who could satisfy it? The answer is simple: if the reference of the bare plural argument of the predicate includes objects from the past and present (and, though not relevant in this case, the future too), it is straightforward to see how a plurality may satisfy the predicate 'be extinct' at t, even if there are no members of the plurality at t .

We are then left with a question about the lexical semantics of 'be extinct'; following Koslicki (1999: 457) and Izumi (2012: 214), some things are (collectively) extinct at time t when there were some of them before t and there are none at t. A defender of kind-level predication might respond: you have missed the fact that 'be extinct' can't be true of just any old group of objects, but only kinds. For example:

(156) #Donny and Danny are extinct.

My response is that a restriction to kinds won't explain the distribution of 'be extinct', since many bare plurals which plausibly denote kinds are not acceptable as its arguments:

(157) a) #Betamaxes are extinct.
b) #Diamonds are extinct.
c) #Black cabs are extinct.

It is not just that these sentences are possibly false: they seem just as bad as (156). But a restriction to kinds can't explain this fact. On a kind-predication theory, the subject NPs in (157) refer to kinds and so the restriction would be met, yet they are still unacceptable. Instead, just as some predicates have a selection restriction for animacy, as in (158) (Cohen, 2020: 122), we can assume that 'be extinct' has a selection restriction for genera, species or sub-species.[46]

46 Perhaps 'be extinct' is felicitous with a wider range of arguments than this; according to the Oxford English Dictionary, it is possible to accurately describe a volcano as extinct, although there may be some

(158) #This rock is sad.

That will be sufficient to explain the infelicity of (156); unlike a restriction to kinds, it also explains the infelicity of (157). So the species-restriction is a better explanation of the linguistic data than a kind-restriction. Thus there is no motivation to view 'be extinct' as a kind-level predicate.

3.9.4.2 *'Be invented'*

Now let's turn to an apparent problem for the proposal that another kind-level predicate, 'be invented', is really a collective predicate. As Liebesman and Sterken (2021) have argued, it is not clear that '*a* was/were invented to...' entails that there ever have been or ever will be any instances of *a* in existence. They give the following example:

> A bicycle company is about to launch a new model called the Pavement Chewer in an attempt to corner the middle-aged cyclist market. In the same launch, they're also going to release a tricycle, called the Pavement Nibbler to try to capture the preschool market. However, they have yet to produce anything – no bicycles, nor tricycles. Now consider:
>
> (159) Pavement Chewers were invented to appeal to middle-aged cyclists.
>
> (160) Pavement Nibblers were invented to appeal to middle-aged cyclists.
>
> (159) is true while (160) is false (Liebesman & Sterken, 2021: 3, numbering amended).

They argue that this shows that both 'Pavement Chewers' and 'Pavement Nibblers' refer to kinds, which need not be realised in the actual world. If these expressions do not have different referents, they ask, how could the two sentences differ in truth value, given that the same predicate applies to both arguments? Liebesman and Sterken conclude that bare plurals do not refer to pluralities, since the two sentences should have the same truth value if the NPs refer to the same thing (i.e. nothing).

motivation to treat this as a case of polysemy or even ambiguity. But if 'be extinct' is not ambiguous, this is further evidence against a type-restriction view of the predicate.

INTENTIONAL OBJECTS To resolve this issue, I propose that the predicate here is *intentional* and takes what Moltmann calls an "intentional object", an object whose existence is dependent on a particular intentional act (Moltmann, 2015: 145). Intentional objects need not be 'real', yet they may play the role of referents or objects in the ontology natural language, for instance:

(161) a) The planned hike was delayed by the rain. In the end it never took place.
b) The woman that John imagined never materialised.
c) The catalogue describes some buildings which don't exist.

In these examples, a predicate of intentionality (planning, imagining, describing) is said to be true of some object which cannot plausibly be a kind and which is consistently asserted not to exist. The predicate 'be invented to' clearly involves intentionality of some kind. And if we need intentional objects in cases where kinds have nothing to do with the interpretation of the sentence, we can extend this analysis to cover Liebesman and Sterken's examples, where some objects are said to have been invented for some purpose yet can consistently be said not to exist.

(162) Pavement Chewers were invented to appeal to middle-aged cyclists, but since the company went bankrupt before any were built, there are no Pavement Chewers.

On Moltmann's proposal, the bare plural NP in (162), 'Pavement Chewers', can be indexed to the intentional act of inventing, so that the existence of some (abstract) Pavement Chewers is dependent on that particular act. The counterexample therefore fails because the two bare plurals in (159) and (160) do not, after all, denote the same objects. Instead, each bare plural refers to different intentional objects, which are dependent on the two distinct intentional acts of invention described in the sentences.

3.10 CONCLUSION

The objectives of this paper have in one sense been quite modest – to establish the viability of TPRT – but in another, ambitious – to overturn the stubborn consensus in favour of relational and kind-predication theories that has dominated this debate for

nearly five decades. I hope to have achieved at least the modest goal and in so doing helped to explain why some of the problems associated with bare plural generics seem so intractable. The mistake that has been made is to suppose that bare plural generics are special, when in fact they are more amenable to semantic analysis when we treat them like any other plural sentences.

There are many issues relating to generics that I have failed to discuss here, most notably the relationship between bare plural generics and other linguistic forms capable of expressing generic meanings, including indefinite and definite singular NPs. That omission has been necessitated by the time required to establish the viability of TPRT within its own limited domain. Unlike relational theories, which can naturally be extended from bare plural generics to indefinite singulars, TPRT has no obvious analogue for indefinite singulars. On my theory, plurality is taken as central to the generic meaning of generic bare plurals, yet it clearly is not involved in the meaning of indefinite singulars. As I have suggested, the task is to deliver compositional semantic theories for each linguistic form which assign these different linguistic forms meanings which do justice to our linguistic intuitions. There is no reason to think that this task must be achieved by proposing that all the distinct linguistic forms capable of expressing generic meanings contain the same phonologically null constituent, 'GEN'. So the task of giving an adequate theory of the meaning of indefinite singular generics remains incomplete.

4

HOW TO MAKE A GENERIC TRUE

4.1 INTRODUCTION

As a matter of fact, every American President in the actual world has been male. Yet the following bare plural generic sentence is not true:[1]

(1) American Presidents are male.

Intuitively, for (1) to be true, there must be some rule in force that states that in order to be an American President, a person must be male. This is the intuition that bare plural generics express nomic or non-accidental generalisations.[2] According to this intuition, generics say not only how things are, but how things would be even if things were different in many respects. This paper will address the question: what about the meaning of generics explains this intuition?

I will argue that the strict truth-conditions of a generic like (1) are satisfied iff the plurality of objects referred to by the bare plural noun phrase satisfy the predicate. This might seem like a non-starter in light of the fact that generics are nomic. But I will argue that nomicity arises because there are many ways that these conditions may be satisfied, some in virtue of facts about the objects that make up that plurality and some in virtue of the properties that those objects have in common in virtue of satisfying the bare plural description (and some in virtue of a mix of the two). It is the facts about the properties the objects have in common which account for the nomicity of generics. To make this claim precise, I will present a novel truthmaker semantics

1 An early version of this paper appeared in the *Proceedings of the 24th Amsterdam Colloquium* as (Ralston, 2024a).

2 As my discussion will focus primarily on bare plural generics, I will drop the 'bare plural' qualification where no confusion will arise.

for plural predication which will allow me to show in concrete terms what it takes for a generic to be true in virtue of facts about the categories that the generic describes.

I begin by discussing the nature of nomicity and some of the linguistic features of generics. Next, I will introduce a novel semantic theory for generics, the *temporal plural reference theory*, according to which generic bare plurals refer to the objects which satisfy a description within a temporally extended domain. I then discuss the different ways that a generic may be made true when it is not determinate which objects will satisfy its bare plural description in the future. Next, I give a truthmaker semantics for plural predication which permits plural predications to be made true even in the face of such indeterminacy. I show that a consequence of the truthmaker semantics in combination with the temporal plural reference theory of generics is that temporally unrestricted generics cannot be made true only by facts about their instances, and that this is sufficient to explain the nomicity of generics. I describe what these facts, facts about categories, are like and show that this explains a contrast between the nomicity of temporally restricted and unrestricted generics. Next, I consider and respond to two objections. In the final section, I show how my analysis of the nomicity of generics helps to explain why generics have been associated with the formation of essentialist beliefs about the nature of the categories they describe.

4.2 PRELIMINARIES

4.2.1 *Nomicity*

I begin with some brief remarks on what I take nomicity to be. I will use 'nomic' as a synonym for 'non-accidental' and 'law-like'. I will generally talk about nomic *sentences*: linguistic entities which, if true, are made true not only by how the world happens to be, but how things would be if things in the world were different, holding fixed some set of rules of laws. The states of affairs in virtue of which these sentences are true (if true at all) have the following characteristic:

COUNTERFACTUALS They support counterfactuals (Goodman, 1983). Suppose there is a law which makes it true that all copper conducts electricity. Then, if I were to connect some particular strip of copper wire to a power source in the right way, it would conduct electricity. That is true, even if I never in fact do such a thing.

As the focus of this paper is the meaning of generics, not the metaphysics of laws, I don't want to take on too many commitments about the nature of these states of affairs. All I will need for my purposes will be *facts about categories*, which I argue are sufficient to explain the nomicity of generics. What I will claim really matters for the nomicity of bare plural generics is that their truth does not depend on which objects turn out to be instances of the bare plural description.

4.2.2 *The meaning of generics*

Now I turn to generics. Linguistically, generics are characterised by the absence of an overt quantifier or adverb of quantification, such as 'most', 'generally', or 'normally'. I will focus here on bare plural *characterising sentences*, such as the following:[3]

(2) a) Ravens are black.
b) Beavers build dams.
c) Cryptocurrencies are digital assets.
d) Logicians are intelligent.

3 Characterising sentences come in other linguistic forms too, such as indefinite singulars and definite singulars:
(i) The lion has a mane.
(ii) A beaver builds dams.
It has been observed that indefinite singular generics exhibit greater nomicity than bare plural generics (Burton-Roberts, 1977; Cohen, 2001; Greenberg, 2003), which raises the question: why not focus on the nomicity of indefinite singular generics too? My response is that, whatever indefinite singular and bare plural generics have in common, they are different constructions and will ultimately require distinct compositional semantics. As Collins (2015) has argued, the linguistic phenomena which fall under the heading of 'generics' are so diverse that it is not reasonable to expect that these sentences having anything in common, as far as the compositional semantics go. It may turn out, as I tentatively suggest in §4.5.2, that different forms of generics can be made true by the same facts, but this should fall out as a consequence of a compositional semantic theory, rather than being hardwired in from the start.

These sentences say something about what is generally true of the category they describe without specifying proportion or frequency.

Although characterising sentences can take a bare plural form, not all bare plurals are characterising sentences. Some bare plurals ascribe a property to a kind which cannot be true of individual members of the kind:

(3) a) Dodos are extinct.
b) Rhinos are rare.
c) Chillis were introduced into India by the Portuguese.
d) Problem gamblers are abundant.

These are known as *kind-predication sentences*. These sentences do not generalise across members of a kind, but say something about a kind as a whole. Characterising sentences, but not kind-predication sentences, give rise to intuitions of nomicity. Because of this, and because I will generally be limiting the scope of this paper to bare plurals, I set kind-predication sentences aside for the purposes of the discussion of the sources of the nomicity of characterising sentences and will henceforth use 'generics' as shorthand for 'bare plural characterising sentences'.

Finally, generics must be distinguished from episodic sentences. Intuitively, episodic sentences say something are about particular objects or situations.

(4) a) Ravens are nearby.
b) Beavers are building dams.
c) Cryptocurrencies are sold here.
d) Problem gamblers are betting.

In each case, the sentence is about some particular object(s) or situation. As my focus is on generics, I will use the following test to distinguish generic bare plurals from existential ones: existentially interpreted bare plurals are upward-entailing, whereas generic ones are not (Krifka et al., 1995).

(5) a) Ravens are black. $\nRightarrow$ Birds are black.

b) Beaver builds dams. ⇏ Mammal builds dams.

(6) a) Ravens are nearby. ⇒ Birds are nearby.

b) Beavers are building dams. ⇒ Mammals are building dams.

Next, I turn to the question of what generics mean.

4.2.2.1 *Non-maximality*

Generics admit exceptions; (2a) is true even though not every raven is black. So they are weaker than unrestricted universal generalisations. I will not discuss this feature of generics in detail here, but I note that this phenomenon is not unique to them: variable admissibility of exceptions is characteristic of plural predication, in particular sentences with plural definite descriptions. This phenomenon is known as *non-maximality* (Brisson, 2003; Malamud, 2012; Križ, 2016; see also Chapter 3 of this thesis). The relevance of non-maximality to nomicity is discussed in §4.4.2.2.

4.2.2.2 *Temporary generalisations*

Generics are not felicitous when used to describe states of affairs which hold only temporarily or accidentally. Cohen illustrated this with his famous example:

(7) a) *Context:* By chance, every current Supreme Court judge has a prime Social Security number.

b) Supreme Court judges have a prime Social Security number (Cohen, 1999a: 35).

As Cohen argues, the sentence sounds false in this situation, since it is not sufficient that every Supreme Court judge just happens to have a prime Social Security number for (7b) to be true. Whenever we construct a suitable context in which, for only a short period of time, every instance of the bare plural description has some property, a generic which ascribes that property to that category will sound false in that context (see also Greenberg, 2003: 140; Nickel, 2016: 16). But as the example (1) suggests, the point is not just that generics can fail to be true in virtue of 100% of the instances of

the category at the time of evaluation satisfy the predicate. In that example, 100% of the American Presidents *up until now* have been male, yet the sentence still does not sound true. What seems to matter is that there is nothing in the rules that constitute what it is to be an American President that prevent a woman from being one.

4.2.2.3 *Temporally restricted generics*

Although some generics can naturally be interpreted as expressing 'timeless' truths, generics can be felicitously restricted to an explicitly bounded period of time (Musan, 1997; Carlson, 1999; Greenberg, 2003).[4]

(8) a) In Roman times, senators wore togas.
b) In the dark ages, poor people went hungry.

Temporally unrestricted generics differ with respect to nomicity when compared with restricted generics. The following minimal pair is inspired by Dahl (1975).

(9) a) Superpowers are referred to by abbreviations the first letter of which is 'U'.
b) In the 1970s, superpowers were referred to by abbreviations the first letter of which is 'U'.

As Dahl argued, (9a) generates a very strong intuition of nomicity and seems to entail the counterfactual: had Andorra were a superpower, Andorra would have been referred to by an abbreviation the first letter of which is 'U', which is clearly false. The unrestricted generic (9a) can only be true in virtue of some non-accidental state of affairs, so even as uttered in the 1970s when every actual superpower (the USA and USSR) was referred to by an abbreviation the first letter of which is 'U', (9a) is false.

(9b), on the other hand, seems consistent with an accidental understanding.[5] To verify this, consider that (9b) is much better with a continuation which asserts that the connection is a mere accident than (9a):

4 As Carlson (1999) showed, generics can be restricted in a wide variety of ways, not just temporally. What matters for my purposes is the possibility of temporal restrictions.

5 I will use the term 'understanding' instead of 'reading' or 'interpretation' to avoid making a commitment to there being a genuine ambiguity between an accidental and nomic understanding of a sentence.

(10) a) Superpowers are referred to by abbreviations the first letter of which is 'U', ?but that is only a historical accident.

b) In the 1970s, superpowers were referred to by abbreviations the first letter of which is 'U', but that was only a historical accident.

(11) a) Supreme Court judges have a prime Social Security number, ?but that is only a historical accident.

b) In the 1970s, Supreme Court judges had a prime Social Security number, but that was only a historical accident.

So it is sufficient that the actual superpowers in the 1970s were referred to by such abbreviations for (9b) to be true. On the other hand, (9b) can *also* be interpreted as nomic. To verify that temporally restricted generics are consistent with a nomic understanding, we see that they may naturally support counterfactuals in the same way as unrestricted generics:

(12) a) In 2020, non-essential workers stayed at home. So, if Bea had been a non-essential worker in 2020, she would have stayed at home.

b) Before 1870, Popes were sovereign of the Papal States. So if Pope Leo had been Pope before 1870, he would have been sovereign of the Papal States.

We can therefore conclude that (i) temporally unrestricted generics are consistent only with a nomic understanding and (ii) temporally restricted generics are consistent with both a nomic and an accidental understanding.[6]

6 One point of confusion should be cleared up here: it has been claimed that there are temporally unrestricted generics which are true in virtue of only statistical connections, such as 'Cars have radios' or 'Barns are red' (Prasada & Dillingham, 2006; Prasada et al., 2013). Those cases might seem like counterexamples to the claim that temporally unrestricted generics admit only a nomic understanding. However, my use of the term 'nomic' is quite weak, since I am using it as a synonym for 'non-accidental'. It is difficult to see how to interpret these putative counterexamples as not supporting the corresponding counterfactuals, e.g. if this bus had been a car, it would have had a radio, in which case they qualify as nomic in my terms.

4.2.3 *Existing theories of generics*

Since my objective in this paper is to establish that a theory that analyses generics in terms of plural predication is viable, I will not devote much time to the well-established alternatives. I will briefly note here how the nomicity of generics is explained by the two main families of theories and raise some issues with these explanations, relating to their treatment of temporally restricted generics.

The dominant view in the literature is that generics have a modal component to their semantics (Dahl, 1975; Pelletier & Asher, 1997; Thakral, 2018), usually contributed by the phonologically null operator, 'GEN' (e.g. Krifka et al., 1995; Eckardt, 1999; Greenberg, 2003; Nickel, 2016; Kirkpatrick, 2023). Simplifying a little, on this view, the logical form of (1) is as follows, where *gen* is the metalanguage translation of 'GEN':

(13) $gen(x)$: american.president(x)[male(x)]

I'll call theories which assign a relational form logical like this to characterising sentences, *relational* theories. The standard view of the semantic contribution of *gen* is that is a modal quantifier over objects in possible worlds or situations. The nomicity of generics falls out as a consequence of this semantics, since the truth-conditions of generics depend on what happens in non-actual worlds or situations.

The major competing theory is the *kind-predication* theory, according to which generics have a subject-predicate logical form and bare plural noun phrases refer to kinds. On the simplest version of the theory, (1) can be analysed as follows (Liebesman, 2011; Liebesman & Magidor, 2025):

(14) male(AMERICAN.PRESIDENT)[7]

The nomicity of generics is explained in terms of the metaphysics of kinds. Kinds are more than just their actual members, so even though all actual American Presidents have been male, it does not follow that the kind AMERICAN.PRESIDENT has the property of being male. Things are a little more complex for 'sophisticated' kind-predication

7 I use CAPITALS for expressions in the metalanguage which denote a kind.

theories, according to which there is a predicate modifier which takes properties of individuals into properties of kinds (Carlson, 1977; Teichman, 2023). But the same basic mechanism accounts for the nomicity of generics, as this modified predicate expresses a kind-level property which is sensitive to instances of kinds in non-actual worlds, so it is possible for all the actual instances of a kind to have a property without the kind itself having the (modified) property.

It is not obvious how either of these approaches would explain the diminished nomicity of temporally restricted generics discussed in the previous section. The problem is that their explanations for why an unrestricted generic like (9a) is nomic apply just as well to restricted generics like (9b). For example, on the relational theory, (9b) is plausibly analysed as follows, where the temporal restriction is added to the restrictor of *gen* (here s is a situation variable and *loc* is relation between situations and times) (Musan, 1997; Carlson, 1999):

(15) $gen(x, s) : loc(s, 1970s) \wedge \text{superpower}(x)(s)[\text{abbreviate.with.'U'}(x)(s)]$

But, because *gen* is a modal quantifier, this gives no explanation of the fact that (9b) more readily admits an accidental understanding than the temporally unrestricted (9a). That is because the sentence still quantifies over possible situations involving superpowers in the 1970s. The effect of the temporal modifier is merely to narrow the domain of quantification. On this analysis, (9b) should require that non-actual superpowers in the 1970s were referred to by abbreviations the first letter of which is 'U'. But it does not, so the analysis is flawed.

Now consider a (simplified) kind-predication analysis of (9b) as uttered at time t, where the past tense is interpreted as an existential quantifier over times and the temporal adverbial restricts the domain of quantification:

(16) $\exists t'(t' < t \wedge loc(t', 1970s) \wedge \text{abbreviate.with.'U'}(\textsc{superpower})(t'))$

Again, this analysis offers no explanation as to why the temporally restricted (9b) easily admits an accidental understanding while the unrestricted (9a) does not. That is because what it takes for the kind SUPERPOWER to have this property at some time

in the 1970s is just the same as what it takes for the kind to have the property now; the temporal restriction should make no difference to nomicity. But it does, so there is no explanation for the asymmetry between (9a) and (9b).

These brief remarks are not knock-down arguments against either analysis. Nothing I have said rules out some more complex explanation of the asymmetry. But there are well-known objections to both relational and kind-predication theories, and my objective is to establish that an alternative semantic theory, based on plural reference, is viable, so I will now move on to my original proposal (for discussion of these criticisms, see Chapter 1).

4.3 THE TEMPORAL PLURAL REFERENCE THEORY OF BARE PLURAL GENERICS

My proposal is to treat generic bare plurals as expressions which refer to more than one object. Unlike versions of this theory according to which generic bare plurals refer to a particular type of object, a *kind*, the version of the referential approach I propose does without with an ontological commitment to kinds or objects in possible worlds other than our own. Call this the *temporal plural reference theory* of generics (TPRT). The motivation for TPRT stems from two methodological principles. First, all things being equal, a semantic theory is better if it avoids positing unarticulated constituents which are not articulated in any known language (i.e. 'GEN') (Liebesman, 2011: 415). Second, all things being equal, a semantic theory is better if it avoids making ontological commitments to new types of objects. In the context of the semantics of generics, these two methodological principles give us a reason to prefer theories which posit neither an unarticulated generic operator nor a domain of kinds.[8] Of course, all things need not be equal, but my task in this paper is to sketch out a semantic theory for generics, TPRT, which fits the linguistic data and respects these two methodological preferences.

8 It is worth noting here that the notion of a kind required for a workable kind-predication theory of generics goes well beyond the notion of a natural (or even a social) kind familiar from metaphysics and the philosophy of science, since almost any noun phrase can be felicitous in a bare plural generic. This means that a kind-predication theory of generics will bear significant ontological commitments over and above those of these philosophical theories of kinds.

The main contribution of this paper will be to show that TPRT can account for the nomicity of generics while respecting both of these principles.

According to plural reference theories, in a generic, ⌜Fs are G⌝, the bare plural ⌜Fs⌝ just refers to the objects which are Fs within some domain.[9] What is that domain? According to one version of a plural reference theory, the domain consists of all possible worlds (Moltmann, 2013). Generic bare plurals would then refer to all the objects in any possible world which satisfy the description – Moltmann calls these *modalized pluralities*. This would explain the nomicity of generics in much the same way as the kind-predication theory: what it takes for a modalized plurality to satisfy a predicate does not only depend on how things are in the actual world.

One problem with this version of the theory is that a generic ⌜Fs are G⌝ becomes equivalent to ⌜The possible Fs are G⌝. For example:

(17) a) Ravens are black.
b) The possible ravens are black.

Moltmann is aware of the obvious differences in meaning between (17a) and (17b) and suggests the following explanation for the differences in meaning:

> Definite NPs involving [*sic*] a single domain of entities, which in the presence of *possible* will be extended to a domain of merely possible entities. By contrast, kind terms [i.e. generic bare plurals] involve plural reference to different individuals in different domains, belonging to different circumstances (Moltmann, 2013: 39).

This is intended to defuse the objection by showing that the two expressions do not really co-refer, which would explain the substitution failure in (17). The problem with Moltmann's reply is that she gives no explanation of what it means for an expression to refer within different domains, rather than to refer within one large domain which includes those domains. This objection can be pressed by asking: what are the semantic values of 'ravens' and 'the possible ravens' in (17) according to the theory? If both refer

9 For reasons of style, in what follows I will use the singular noun 'plurality' as a stand-in for a plural when the plural is awkward. This should not be read as bearing an ontological commitment to plural objects, but merely as a way of speaking.

to the same things, then, no matter how they refer to those things, the two expressions should be substitutable (at least in simple sentences like (17a)).[10]

This version of the plural reference theory also makes the wrong predictions for generics with explicit temporal restrictions, such as those in (8):

(18) a) In Roman times, senators wore togas. = (8a)
b) In Roman times, the possible senators wore togas.

(19) a) In the dark ages, poor people went hungry. = (8b)
b) In the dark ages, the possible poor people went hungry.

The (a)-sentences clearly are not equivalent to the (b)-sentences, since the (a)-sentences are only about the actual instances of the description. The (a)-sentences would not have been true if there were no actual senators/poor people during the periods of time described, whereas the (b)-sentences could be true provided there were some objects which might have been senators or poor people. I conclude that generic bare plurals do not refer to modalized pluralities.

4.3.1 *Plural reference and temporal domains*

In the simplest version of the plural reference theory, generic bare plurals would refer to whatever objects satisfy the description at a particular moment in time. On this view, a model for a bare plural generic would represent a snapshot of how things are at that moment. But, as the discussion in §4.2.2.2 showed, that cannot be right, since generics depend for their truth not only how things are at the time of evaluation, but at other times too. The most minimal modification that we can make to this simple version of the plural reference theory that can accommodate this fact is to require that the domain be extended in its temporal dimension.

10 One way of making sense of Moltmann's distinction is in terms of superplural reference, that is, reference to pluralities of pluralities (see Oliver & Smiley, 2016: 133-139). Then the substitution would be illegitimate, since the two terms would not co-refer because the bare plural in (17a) would refer to a superplurality while the definite description in (17b) would refer to a plurality. However, it is not clear what motivation there would be for treating generic bare plurals as superplurals except to avoid the substitution in (17), so I set this possibility aside.

4.3.1.1 *Temporal indices*

The idea that predicates or determiners can be indexed to times is well-established in linguistic semantics, because of evidence that noun phrases can be interpreted relative to a temporal index which is distinct from the time of evaluation of the main verb of the sentence. The classic example is from Enç (1986: 409):

(20) Every fugitive is now in jail.

For this sentence to be true, it must be evaluated relative to a temporal index in which everything in the extension of 'fugitive' is in jail. But of course nothing is both a fugitive and in jail at the same time. If the noun phrase had to be evaluated at the same temporal index as the main verb, we would predict (20) to be contradictory. But it is not, so we have evidence that the noun phrase is evaluated relative to a distinct (earlier) time from the main verb of the sentence. To account for this, we can take the semantic values of noun phrases to be functions to times to extensions, i.e. to the set of things which satisfy the noun phrase at that time.

4.3.1.2 *An unpronounced determiner*

The proposal here is that generically interpreted bare plurals include a phonologically null determiner very close in meaning to 'the' (for details of the proposal, see Chapter 3). Following proposals by Schwarz (2009) and Elbourne (2013), both 'the' and the null determiner, '$\varnothing_{the}$', take a silent pronoun as an argument which fixes their domain.[11] Formally, the domain is set by an index, i. For an index i, a definite description ⌜The Fs⌝ or a bare plural ⌜$\varnothing_{the}$ Fs⌝ refers to the objects all of which are F at i and which everything that is F at i is among.[12]

(21) a) $[\![\text{the}]\!]^g = \lambda P_{\langle e,st\rangle}.\lambda i_s : \exists xx^{\text{MAX}} P(xx)(i).\imath xx^{\text{MAX}} P(xx)(i)$

b) ${}^{\text{MAX}}P =_{df} \lambda xx_e.\lambda i_s.P(xx)(i) \wedge \forall yy(P(yy)(i) \rightarrow yy \preceq xx)$

Call the index argument of this function the *domain index* of the determiner. Here, I will focus only on the temporal aspects of generics, so the value of the domain

11 The idea that determiners take an index which restricts their domain as an argument is also proposed by von Fintel (1994), Cooper (1996) and Musan (1997).

12 Here I use plural variables, xx, yy and the plural inclusion relation $\preceq$ (see Rayo, 2002; Oliver & Smiley, 2016).

index will be a time. The simplest theory would be to fix the domain once and for all, perhaps requiring that the value of the domain index must always include all of time. But as shown above, generics may be temporally restricted, either because they are about an explicitly restricted period of time – as in (8) – or because there is some implicit restriction:

(22) a) (These days) software developers are highly valued.
b) Police officers wear bodycams.

On the assumption that generic bare plurals contain a covert temporal pronoun, the syntax of (22a) will be as follows:

(23) [S [DP [$\varnothing_{the}$ [NP software developers]] i_1] [VP are highly valued]]

According to TPRT, the DP in (23) refers plurally to anything which is a software developer within the period of time denoted by the index pronoun 'i_1'. In this case, the index pronoun is not bound, so it is assigned a value by the variable assignment, g. But, like other pronouns, it may be be bound (Tonhauser, 2020: 12).

(24) a) Whenever inflation is high, cryptocurrencies are popular.
b) When the sun rises, roosters crow.
c) When a hurricane hits Florida, flying houses are common.

In (24a), the value of the index varies according to distinct times at which inflation is high, so which cryptocurrencies are popular can vary from case to case.[13] When the index is not bound in the structure, its value must be resolved pragmatically (like other pronouns). What is distinctive about the temporal index argument of '$\varnothing_{the}$' is that it takes the broadest possible value in the context, unlike determiners such as 'the' or 'every' which are easily restricted to a very short temporal index. This explains why generic bare plurals, unlike other DPs, resist contextually restricted interpretations:

(25) I am meeting some linguists and philosophers for lunch today…

13 Although their implementation differs from my own, following the line of argument of Stanley and Szabó (2000), the fact that these pronouns exhibit bounds interpretations can be seen as evidence in favour of their syntactic reality.

a) Every linguist is a semanticist. *contextual restriction preferred*

b) The linguists are semanticists. *contextual restriction preferred*

c) ?Linguists are semanticists. *contextual restriction unavailable*

As was already observed by Enç (1986), the value of the index pronoun may include times in the past, present and future. It may be helpful to think of the pronoun as a plural pronoun whose reference is to many (possibly discontinuous) times. Like any referential use of a pronoun, the value of the index should be recoverable in the context. I propose the following generalisation for resolving free temporal indices on generic bare plurals:

RESOLVE T-INDEX Assign a free temporal pronoun i whose sister is a node $\{\varnothing_{the}$, NP$\}$ the most temporally inclusive value which is plausible in the context of utterance.

An index is temporally inclusive to the extent that it includes a longer a period of time. So the most inclusive value for an index will be the one which includes the longest period of time. A consequence of RESOLVE T-INDEX is that for sentences like (7b), (22a) and (22b), for which the only plausible value for the domain index is some relatively short period of time, the bare plurals will be interpreted as being about Supreme Court judges/software developers/police officers only during that period. In contrast, for generics about biological categories like (2a), it is more plausible that the value for the domain index might be all of time, in which case the reference of the generic bare plural, 'ravens', would be the plurality of objects which are ravens at some time. Finally, for a sentence with an explicit temporal restriction like (8a), the reference of the bare plural will be just the plurality of senators within a domain that is temporally located during Roman times. We can also allow that, when the restriction is implicit, the value of the domain index is underdetermined by the context of utterance. This underdetermination won't hamper successful communication, provided that speaker and hearer can agree on some set of propositions which are expressed by the sentence (or, more precisely, on a range of possible values for the domain index); for further discussion of this phenomenon, see Buchanan (2010).

4.3.1.3 *Ruling out temporary generalisations*

Since generics say something about objects which satisfy a description within a temporally extended domain, it immediately follows that temporary accidents like Cohen's Supreme Court judges example in (7b) may not be enough to make a generic true. This is because for that sentence to be true, it is not enough that the *current* Supreme Court judges have prime social security numbers. RESOLVE T-INDEX tells us that the domain of the bare plural should be more inclusive than this, so the sentence can (when uttered out of the blue) only report a regularity which extends through time.

4.4 HOW TO MAKE A GENERIC TRUE

I have argued that generics involve plural reference to objects within a domain which is extended through time and that this explains why they cannot be true in virtue of merely temporary states of affairs. In what follows, I will show that another consequence of this theory is that when a generic contains a bare plural noun phrase in which the domain index of '$\varnothing_{the}$' includes times in the future, an accidental understanding is not possible. I do this by giving an intrinsic truthmaker semantics for plural predication (Linnebo, 2022).

4.4.1 *Generic and instance-based explanations*

The idea is that when the index on the unpronounced determiner extends into the future, it may be indeterminate which plurality a generic bare plural refers to. There are two sources of indeterminacy: first, if it is not determinate what objects there will be in the future, then it will be indeterminate which objects are members of a domain which extends into the future. Second, even if it is completely determinate what objects there will be in the future, it may be indeterminate which objects will have which properties at any given time; for example, it might be indeterminate which

individual will be the President of the USA on 25 February 2039. So there are two possible sources of indeterminacy: (i) the makeup of the domain and (ii) the properties of the members of the domain.

But it seems that even when this indeterminacy is present, a sentence may still be true. In such cases, the nature of the categories described in a sentence might be enough to guarantee that the sentence is true, regardless of which objects are members of the domain of the sentence and what properties they happen to have (Hale, 2020; Linnebo, 2022). For example, the universal generalisation 'All humans are mortal' might be true in virtue of the fact that part of what it is to be human is to be mortal (Fine, 2017: 568-569). This is independent of which objects are in the domain of the universal quantifier and which of those objects are humans, so the truth of the sentence is independent of which things turn out to be humans. Recall (1), repeated below:

(26) American Presidents are male. = (1)

Now consider the application to TPRT: if the domain index on the generic bare plural 'American Presidents' extends into the future it will be indeterminate to which objects it refers. How then, might (26) be true? It could be true in virtue of facts about what it is to be an American President. These facts are independent of facts about who ends up being an American President, so the truth of (26) would not be dependent on who in fact becomes the President. Linnebo calls this kind of explanation a *generic explanation*:

GENERIC EXPLANATION An explanation is *generic* insofar as it... is based on general facts about the properties or operations involved in the claim that is generalized (Linnebo, 2022: 349).

An explanation is *instance-based* to the extent that it is not generic.[14] When a sentence includes an expression whose domain is indeterminate for either of the reasons described above, completely instance-based explanations will not be available.

14 A completely instance-based explanation for a generic ⌜Fs are G⌝ would contain only information about individual Fs.

For example, if it is not determinate which objects there will be, it is not possible for the universal generalisation 'All humans are mortal' to be explained by a conjunctive fact, each of whose conjuncts is a fact about a particular human, since it is not determinate which objects there will be and which of those will be humans. But it may still be explained by the generic fact that to be human is to be mortal, independently of which objects there are and which are humans, so the generalisation may still be true in the absence of an instance-based explanation. Returning to the generic (26), while that sentence cannot be made true by an instance-based explanation, it might be made true by a generic fact about what it is to be the President of the USA which is independent of facts about particular individuals who are or have been the President (although of course it is not, since there is no such fact).

My claim, which is in the spirit of Linnebo's proposal for universal generalisations, is that when the domain of a generic bare plural extends into the future, instance-based explanations are not possible, because it is not determinate which objects are referred to by the bare plural. It follows that sentences with a bare plural whose domain index extends into the future and which are not supported by a generic explanation of some kind are interpreted as not true. This was witnessed by Cohen's (7b) and Dahl's (9a). It can also be seen in the following contrast:

(27) a) American Presidents are male.
b) Until now, American Presidents have been male.

Intuitively, (27b) is true: the fact that the actual American Presidents until now have been male is sufficient for it to be true. By RESOLVE T-INDEX, the index on the null determiner in (27b) cannot extend into the future. There is, therefore, an instance-based explanation for (27b), which is why (27b) is true. In contrast, again by RESOLVE T-INDEX, the temporal index in (27a) must extend into the future (unless there is some very strong contextual cue which makes clear that future Presidents cannot be under consideration). And since it is not determinate which objects are in the extension of 'American Presidents' in the future, no instance-based explanation is possible. Moreover, since there is nothing about the rules of the American Presidency which

would explain why any future American President would be male, there is no generic explanation either. This explains why (27a) is not true.

One clarification: since it is determinate which objects from the past are members of the domain, what really matters is that there be a generic explanation which concerns future members of the category (recall that being a generic explanation is a matter of degree). But typically generic explanations will not discriminate between past, present and future members.

To summarise, this, I argue, is the sense in which generics may be nomic: they can be made true by facts about categories rather than individuals. Facts about categories are 'non-accidental' only in the sense that they do not depend on which objects happen to be members of a category. When there is no such fact, the sentence is interpreted as not true.

4.4.2 *Intrinsic truthmaker semantics*

Next, I will give a truthmaker semantics for plural predication which explains how generics with indeterminate domains can be made true in the absence of an instance-based explanation. The central idea of truthmaker semantics is that the meaning of a sentence is identified with the set of states which make it true (or verify it). The verification relation, $\Vdash$, holds between a state and sentence: $\ulcorner s \Vdash \phi \urcorner$ means that s verifies ϕ. Call the set of states which make a sentence ϕ true the set of ϕ's verifiers.

Given the identification of the meaning of a sentence and the set of its verifiers, it is possible for there to be differences in meaning between two sentences which are true in all and only the same worlds, since they may differ in their verifiers (Fine, 2017). Such distinctions are not possible in standard possible worlds semantics, because, on that approach, two sentences are identical meaning iff they are true in all and only the same worlds. For example, consider two instances of the Law of Excluded Middle, $A \vee \neg A$ and $B \vee \neg B$. These are true in all the same possible worlds (i.e. all worlds), but they differ in their verifiers, since the states that verify the former are verifiers of

either A or $\neg A$, while the states which verify the latter are states which verify B or $\neg B$.

Here, I will adopt a revision of Linnebo's *intrinsic* truthmaker semantics, which he describes as motivated by the following idea: "My target idea is that a state s verifies a statement ϕ just in case material intrinsic to s suffices to explain ϕ, leaving no need to 'look beyond' s to account for the truth of ϕ" (Linnebo, 2022: 357). Unlike in Fine's (2017) truthmaker semantics, the notion of verification here is monotonic, i.e. for any state s which verifies ϕ, any state which is at least as informative as s verifies ϕ. Formally, we begin with a set of states or *state space* S and a weak partial order on S, $\leqslant$. The empty state $\mathbb{0}$ is the bottom of this order – every state is more informative than $\mathbb{0}$ – and the inconsistent state $\mathbb{1}$ is at the top – every sentence is verified by $\mathbb{1}$. Finally, there is a fusion operation on S, $\sqcup$, which conjoins states. It is constrained so that $s \leqslant s'$ iff $s \sqcup s' = s'$.

To state the rules for the meanings of atomic and complex sentences, we need an interpretation function, $[\![\zeta]\!]_\xi$, which is a function from a pair of an n-place predicate P and a state s to the n-tuple of objects in the domain of s which satisfy P according to s.

I will take an existence predicate, E, as a primitive, where $s \Vdash Ea$ iff s includes the information that a exists.[15] We can then define existence for pluralities in terms of this existence predicate (Linnebo, 2022: 370), where $\prec$ is the proper plural inclusion relation, and finally define the *domain* of a state s, $D(s)$ as the set of objects which exist according to s.

(EE)	$s \Vdash Eaa$	iff	if $a \prec aa$, then $s \Vdash a \prec aa \wedge Ea$ and if $\neg a \prec aa$ and $t \Vdash Ea$, then $s \sqcup t \Vdash \neg a \prec aa$
(D)	$aa \in D(s)$	iff	$s \Vdash Eaa$

The semantic clauses for atomic and complex sentences are as follows, where $\bot$ is a sentential constant which is false on every interpretation:

15 This differs from Linnebo's definition of the existence predicate in terms of the identity relation. The reason for this is that I will be using a different clause for predication, which allows predications to be verified by a state even when that state does not include the information that the argument(s) of the predicate exist.

$(\Vdash \mathit{At})$	$s \Vdash P(aa_1, \ldots, aa_n)$	iff	$\langle aa_1, \ldots, aa_n \rangle \in [\![P]\!]_{s \sqcup t}$ for every t such that $aa_1 \in D(t)$ and … and $aa_n \in D(t)$
$(\Vdash \mathbb{1})$	$\mathbb{1} \Vdash \phi$	if	$s \Vdash \bot$, then $s = \mathbb{1}$
$(\Vdash \wedge)$	$s \Vdash \phi \wedge \psi$	iff	$s \Vdash \phi$ and $s \Vdash \psi$
$(\Vdash \vee)$	$s \Vdash \phi \vee \psi$	iff	$s \Vdash \phi$ or $s \Vdash \psi$
$(\Vdash \rightarrow)$	$s \Vdash \phi \rightarrow \psi$	iff	for every t, if $t \Vdash \phi$, then $s \sqcup t \Vdash \psi$
$(\Vdash \neg)$	$s \Vdash \neg\phi$	iff	for every t, if $t \Vdash \phi$ then $s \sqcup t \Vdash \bot$
$(\Vdash \exists)$	$s \Vdash \exists x \phi(x)$	iff	$s \Vdash \phi(a)$ for some $a \in D(s)$
$(\Vdash \forall)$	$s \Vdash \forall x \phi(x)$	iff	$s \sqcup t \Vdash \phi(a)$ for every t and every $a \in D(t)$

These clauses are borrowed from Linnebo (2022: 373, 376), with the exception of $(\Vdash \mathit{At})$. This clause says that a state s verifies an atomic predication iff any way of adding the information that the objects that are the argument of the predicate exist to s yields a state which contains the information that those objects exist and satisfy the predicate. (Since verification is monotonic, if $s \Vdash E(aa)$ then $s \sqcup t \Vdash E(aa)$ for all t.) This allows atomic predications to be made true by states which don't themselves contain the information that the subjects of the predicate exist. This means that the verifiers of Pa may not all verify $\exists x Px$, since $(\Vdash \exists)$ says that for a state to verify an existential generalisation, it must verify that the object that is the witness to the generalisation exists. Moreover, a verifier s for $\exists x P(x)$ need not be such that there is an $a \in D(s)$ such that $a \in [\![P]\!]_s$, since it can be true that $s \Vdash P(a)$ without it being true that $a \in [\![P]\!]_s$.[16]

4.4.2.1 *The nomicity of generics*

To see how $(\Vdash \mathit{At})$ helps to explain the nomicity of generics, notice that a plural predication can be verified by a state s which does not itself verify that the things which are the argument of the predicate exist, provided that any way of adding that

16 These remarks suggest that the correct logic given these semantics is a free logic, since atomic predications can be true even when nothing in the domain satisfies the predicate. In free logic, the classical rule of existential generalisation is replaced with: $A, E(b) \vdash \exists x A[x/b]$, which would be valid on these semantics. However, in contrast to many free logics, the classical rule of universal instantiation, $\forall x A \vdash A[b/x]$, will be valid on these semantics, since a universal generalisation $\forall x \phi(x)$ is made true by a state s iff $\phi(x)$ is true of every object at every extension of s; that is, whether a universal generalisation is made true by a state depends not only on what objects exist at that state but at extensions of it too. Moreover, since the inference $\neg\exists x \neg F(x) \vDash \forall x F(x)$ is not valid (because absence of a counterexample at a state is not sufficient for the truth of the corresponding universal generalisation at that state), a logic weaker than classical logic will be required; Linnebo (2022) suggests that intuitionistic logic will be sufficient for this purpose.

information to s meets the condition in $(\Vdash At)$. To satisfy this condition, s would need to include information about the things that satisfy the description which is independent of which objects they are. That kind of information will be information about the *category* that corresponds to the description, e.g. the category of American Presidents in (26). In that example, there is no fact about the category which would satisfy the relevant instance of $(\Vdash At)$, which explains the intuition that (26) is not true.

Returning to the distinction between generic and instance-based explanations, we can now say that a bare plural generic has an instance-based explanation to the extent that its truthmaker verifies the existence of each instance of the bare plural description. A completely instance-based verifier would include the information that the plurality denoted by the bare plural description exists and that it satisfies the predicate. As I have argued, temporally unrestricted generics do not admit completely instance-based explanations because there are (currently) no states which verify the existence of that plurality. But these generics may still be verified by states which include facts about the category that they describe (e.g. about the American Presidency, the Supreme Court, being a raven) when these states are such that any way of adding the information that the plurality exists yields a state in which the plurality exists and satisfies the predicate. So temporally unrestricted generics may have generic explanations.

4.4.2.2 *The underspecification of plural predication*

So far, I have ignored the important question of exception-tolerance: not every F must be G for ⌜Fs are G⌝ to be true. My explanation for this phenomenon will follow from the hypothesis that generics are instances of plural predication. Following Malamud (2012), Križ (2016) and Križ and Spector (2021), plural predications may be underspecified with respect to the proportion of the argument which must satisfy the predicate for the sentence to be true. Adding 'all' to the sentence eliminates the underspecification.

On this view, there is no fact of the matter as to how many Fs must be G in order for ⌜Fs are G⌝ to be true.[17] But the context of utterance and the purposes of

17 For discussion of the claim that the meaning of generics is underspecified or underdetermined, see Declerck (1986), Sterken (2015a), Nguyen (2020) and Bosse (2021).

conversation can help to narrow down the intended meaning of a sentence, by filtering out unintended interpretations.[18]

With respect to $(\Vdash At)$, this means that a plurality xx may satisfy a (distributive) plural predicate Fs even when not every y among xx individually satisfies F - this is known as a *non-maximal* reading of the predicate. Which exceptions a given utterance of a plural predication tolerates will depend on the purposes of conversation.[19] The availability of non-maximal readings of generics is therefore not an objection to TPRT – in fact the analysis of generics in terms of plural predication predicts that they should exhibit non-maximality.

4.4.3 *Facts about categories*

Next I turn to the nature of the states which would be suitable verifiers for temporally unrestricted generics. As Fine (2017: 557) has observed, it is not strictly speaking necessary, in giving a truthmaker semantics for a fragment of a language, to say exactly what the verifiers are (see also Linnebo, 2022: 457-458). That is a question of metaphysics, not semantics. However, I have argued here that it is in virtue of being made true by certain states that (some) generics exhibit nomicity. The truthmaker clause for atomic predications, $(\Vdash At)$, may not on its own be enough to convince a skeptic that there really are states that could play the role that I need them to.

I have claimed that the right sorts of states are facts about categories. What is a fact about a category? And what is the relationship between facts about categories and facts about individuals? First, given that plural predications admit exceptions, the relevant facts need not be ones which are true of all possible members of a category, although they may be (as, for example, would be the case for 'Horses are mammals'). Second, I will distinguish between two types of facts about categories, both of which can make a generic true: (i) those which explain why an instance of the category

18 An additional benefit of this approach is that it would explain why generics and other plural sentences, but not universal generalisations, exhibit what is known as *homogeneity*, i.e. it is not always that the case that either Fs are G or Fs are not G (Križ, 2016). See Kirkpatrick (2025) for discussion of homogeneity in generics. I discuss this point in detail in Chapter 3.

19 For discussion of this phenomenon for other plural expressions, see references above, especially Malamud (2012) and Križ (2016) as well as Bar-Lev (2021).

belongs to the category, and (ii) those which follow from an object's being a member of some category. Since the former stem from what determines whether some object is a member of a category, call these *determinant facts*. Since the latter concern conclusions what follows from an object being a member of a category, call these *consequent facts*.

4.4.3.1 *Determinant facts*

When an object is a member of a category, there is something in virtue of which it is a member of that category. A determinant fact about a category concerns the conditions that an object must satisfy in order to be a member of the category. So, for example, it is a fact about the category of American Presidents that a person must be aged 35 or above to become a member of that category, since the rules for membership of the category state that this the case, even though it is metaphysically possible for a person to become the President of the USA aged 34 (that is, the rules might have been different). Determinant facts depend on what makes something a member of a category, which may itself depend on the way that membership of a category has been explicitly defined. For biological categories, the relevant facts will be ones about the genotype of the organism and its relation to the phenotype. That is, what makes it the case that ravens are black is the fact that to be a raven is to have certain genetic characteristics which are expressed in such a way that ravens appear black.

4.4.3.2 *Consequent facts*

In addition to facts about what determines whether something is a member of a category, there are facts about what *follows* from something being a member of a category. These are consequent facts. Consider the following example:

(28) Single-use plastics bags sold by retailers in the UK with more than 250 full-time equivalent employees cost at least 10p.

There is a law in force in the UK which states that if something is a single-use plastic bag sold by a retailer with more than 250 full-time equivalent employees, that thing must cost at least 10p (with certain exceptions). This is not a fact about

what determines whether an object falls into that category, but about what follows from being a member of that category, without containing on any information about particular members.

Compared to determinant facts, the direction of explanation is reversed, since it is being a member of the category which explains why the object has some property, rather than having some property which (partially) explains why some object is a member of the category.

To see the distinction, imagine that there is a rule stating that, in order to become a Supreme Court judge, a person must have been assigned a prime Social Security number. Then it would be a determinant fact about Supreme Court judges that they have a prime Social Security number. But if there were a rule stating that on becoming a Supreme Court judge, a person's Social Security number was amended to the nearest available prime number, it would be a consequent fact about Supreme Court judges that they have a prime Social Security number.

4.4.3.3 *Relations between universals?*

One point which bears emphasising here is that facts about categories need not be facts about relations between universals, in the sense of Armstrong (1983) (although they could be). Armstrong argued that laws should be analysed in terms of one universal necessitating another; that is, it is a law that Fs are G iff N(F,G), where N is the relation of necessitation that holds between universals (Armstrong, 1983: 78). However, the sorts of facts I am discussing need not be associated with universals, since generics are felicitous with noun phrases which are clearly not associated with universals. This was the case for (28): there surely is no universal relating to the category of single-use plastics bags sold by retailers in the UK with more than 250 full-time equivalent employees.

To summarise: both determinant and consequent facts about categories are able to verify generics. Returning to the 'bad' generics (26), Cohen's (7b) and Dahl's (9a) we can see that they are not supported by either kind of fact about the categories of American Presidents, Supreme Court judges and superpowers respectively, which is

why they are not true. So we have explained the nomicity of generics in terms of the states of affairs that make them true, rather than in the nature of the generalisations they express.

4.4.4 *Temporally restricted generics*

We are now in a position to explain the puzzling contrast between between temporally restricted and unrestricted generics introduced in §4.2.2.3. Recall that temporally restricted generics elicit diminished intuitions of nomicity, since they are compatible with both accidental and nomic understandings. This is explained as follows: when a generic is temporally restricted to a period of time which occurs wholly in the past, the domain of the bare plural will be determinate and so an instance-based explanation will be possible. But a generic explanation is also possible, since it might be the case that facts about categories are sufficient to verify the predication, even though an instance-based explanation is also available.

To verify this, consider again the minimal pair repeated from (9).

(29) a) Superpowers are referred to by abbreviations the first letter of which is 'U'.
b) In the 1970s, superpowers were referred to by abbreviations the first letter of which is 'U'.

As Dahl observed concerning (29a), the temporally unrestricted generic asserted in 1975 suggests that being referred to by an abbreviation the first letter of which is 'U' was a prerequisite for being a superpower (Dahl, 1975: 109). Yet it is straightforward to understand the temporally restricted (29b) as an accidental generalisation, according to which it was simply true of the superpowers in the 1970s that they were referred to by such an abbreviation. There is a state which verifies the existence of the USA and USSR and which also verifies that they were referred to by abbreviations the first letter of which is 'U'. So (29b) admits a completely instance-based explanation because it is temporally restricted.

(29b) is also consistent with a nomic understanding. That is, one could understand (29b) as being true in virtue of a generic explanation – a fact about the category of superpowers. If there were such a fact, it would support the corresponding counterfactual: if China had been a superpower in the 1970s, it would have been referred to by an abbreviation the first letter of which is 'U'. Of course, there is no such fact. But since (29b) can be made true by either a generic or an instance-based explanation, it is still acceptable even though we know that there is no generic explanation.

Finally, consider (30), which has a temporal restriction wholly located in the future.

(30) In the 2050s, superpowers will be referred to by abbreviations the first letter of which is 'U'.

It is extremely difficult to understand (30) as accidental because the domain of the bare plural is located entirely in the future. Again, this follows from TPRT, since it is indeterminate which objects will be superpowers in the 2050s and so an instance-based explanation of (30) is not possible. This is further evidence that it is the temporal location of the domain which determines whether an accidental understanding of a generic is possible.

I conclude that generics with temporal restrictions in the past can be made true by either generic or instance-based explanations, whereas unrestricted generics or generics with temporal restrictions in the future cannot be made true by instance-based explanations. The availability of both generic and instance-based explanations explains why it is possible to interpret temporally restricted generics about the past as either nomic or accidental. The impossibility of an instance-based explanation explains the lack of an accidental understanding for temporally unrestricted generics.

4.5 TWO OBJECTIONS

4.5.1 *What if determinism is true?*

I have assumed throughout that it is not determinate which objects there will be in the future and what properties they will have. That explained why generics whose

domains extend into the future are not compatible with instance-based explanations and hence accidental understandings. But what if that assumption is incorrect and the future is completely determined by the current facts and the laws of nature? That is, what if *determinism* is true? (Similar remarks will apply if eternalism, the view that everything that exists at any time exists eternally, is true.) It might initially seem as though my explanation of the nomicity of generics hangs on the question of the determinacy of the future, since, if determinism is true, then there is no metaphysical barrier to an instance-based explanation for a generic whose domain extends into the future. It would appear, then, that if determinism is true, instance-based explanations of temporally unrestricted generics are available and hence it should be possible to understand them as accidental.

What really matters for my account, however, is not whether determinism is true. That is because even if it is in fact determinate what objects there will be and what properties they will have, in our ordinary linguistic practices we do not behave as though any person could have access to that information. That is, we *behave* as though instance-based explanations for temporally unrestricted generics are impossible, even if it turns out that they are.

If the temporally unrestricted (2a) does indeed have an instance-based explanation – consisting of a fact or facts about a particular plurality of ravens – I could hardly be acquainted with it now, since many ravens are located in the future. Perhaps Laplace's demon could make accurate predictions about the properties of individual ravens in the future, based on its knowledge of the past and present, but in practice such knowledge is not attainable.

(31) Ravens are black. = (2a)

If determinism is true, the domain of any generic bare plural is determinate and an instance-based explanation is always available. But it is not in practice possible to learn about the objects referred to by a bare plural as particulars. So even if instance-based explanations for generics are metaphysically possible, they are, as far as ordinary users of language are concerned, unknowable.

So how could a person be in a position to know (31)? This would be possible only if they possess knowledge of a generic explanation for (31), that is, a fact about what it is to be a raven. So even if it is metaphysically possible for (31) to be made true by an instance-based explanation, we would interpret a speaker of (31) as intending to communicate that it has a generic explanation.

I conclude that even if determinism is true, generics whose domains extend into the unknowable future will be understood as nomic, since instance-based explanations will not be epistemically accessible. That is independent of whether a generic really does have an instance-based explanation.

4.5.2 *Ambiguity or underspecification?*

The opposition between generic and instance-based explanation roughly corresponds to Carlson's (1995) distinction between two approaches to the semantics of generics:

RULES-AND-REGULATIONS (R&R) A generic is true in virtue of the existence of some structure, such as a rule or a regulation, in the world.

INDUCTIVE A generic is true once 'enough' instances have been observed.

Carlson suggests that a generic like 'Bishops move diagonally' is easily explained on the former approach – it is true in virtue of the rules of chess – while a generic like 'Dogs bark' is explained by the latter – it is true in virtue of enough instances of dogs barking. 'Rules-and-regulations' is perhaps not the best way of describing the first approach, however, since the structure which makes a generic true need not be anything as formal as a rule or a regulation, although these would be paradigm cases.

Cohen (2001) argues that these two approaches actually correspond to truth-conditionally distinct *readings* of generics (see also Guerrini, 2024). Cohen observes that there appear to be two types of generalisation that generics can express. This is based on Lawler's (1973) classic minimal pairs:

(32) a) Madrigals are polyphonic.

b) A madrigal is polyphonic.

(33) a) Madrigals are popular.

b) ?A madrigal is popular. *only acceptable as existential*

The idea is that being polyphonic is an 'essential' property of madrigals, whereas being popular is not (that is, something would not be a madrigal if it were not polyphonic, but it could be a madrigal if it were not popular).[20] The generalisation is that indefinite singulars like the (b)-sentences cannot admit a generic reading when the predicate is understood as being somehow accidental to the subject, whereas bare plural sentences like the (a)-sentences can receive a generic reading in that context.

Cohen argues that a bare plural generic like (33a) is ambiguous between a reading which can be true even when there is nothing more than a statistical connection between the subject and the predicate (the inductivist reading) and a reading which can be true only if there is a rule in effect which requires that things be a certain way (the R&R reading) (Cohen, 2001: 195). According to Cohen, the R&R reading, but not the inductivist reading, has normative force. Indefinite singular generics can only express the R&R reading, whereas bare plural generics may express either. That explains the more limited availability of generic readings of indefinite singulars. This is a competitor to my analysis, since I have identified (at least) two ways that a generic might be true, but claimed that these do not correspond to truth-conditionally distinct readings.

Many of Cohen's observations are sound. In particular, he shows that by manipulating context so that it is plausible that there is some nomic regularity that supports their truth, indefinite singular sentences which only have an existential reading when uttered out of the blue can get a felicitous generic reading (Cohen, 2001: 196-197). However, if he is right that generics have two truth-conditionally distinct readings, we would expect to be able to truthfully assert conjunctions like the following (Zwicky & Sadock, 1975):

20 I am informed by Ian Rumfitt that it is not in fact true that madrigals must be polyphonic. Instead, what is distinctive of madrigals is that they are secular, rather than religious. But for the sake of the example, we shall ignore this detail.

(34) Madrigals are popular but it is not the case that madrigals are popular.

I can see no way of constructing a context where this sentence is not false (see also Sterken, 2015a: 7). To avoid the impression that (34) sounds bad because it has two conjuncts that sound like the negations of each other, on Cohen's analysis we should also be able to get a true reading of the following:

(35) Madrigals are polyphonic but it is not the case that a madrigal is polyphonic.

As with (34), there seems to be no true reading of this sentence. According to Cohen, a true reading should be possible if the first conjunct has the inductivist reading and the second the R&R reading. That would be true when it is accidentally true that madrigals are generally polyphonic, but there is no rule or regulation requiring this to be the case. Such a situation is clearly possible, but this sentence sounds as bad as (34) (unless the second conjunct is read as an existential). Contrast that with a case of lexical ambiguity:

(36) That's a bank, but it isn't a bank.

This sentence has a true reading on which the first instance of 'bank' refers to the bank of a river and the second a financial institution which is good evidence of a genuine ambiguity.

My explanation of these facts is similar to Cohen's except that I do not posit an ambiguity. Generics in principle admit either instance-based or generic explanations (or a blend between the two) and hence can be understood as either accidental or non-accidental. But since they are not ambiguous, we would, correctly, expect them to fail the ambiguity tests above since there is no way to make both conjuncts true at the same time.

4.5.3 *Interim summary*

I conclude this part of the paper with a brief summary of the proposal so far. Generics are instances of plural predication. Generic bare plural noun phrases refer to the

objects which satisfy a description within some temporally extended domain. Generics may in principle be made true by generic or instance-based explanations, although completely instance-based explanations are not available when their domain extends into the future. When no instance-based explanation for a generic is available, it will not be understood as expressing an accidental generalisation. And so generics for which no instance-based explanation is possible and no generic explanation is (in fact) available are judged not true.

4.6 GENERICS IN PSYCHOLOGY

In this final section, I will show how my proposals can help to make sense of some of the empirical evidence that generics are associated with the formation of psychological essentialist beliefs about social kinds. This issue has been widely discussed because of a correlation between essentialist beliefs and a range of intergroup biases.

I begin with a brief summary of the empirical evidence, before going on to give an account of how inferences which it is natural for hearers of generics to make can lead them to form beliefs about categories, rather than individuals. These beliefs have typically been interpreted in the empirical literature as essentialist beliefs, but I will argue that only a weaker interpretation is warranted. Finally, I argue that this same reasoning applies equally to any generalisation which cannot have an instance-based explanation, which suggests that the focus in the literature on generics at the expense of other linguistic constructions is a mistake.

4.6.1 *Generics and essentialism*

Psychological essentialism is the belief that a kind has some "hidden, nonobvious, and persistent property or underlying nature shared by members of that kind that causally grounds their common properties and dispositions" (Leslie, 2017: 405) (from here on, I drop the qualifier 'psychological'). According to recent empirical research, exposure to even the most apparently innocuous generic language about novel kinds – sentences

like 'Zarpies have striped hair' – is associated with the formation of essentialist beliefs about those kinds (Gelman et al., 2010; Cimpian and Markman, 2011; Rhodes et al., 2012; Leshin et al., 2021; Benitez et al., 2022; see also Rhodes et al., 2025 for a recent opinionated summary of this evidence).[21]

Of particular interest are essentialist beliefs about social groups, including those based on sexual or gender identification, religion, class or nationality. Essentialist beliefs about social groups have been found to be correlated with stereotype endorsement (Bastian & Haslam, 2006) and been blamed as a trigger for people to hold prejudiced beliefs about those kinds (Allport, 1954), which may in turn have detrimental effects on inter-group relations (Keller, 2005; Rhodes et al., 2018).

Theoretical work has built on the empirical evidence of the essentialising effects of generic language, suggesting that we should avoid using generics about social groups, because of the potential of our utterances to lead people to essentialise those groups (Haslanger, 2012; Langton et al., 2012; Leslie, 2014, 2017; Wodak et al., 2015; Wodak & Leslie, 2017). This prescription has even been extended to cover "counterstereotypic" generics, such as 'Girls are good at math' (Rhodes et al., 2025: 8; Benitez et al., 2025), on the basis that even these well-intentioned generics can promote essentialism.

4.6.1.1 *Objections*

A number of objections have been raised to the conclusion that hearing generics causes the formation of essentialist beliefs.[22] First, it has been observed that there is no direct link between hearing generics and developing essentialist beliefs, because belief in a hidden essence is just one type of belief that hearing generics about a kind can promote. Noyes and Keil (2019, 2020) found evidence that both children and adults could distinguish inherent or biological explanations from structural or social ones when hearing generics. The production of essentialist beliefs was dependent on the framing of the generic. Other studies found that the production of essentialist beliefs was lower when the sentence was heard in the context of a structural or extrinsic

21 For some general discussions of psychological essentialism, see Neufeld (2022), Ritchie and Knobe (2020) and Rhodes and Moty (2020).

22 For a recent critical discussion of the evidence, see Neufeld et al. (2025).

frame (Vasilyeva et al., 2018; Vasilyeva & Lombrozo, 2020; Vasilyeva et al., 2020). This suggests that the relationship between hearing generics and forming essentialist beliefs is at best indirect.

There is also a concern that the empirical studies cited above classify too broad a range of responses as evidence of 'essentialist' beliefs. They don't distinguish between beliefs which merely generalise across members of a kind, as opposed to essentialising the kind, so the evidence in those studies is consistent with participants in those studies holding beliefs about kinds that fall short of genuine essentialist beliefs (Hoicka et al., 2021).

Saul (2023) argues that the evidence that generic language *in particular*, as opposed to general language or language indicating high-proportions, causes people to hold essentialist beliefs is weak, since the studies cited above only compared generics against "specific" language such as 'This F is G' or 'This is G'; when the effects of hearing generics were compared to the effects of hearing other language indicating high-proportions (e.g. 'Most Fs are G'), generics no longer appeared distinctive (Hoicka et al., 2021). This suggests that the focus on generics at the expense of other linguistic forms is unwarranted.

While the evidence does not support a direct causal link between hearing generic language and the development of essentialist beliefs, it is clear that there is some association between hearing generic language about a category and forming beliefs about that category, even if those beliefs fall short of being essentialist beliefs and other linguistic forms have similar effects. What we want, then, is an explanation of the empirical data that is consistent with these two facts.

4.6.2 *Inferences about explanations*

In what follows, I will argue that the association between generics and essentialism is based on an inference that a hearer is disposed to make: that generics, when true, are true in virtue of generic explanations. Psychological essentialism amounts, on this interpretation, to the belief that categories have essences and that these essences

are suitable facts about categories which can act as generic explanations for truth of generics. Because essences are not the only possible facts about categories that could make a generic true, the association between essentialism and hearing generics is sensitive to the way that the generic itself is presented. In particular, if a generic is framed in a way that undermines the inference that the explanation is an essentialist one, essentialist beliefs should be less likely to develop. That would explain both the evidence that hearing generics can cause people to develop psychological essentialist beliefs and the evidence which shows that the connection is only indirect. In addition, since the mechanism explaining the association between essentialism and generics is not unique to generics, any sentence which can only have a generic explanation is liable to communicate essentialist beliefs in the same way, so it is predicted that not only generics but other linguistic forms should have this effect too.

4.6.2.1 *The Generic Explanation Inference*

Consider the question: how can a speaker of a temporally unrestricted generic ϕ know that ϕ is true? The truth of ϕ cannot be established on the basis of examining instances of the generalisation, since an instance-based explanation will always run up against the problem that instances in the future are not accessible to observation. In contrast, a speaker could have knowledge about category even if they are not directly acquainted with its members. So, assuming that users of a language operate on the presumption that speakers try to speak the truth (Grice, 1975; Lewis, 1975b), a natural inference for a hearer of ϕ to make is that the speaker has or purports to have knowledge of a generic explanation which explains the truth of ϕ. Call this the *Generic Explanation Inference*. To see how this inference operates, consider the following example.

(37) Women are nurturing.

A hearer, h, hears a speaker, s, utter (37). Operating on the assumption that s has good evidence for the truth of what they say, h assumes that s possesses knowledge of a fact about the category of women in virtue of which (37) is true. It is obvious to h that s does not possess knowledge of enough facts about individual women that

would constitute good evidence for (37). So *h* concludes that (37) is true in virtue of a generic explanation.

This procedure may, but need not, lead to a belief that the generic explanation for (37) involves an essence shared by all women. If we vary the context by introducing a structural frame for a sentence, it is easier to see how a generic could be understood as being made true by an extrinsic generic explanation:

(38) For many years, women have been subject to a patriarchal social order which impresses upon the importance of reproduction and motherhood. As a result, women are nurturing.

In other cases, the assumption that the explanation involves a shared essence is less plausible in the first place; a generic explanation for (39) may involve facts about a labour market, without making any reference to the nature or essence of Premier League footballers.

(39) Premier League footballers are well-paid.

This explanation invokes a consequent fact about Premier League footballers: being well-paid is not what makes a person a Premier League footballer, rather being a Premier League footballer is (part of) what explains why a person is well-paid. This is also an example of an extrinsic or structural explanation, since it does not involve any inherent features of Premier League footballers.

So extrinsic generic explanations are possible. But there is some reason to think that people have a preference for inherent explanations – explanations which only make reference to the nature of objects in themselves.

4.6.3 *The inherence heuristic*

Cimpian and Salomon (2014) argue that people employ an "inherence heuristic" when explaining regularities. According to this model, when trying to explain a regularity, people perform a quick but shallow search for information which might constitute

an explanation. Because inherent features of entities are more readily available than extrinsic features, these can more easily be put to work in formulating an explanation. For example, if I am asked why people drink orange juice at breakfast, the heuristic might work as follows:

(i) I search for an immediately available answer and find none.

(ii) I search for information about immediately available categories, i.e. orange juice and breakfast.

(iii) The information that orange juice is tangy and refreshing and breakfast is in the morning is available.

(iv) I conclude that we drink orange juice at breakfast because its tangy and refreshing taste helps to wake us up.

Cimpian and Salomon argue that historical or extrinsic features of entities are generally much harder for our cognitive system to access (if they are known at all), and so these features are not generally under consideration when the heuristic is in use. It turns out, they claim, that many people now drink orange juice for breakfast because of a successful marketing campaign in the early 1900s by the California Fruit Growers Exchange (Cimpian & Salomon, 2014: 462). But the explanation in terms of inherent characteristics of oranges is much easier to access and so is more likely to be endorsed. Moreover, the inherent explanation is likely to be arrived at before any extrinsic one, increasing the likelihood that it is endorsed (Hussak & Cimpian, 2018). Another relevant factor is the relative simplicity of inherent explanations: extrinsic explanations are likely to be mediated by a huge variety of complex causal factors which are not relevant for inherent explanations. This means that inherent explanations require less cognitive effort to form, feel 'easier' to produce and so are preferred.

Salomon and Cimpian (2014) found that participants who were more likely to employ inherent reasoning exhibited exhibited more evidence of holding essentialist beliefs. That is easily explained on this model, since essences are paradigm examples of inherent features of entities. They are therefore suitable candidates to serve as the

basis for inherent explanations, making explanations in terms of essences cognitively easier than explanations in terms of extrinsic features.

Returning to (39), while a structural explanation is certainly plausible, the employment of the inherence heuristic explains why it may instead be interpreted as holding in virtue of an intrinsic explanation, especially by someone who does not understand how labour markets work. They may reason that the nature of Premier League footballers – their physical ability or their talent – explains why they are well-paid.

Together with the Generic Explanation Inference, this yields a good explanation for the association between hearing generics and the formation of beliefs about the inherent essences of social kinds. The Generic Explanation Inference means that hearers are likely to seek explanations for the generics they hear; the inherence heuristic means that they are likely to settle on explanations that are inherent to the categories that the generic describes.

This explanation meets the first of the two desirable criteria above: it is consistent with the evidence that the association between hearing generics and developing essentialist beliefs is mediated by the context in which an utterance is framed, since the inherence heuristic is only a heuristic: it can be overridden. As the results of Vasilyeva et al. (2018), Noyes and Keil (2019, 2020), Vasilyeva and Lombrozo (2020) and Vasilyeva et al. (2020) suggest, that can be achieved by providing a structural frame for a generic which makes extrinsic explanations more salient or plausible. This may be because the frame explicitly challenges the inherence heuristic, since the frame tells the hearer to look for extrinsic explanations which could serve as a generic explanation for a generic (for some empirical evidence supporting this claim, see Salomon & Cimpian, 2014).

4.6.4 *Other forms of generalisations*

Finally, I will argue that this explanation meets the second criterion above: it is consistent with the same mechanism driving the formation of essentialism for other kinds of sentence, including explicitly quantified generalisations.

What matters for the employment of the Generic Explanation Inference is that the domain of objects under consideration is understood to be indeterminate. When this is the case, a speaker can only know that what they say is true if they possess knowledge of a generic explanation. If there are other linguistic forms which exhibit this feature, then there is nothing about generics which makes them distinctive in this regard. The proposal makes the same predictions about the employment of the Generic Explanation Inference for any sentence which has as a constituent an expression whose domain extends into the future.

If quantifiers like 'most', 'many' or 'all' are interpreted relative to a temporal index which can extend into the future (as suggested by Linnebo, 2022; see also the suggestion in §3.5.7 of this thesis), we would expect that sentences with these quantifiers to behave like generics with respect to nomicity. This suggestion is in line with the empirical evidence that hearing some explicitly quantified generalisations had the same effect as generics on the formation of essentialist beliefs (Hoicka et al., 2021). What was important about the study by Hoicka et al. is that, unlike the others cited above, they compared the effects of hearing generics with the effects of hearing other linguistic forms which signalled high proportions:

(40) a) Zarpies love to eat flowers. *generic*
b) Most Zarpies love to eat flowers. *most*
c) Many Zarpies love to eat flowers. *many*
d) Some Zarpies love to eat flowers. *some*
e) This Zarpie loves to eat flowers. *specific*

This study found that adults essentialised or generalised more when they heard the *many*, *most* or *generic* sentences, compared to the *specific* and *some* sentences, which

lends some initial support to this hypothesis. However, more empirical work would be needed to establish the conclusion. For example, one prediction that could be tested empirically is that sentences with temporal restrictions wholly located in the past should be less likely to trigger the Generic Explanation Inference than sentences with temporal restrictions located in the future. I leave this to future work.

4.7 CONCLUSION

I have argued that the nomicity of generics arises from the nature of the states of affairs that make them true. According to my proposal, generic bare plurals refer plurally. Their reference is restricted by a temporal index which can extend into the future. When it is not determinate which objects are referred to by the bare plural, it is not possible for the sentence to be made true by facts about the objects that satisfy the bare plural description (i.e. by an instance-based explanation). Instead, it can only be made true by a generic explanation, one that contains facts about a category that are independent of facts about which objects are the members of that category. Bare plural generics are nomic only in the sense that they can be made true by facts which are independent of the actual objects that satisfy their bare plural description.

I then showed how this analysis can be brought to bear on the question of the association between generics and psychological essentialism. Essences, if they exist, constitute good candidates for generic explanations, so if generics are interpreted as being made true by generic explanations, it is not surprising that hearing generics about social kinds can lead people to believe that those kinds have essences. On the other hand, the connection is only indirect, since structural or extrinsic explanations can also serve as generic explanations. Any sentence whose domain extends into the future is expected to trigger this same pattern of inference and so the narrow focus on generics prevalent in the literature on this topic is not warranted.

5

RECONCEPTUALISING THE PSYCHOLOGICAL THEORY OF GENERICS

5.1 INTRODUCTION

One puzzle about the semantics of generics is that our intuitive judgments of their truth values do not seem to track facts about the prevalence of a property among members of a category, which has made providing an account of the semantics of generics in set-theoretic terms extremely challenging. The Psychological Theory of Generics (PTG) was developed by Sarah-Jane Leslie (2007, 2008) to account for this fact. According to PTG, generics express a cognitively default mode of generalisation, which is why our intuitions regarding the truth values of generics are so varied.[1]

Picking up on a line of argument raised by van Rooij and Schulz (2020), I will argue that PTG is best interpreted not as a theory of the meaning of generics, but as a theory of the psychological heuristics that we use to judge whether or not generics are true.

The paper begins with a description of PTG and a discussion of the risks associated with interpreting PTG as a theory of the meaning of generics. I then argue that Leslie's methodology is not well-suited to producing a theory of the meaning of generics, on the basis that it takes speakers' judgments at face value and ignores the non-semantic factors that might affect these judgments. Finally, I present a reconceptualised version of PTG as a theory of how heuristics affect our judgements of the truth values of generics and discuss the application of this reconceptualised theory to some of the puzzles posed by generics, including their apparent content-sensitivity, their inferential asymmetry and their association with stereotyping and prejudice.

1 A version of this paper was published in *Philosophical Studies* as (Ralston, 2024b).

5.1.1 *The Psychological Theory of Generics*

I focus here, as Leslie does, on generic readings of English bare plural sentences ('generics'). Generics have historically proven difficult to analyse using the tools of formal semantics. Because of this, Leslie takes a heterodox approach to the study of the meaning of generics. She is, she explains, not interested in the traditional semantic question of how our words relate to the world, but in the psycholinguistic question of how "human beings actually understand their natural languages" (2008: 20). She therefore forgoes the standard tools of formal semantics in favour of her psychological approach.

Eschewing a metaphysical account of genericity, Leslie builds PTG around the idea that generics are the expression of the cognitive system's "default" mode of generalisation. The central idea behind PTG is that humans are hard-wired to generalise in a certain way and that this mode of generalisation is not particularly systematic. Generics give voice to this mode of generalisation, which explains the perplexing features of their semantics. This mode of generalisation is not linguistic but predates the acquisition of language in cognitive development (Leslie, 2007: 381).

Leslie highlights the "near-paradoxical nature of the acquisition of generics" (2008: 16): children develop the ability to use and understand generics from a very early age, before they learn to use explicit quantifiers (Leslie, 2008: 19). Leslie's hypothesis neatly explains why this should be the case, since children have an innate ability to generalise in this 'default' manner, the linguistic expression of which is a generic.

This mode of generalisation is closely associated with our "System 1" reasoning, which is fast and insensitive to factors such as probability and quantity, drawing conclusions on the basis of heuristics and biases (Leslie, 2007: 396; cf. Kahneman, 2003: 688-689). "System 2", by contrast, is slower but more precise and is, according to Leslie, more closely associated with quantified generalisations whose semantics can be stated in set-theoretic terms.

PTG sorts generics into various categories which have different, though related, truth-conditions. The different categories derive from the content of the generic itself.

PTG therefore attributes highly *content-sensitive* truth-conditions to generics. One class of troublesome examples which are an important motivation for Leslie's theory and have formed the basis of much of her subsequent work on generics (e.g. Leslie, 2017) are referred to as 'striking property' generics. These are generics which attribute a striking or dangerous property to a kind:

(1) a) Mosquitoes carry the West Nile Virus.
b) Sharks attack bathers (Leslie, 2008: 14).

Leslie maintains that these generics are true, even though the vast majority of mosquitoes do not carry the West Nile Virus and the vast majority of sharks never attack bathers. This view receives some empirical support from a truth value judgment task and an implied prevalence task in Cimpian et al. (2010: Experiment 1), although judgments were a little more mixed for striking property generics in the study carried out by Prasada et al. (2013).

Leslie offers the following description of what she calls the "worldly truth specifications" of generics:

> We can describe the circumstances under which a generic of the form 'Ks are F' is true as follows:
>
> The counterinstances are negative, and:
>
> If F lies along a characteristic dimension for the Ks, then some Ks are F, unless K is an artifact or social kind, in which case F is the function or purpose of the kind K;
>
> If F is striking, then some Ks are F and the others are disposed to be F;
>
> Otherwise, almost all Ks are F (Leslie, 2008: 43).

What it takes for a counterinstance to be 'negative' is psychological; a counterinstance having some psychologically salient alternative property would count as positive, whereas merely lacking the relevant property would count as a negative counterinstance (Leslie, 2008: 34).

It is worth noting immediately that Leslie does not present these conditions as part of the compositional semantics of generics. They provide a metaphysical account of "how the world must be for the sentence to be true" (Leslie, 2008: 43). She regards it as unlikely that we will be able to give an informative semantics for generics which does not use the generic operator *gen* – the object-language expression of which is not phonologically realised – in the metalanguage, just as we state the truth-conditions for a sentence like 'Bob is red' by using the predicate 'is red' in the metalanguage: 'Bob is red' is true iff Bob is red (Leslie, 2007: 387-388). So she states the (semantic) truth-conditions of a generic such as (1b) in the following way:

(2) 'Sharks attack bathers' is true iff $gen(x)$: shark(x)[attack.bathers(x)]

The lengthy 'worldly truth specifications' above are an attempt to spell out what the world has to be like in order for the right-hand side of this biconditional to be true, a task which, Leslie argues, takes us beyond semantics. Regardless of whether it makes sense to talk, as Leslie does, about a distinction between "worldly truth specifications" and "semantically derived truth conditions",[2] it is surprising that the *worldly* truth specifications depend not only on mind-independent features of the world, but on psychological phenomena, such as a property's being "striking" or a counterinstance's being "negative". This focus on the psychological, rather than semantic or metaphysical, suggests that the objective of PTG is not to offer a theory of the meaning of generics, since it focuses on the conditions under which we make certain judgments about generics, not the conditions under which they are true.

I won't discuss the accuracy of Leslie's account of the truth-conditions of generics in this paper (see Sterken, 2015c for an extended discussion). Instead, I will argue that Leslie makes an error in taking the judgments of speakers about the truth values of generics at face value, and so presents PTG as a theory of the meaning of generics, when in fact the evidence supports a theory which explains how people make judgments about the truth of generics.

2 Throughout, I will use 'truth-conditions' to mean 'the conditions under which a sentence is true'. I leave it to the reader to decide whether the semantic truth-conditions of a sentence can differ from its metaphysical truth-conditions.

5.1.2 *Theory confusion*

As will become clear, the questions that PTG seeks to answer are quite different from those in the formal semantics literature on generics. This is not always made clear; for example, in Sterken's overview of contemporary theories of generics (2017), "Psychologically Based Theories", in which only Leslie's PTG is included, are presented alongside "Normality Theories" and contextualist approaches, as though they are competing theories attempting to answer the same questions. Nickel does the same in his handbook article on generics (2017).[3] But if those theories are attempting to provide a compositional semantics on the basis of our intuitive judgments about the truth values of generics, whereas PTG gives an account of when and why people are disposed to judge generics to be true, they are not in competition.[4] In fact, PTG may well be compatible with a range of semantic theories for generics, provided that the semantic theories are compatible with the psychological explanations PTG offers for our truth value judgments about generics. Some authors, especially those who focus on the variability of the meaning of generics in the way that Leslie does, explicitly acknowledge that PTG is not a competitor to their own semantic theory of generics (e.g. Sterken, 2015a: 27; Lee & Nguyen, 2022: 1757). Leslie's suggestion that generics are associated with an unsystematic default mode of generalisation might explain why the truth-conditions of generics are so variable, without itself offering a theory of their truth-conditions.[5] Semantic theories of generics need not, however, vindicate the truth value judgments that PTG is based on, as it is quite possible that people are systematically mistaken in their evaluations of generics (Sterken, 2015c: 2508).

There is a danger that, outside of the formal semantics literature, these fine distinctions are not properly understood, and PTG is accepted as a semantic theory on the basis that it handles the linguistic data better than any alternatives. There is a

3 Nickel does, however, acknowledge that "cognitive approaches to generics generally, and Leslie's work in particular, raise a host of foundational questions about the relationship between semantics and human cognition" (2017: 456).

4 In her entry on generics in the Stanford Encyclopedia of Philosophy, Leslie draws the kind of distinction that I propose here and separates her psychological account from "semantic analyses of [the operator] *Gen*" (Leslie & Lerner, 2022). She would, I take it, be sympathetic with the views I put forward here, although the presentation of PTG in her early papers muddies the waters somewhat.

5 I am grateful to an anonymous reviewer from *Philosophical Studies* for raising this point.

burgeoning literature which attempts to reach conclusions about subjects as diverse as Kripkenstein's scepticism about rule-following (Cheng, 2011), epistemic injustice (Rosola & Cella, 2020), ideology (Haslanger, 2012; Langton et al., 2012), decision theory (Schiller, 2023) and scientific communication (Claveau & Girard, 2019; DeJesus et al., 2019; Peters, 2021) on the basis of considerations about the meaning of generics. Those who make normative prescriptions about our use of language based explicitly on the semantics given by PTG, such as Langton et al. (2012) and Claveau and Girard (2019), should be clear that the foundations of their arguments are built on a theory of how we evaluate generics, not what they mean. This need not invalidate their arguments; it is possible that all these arguments require is that we are disposed to evaluate generics as true under certain conditions, not that they actually are true under those conditions.[6]

But failing to heed the distinction between how we think about generics and what generics mean risks placing too much weight on the potential effects of language change and in particular the effects of replacing generics with explicitly quantified generalisations (e.g. Langton et al., 2012: 765). If humans are hard-wired to think in the way that Leslie suggests, then the prospects for meaningful change to be accomplished by adding quantifiers to our generalisations seem slim (cf. Saul, 2023: 1695).

5.2 SEMANTICS BY OPINION POLL

I now turn to my argument that PTG is not a theory of the meaning of generics. While Leslie is not concerned with giving an account of the compositional semantics of generics, I take it that she does want to give a philosophically enlightening account of their meaning, since she begins one of her articles by puzzling over the truth-conditions of generics (Leslie, 2007: 375). But if this is what Leslie is aiming for, then

6 For example, Cheng (2011: 12) explicitly separates Leslie's claims about what it takes for a generic to be true/false from her claims about the cognitive factors which underlie our dispositions to accept or reject generics. Cheng goes on to argue that the meaning of meaning sentences, such as 'I mean *plus* by '+'', should be associated with our dispositions to accept or reject them, and it is in the analysis of these dispositions, rather than the truth-conditions of meaning sentences, that PTG is invoked (Cheng, 2011: 15).

asking only under what conditions ordinary speakers judge certain sentences to be true is not the right way to go about it.

Leslie's favoured approach is founded on the understanding that "empirical work on what regulates people's acceptance of generics can provide evidence in favor of one account of truth-conditions over another" (Lerner & Leslie, 2016: 406). While this generalisation is no doubt true, ordinary speakers' views about the truth values of sentences or utterances are mediated by a wide variety of factors, beyond just their knowledge of truth-conditions. So, while empirical data of this kind is certainly useful in testing hypotheses, it cannot be not the final word in an investigation into semantics.

The claims that Leslie makes in favour of PTG suggest that she takes speaker intuitions at face value. For example, she appeals to the judgments of a hypothetical speaker of English, whose response is "invariably immediate and confident" (Leslie, 2008: 1). Later, she says, "[i]f 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: 16). Leslie goes on to suggest that the mediation of psychological considerations forms part of the truth-conditions of generics, or as she puts it, "how the world must be so as to conform to the claim" made by a generic (Leslie, 2008: 43). But this is an incredibly strong claim; although it is in principle possible that such psychological considerations could enter the metaphysical picture, the more plausible hypothesis is that they mediate our truth value judgments without determining how the world must be in order for the sentences to be true.

5.2.1 *Collecting truth value intuitions*

When considering the empirical approach favoured by Leslie, we should ask how it is possible to pair speakers' intuitions concerning the truth value of a particular generic with the conditions under which a speaker is making that judgment. In an ideal world, an experimental setting would fix the conditions and then elicit a truth value judgment. The truth value judgment would then be paired with the conditions.

A representative example of this methodology is the study in Schwarz (2013a).[7] In this study, which aimed to test a hypothesis concerning the semantics of plural definite descriptions, participants were asked to judge whether sentences with plural definite descriptions were true or false, based on an image. The image consisted of shapes of different colours displayed in an array and participants were asked to judge sentences such as 'The circles were black'. By pairing judgments about sentences with particular states of affairs, e.g. half of the circles being black, it is possible to give a mean truth value judgment relative to a particular state of affairs.[8]

How would one go about performing such a task for generics? One problem with this approach is that generics are usually unbounded, in the sense that they resist contextual domain restriction (Krifka et al., 1995: 45; Drewery, 2005: 376-377; Sterken, 2015a: 13). This means that when we elicit truth value judgments for generics, we are asking speakers to make judgments based on their background conceptual knowledge, not on something controlled for by the experimenter.

Two individuals may therefore disagree over the truth value of a particular generic, not because they differ in their beliefs about its truth-conditions, but because they differ in their beliefs about whether those conditions obtain. For example, two people may disagree concerning the truth value of the following generic:

(3) Ships have sails.

One person may believe that a majority of ships have sails, while the other may believe that a majority of ships do not have sails. The first may therefore judge the sentence true, while the latter false, on the basis of this background information. Alternatively, one might hold that it is part of what it is to be a ship to have sails, while the other may believe that being a large boat is sufficient to be a ship, in which case again the first would judge the sentence true and the latter may judge it false. The two

7 For related examples of this experimental method in semantics, where the experimenter manipulates a carefully controlled context and asks their participants to make a judgment of some kind, see Križ and Chemla (2015) and Augurzky et al. (2023) on plural definite descriptions, Pietroski et al. (2009) on the interpretation of 'most', and Syrett et al. (2006, 2010), Aparicio et al. (2015) and Weicker and Schulz (2020) on gradable adjectives.

8 Schwarz's analysis is more nuanced than this; the experiment aimed to measure reaction times for truth value judgments relative to different states of affairs, in order to test a hypothesis concerning whether 'non-maximal' (i.e. exception-tolerating) readings of plural definite descriptions should be accounted for in their semantics or as a pragmatic phenomenon.

people might therefore disagree on the truth value of (3) because they disagree about the nature of ships, even though, as in the first case, they agree on the truth-conditions of (3). So unlike in Schwarz's experiment described above, where the conditions against which sentences are judged true or false are strictly controlled, it is difficult to control for different background assumptions which may produce different truth value judgments even when participants agree on truth-conditions.

5.2.1.1 *Case study: Prasada et al. (2013)*

To illustrate this difficulty, consider an empirical study by Prasada et al. (2013) which elicited truth value intuitions about a wide range of generics about familiar kinds from ordinary speakers. Participants were asked to judge the truth of sentences on a scale from Definitely True (+3) to Definitely False (–3). This study found evidence of very mixed truth value judgments for many generics. For example:

(4) a) Birds have wings. *Mean truth value judgment*: 2.63
b) Lemons are sour. *Mean truth value judgment*: 2.25

That these apparently obviously true generics received mixed judgments suggests that participants evaluate them relative to different bodies of evidence. This is not surprising; conceptual representations are likely to be highly idiosyncratic and there is no reason to believe that one person's representation of the concept BIRD or LEMON will be identical (or even could be identical) to another's.

Another finding of this study was that participants were unwilling to call any generics completely false, even those often considered paradigm cases of false generics, such as 'Books are paperbacks' (0.21) or 'Canadians are right-handed' (–0.50) (Prasada et al., 2013: Appendix B). Presumably, some of the generics that they tested are indeed completely false, for instance:

(5) a) Plants are ferns. *Mean truth value judgment*: 0.04
b) Novels are mystery novels. *Mean truth value judgment*: 0.04

What are we to make of the result that these sentences were judged more true than false? First and foremost, I contend, it would be a mistake to take this data at face

value and try to construct a semantic theory to vindicate it. We cannot tell from the data what conditions participants were evaluating these sentences relative to,[9] so it is difficult if not impossible to draw conclusions about the truth-conditions of generics based on this data. It may be that participants attempted to accommodate background information in order to make the sentences more plausible, which would explain why so few generics were judged definitely false (the mean judgment for 'minority false' generalisations like (5a) and (5b) was -0.52). This suggests that we must be very cautious when interpreting the results of this kind of study. It is not possible to fix the background conditions against which speakers make their truth value judgments about generics.

5.2.1.2 *Case study: Cimpian et al. (2010)*

One method of countering this shortcoming is to use unfamiliar novel categories to elicit truth value judgments. For example, Cimpian et al. (2010) introduced participants to sentences about animals from a remote island, called "lorches", and asked them to make judgments about sentences such as:

(6) a) Most lorches have purple feathers.

b) Lorches have purple feathers.

(7) a) Most lorches have dangerous purple feathers.

b) Lorches have dangerous purple feathers.

The results for generics were then compared with the sentences featuring the explicit quantifier 'most'. Cimpian et al. showed that, unlike for sentences with the quantifier 'most', participants were more likely to accept a generic when the property was dangerous than when it was not, for a given prevalence of the property among the category.

Admirable though this approach is, it is still not equivalent to the more controlled approach taken by Schwarz in the study concerning plural definite descriptions described above. We can assume that participants were using some information

9 With two exceptions: participants were asked to make prevalence and cue validity estimates.

other than prevalence to make judgments about the truth of the sentences they were presented with. But beyond this, it is difficult to say precisely what information the participants were using to assess the truth values of the generics, since we don't know what background assumptions the participants made about lorches on the basis of the limited information they were given. Indeed there is a growing body of empirical evidence demonstrating that people form beliefs about the nature of novel categories described by generics on the basis of very little information (e.g. Gelman et al., 2010; Rhodes et al., 2012; Noyes & Keil, 2019; Vasilyeva et al., 2020; Leshin et al., 2021; Benitez et al., 2022). So once again, even in a experimental setting, participants are making judgments about generics relative to opaque bodies of background knowledge, which cannot be controlled for. What we learn from this study is that whether a property is dangerous had an effect on truth value judgments, but we cannot say relative to what conditions participants were judging the sentences, so we cannot reach conclusions about people's beliefs about the truth-conditions of generics.

5.2.2 *Top down semantics*

This leads to another methodological question: when we present generics to ordinary speakers, why should we expect that they can reliably and accurately ascertain their truth values? Ordinary speakers are unlikely to attend to fine distinctions between semantic features of a sentence and non-semantic features which do not make a contribution to the truth-conditions of the sentence, but which affect judgments of truth.

To take unreflective truth value judgments as the starting point from which to develop a semantic theory, as Leslie does, is to take a 'top down' approach. This contrasts with a 'bottom up' approach on which we would proceed from the meanings of a sentence's constituents and composition rules to the truth-conditions of the sentence as a whole.[10] The trouble with the top down approach is that it is likely

10 Again, I should emphasise that Leslie explicitly denies that she is giving a compositional semantics for generics. As should be clear, I agree with this contention, but I differ from Leslie in my view of what her theory is really about.

to require a certain gerrymandering of the semantics to fit the data. This is because variations in intuitions may correspond to pragmatic or non-linguistic features of a sentence or utterance, as opposed to semantic features. These features can be subject to systematic error or inconsistency. As Kent Bach has remarked:

> Taking semantic intuitions seriously would make life miserable for semanticists, even more miserable than it already is. It would require doing semantics from the top down... Covert constituents would have to be posited to provide the "residue" of meaning not accounted for by the overt ones. Or, alternatively, special meanings, hence ambiguity, would have to be attributed to certain of the overt constituents, insofar as their ordinary meanings seem not to make the right contribution to what is said. All this can be avoided if we don't take people's seemingly semantic intuitions too seriously (Bach, 2002: 26).

Leslie introduces both a covert constituent and ambiguity into her theory of the meaning of generics: first, she uses the covert operator, 'GEN', which plays an essential variable-binding function; second, she describes truth-conditions which vary radically based on the content of the generic.

These facts do not on their own discredit PTG as a semantic theory, but we should approach PTG with scepticism, because truth value intuitions should not form the basis of semantic theory when they are taken uncritically and at face value, rather than as data against which to test hypotheses, since they may be tracking non-semantic features of sentences. The best methodology for semantics represents a blend between the top down and bottom up approaches, where hypotheses about the contributions of individual words or morphemes and composition rules are assessed against data from truth value intuitions. Discrepancies can be resolved either by amending the hypotheses, or by proposing an error theory to explain away the discrepancy. But in PTG, the bottom up approach is completely absent, especially given the lack of an informative analysis of the object language expression, 'GEN'. The purely top down approach leads to trouble in this particular case because the truth value judgments for generics are so content-sensitive and unsystematic, which provides some at least

prima facie evidence that there is something beyond their semantics mediating the judgments. Moreover, as I argued in the previous section, generics are distinctive in that is difficult to know against what conditions people are making truth value judgments, which makes those judgments especially bad guides to the truth-conditions of those sentences.

The complex and variable nature of Leslie's truth-conditions for generics suggests that they are a result of "overfitting" the data of our intuitions (Williamson, 2020: 264): by failing to recognise that our evaluations of sentences do not perfectly track the truth values of those sentences, Leslie's theory is made unnecessarily elaborate in its content-sensitivity. To put the complexity into the truth-conditions is to be overly credulous in our treatment of speakers' intuitions. Why think that ordinary speakers can simply intuit the truth values of sentences which say something about entire categories, as generics do? Another, more plausible explanation for the variability in truth value judgments is available: speakers use heuristics to evaluate the truth of statements that they are otherwise unable to verify.

5.2.3 *Heuristics and truth value judgments*

We can see the importance of not taking truth value intuitions too seriously elsewhere. A perplexing variability in truth value judgments would be observed for active/passive minimal pairs, since, as Kahneman notes, the grammatical subject makes different information accessible to the cognitive system (2003: 701):

(8) a) Arsenal beat Tottenham.
    b) Tottenham were beaten by Arsenal.

In these cases, judgments of acceptability do not track only semantic features of the sentences. But why would we assume that such divergence amounts to a difference in how the world must be in order for them to be true? To use them as evidence of truth-conditional divergence would be misguided (cf. Frege, 1960: §3).

Another helpful demonstration of this is the "inverse conjunction fallacy", observed by Jönsson and Hampton (2006). This is a phenomenon where people judge universal generalisations about subsets of a category to be less likely to be true than unmodified universal generalisations about the category as a whole. For instance, consider the pair:

(9) a) All sofas have backrests.

b) All uncomfortable handmade sofas have backrests.

In Jönsson and Hampton's studies, they found "that people have a strong tendency to accept unmodified universally quantified sentences [e.g. (9a)] with greater frequency and/or greater confidence than the equivalent modified versions [e.g. (9b)]" (2006: 328).[11] Given a standard account of the semantics of 'all' as a universal quantifier, (9a) entails (9b), since all uncomfortable handmade sofas are sofas. But we cannot recover this fact from experimental data of people's truth value judgments, since some hold (9a) to be true while (9b) is false. If we started from that data, we would be led to posit a variable, highly content-sensitive semantics for 'all'. We might even, in the same vein as Leslie, attempt to provide a complex account of what the world has to be like in order either (9a) or (9b) to be true which vindicates these intuitive judgments.

This would be a mistake. An alternative and much more plausible route is to view the semantics of 'all' as simple and propose a hypothesis about human reasoning which can explain the discrepancy from what we would expect to see given the semantics.[12] For Jönsson and Hampton, this is a heuristic, which is fallible despite being generally useful in evaluating universal generalisations. When evaluating whether a universal generalisation is true or false, if we do not possess "an easily searchable exemplar space, we rely instead on the strength with which the predicate is represented as part of our concept" (2006: 331). In the case of the sofa, there is no easily searchable

11 The same effect was found when the sentences were modified by 'every single' and '100% of' (Experiment 4), which suggests that the effect is not linked to the word 'all'.

12 This follows from the following methodological principle, suggested by Sterken in her own discussion of PTG: "If a cognitive bias affects our intuitions in systematic ways, first try to treat these intuitions as mistaken, and only move to a semantic, metasemantic or pragmatic strategy if that fails" (2015b: 91). The general cognitive bias at work is presumably related to what Leslie calls the default mode of generalisation, but it seems likely on the basis of the evidence discussed here that its application is wider than just the evaluation of generics.

exemplar space of sofas (since we haven't observed all the sofas in the world), so we default to representativeness to decide if the generalisation is true. This results in inconsistent judgments for the two sentences, since the representation or prototype of an uncomfortable handmade sofa is not as closely related with having a backrest as is the case for the prototypical sofa. We can in this way account for the inconsistency of the data without positing an inconsistent semantics for any of the constituents of the sentences. The result is essentially an error theory, which is both simpler than a PTG-style semantic theory and makes falsifiable predictions about the conditions under which people tend to make errors in their evaluation of the truth of universal generalisations.

Reflecting on this result, Hampton elaborates:

> As philosophy and logic students are quickly made aware, universally quantified statements can only strictly be verified by an exhaustive search for counterexamples and the failure to find them. "*All swans are white*" is notoriously much harder to prove true than to prove false. Human language has opted therefore to take a different approach to determining "truth" of such statements, one which frequently leads to judgments that appear fallacious from a logical perspective. Statements of this kind are judged true on the basis of the strength and importance of their association with the content of a concept. To say that "*All S have property P*" in the context of everyday speech is to say something like "*An important and relevant fact about the kind S is that the property P is to be found in individuals of that kind*" (Hampton, 2012: 20).

Leslie herself seems willing to countenance such non-truth-conditional explanations in her own work on what she calls the Generic Overgeneralisation Effect, which is the tendency for people to recall generics as universal generalisations. Regarding Jönsson and Hampton's experiments, Leslie concludes: "Either this constitutes a clear failure to appreciate the logical features of universal statements, or else their participants were interpreting these universals as generics" (Leslie et al., 2011: 28). There is, quite rightly, no suggestion that the semantics of universal generalisations need to be revisited to

account for the truth value intuitions. So why not extend this view to the experimental data concerning generics? It is equally possible to take the variability in truth value judgments of generics as evidence that people do not 'appreciate the logical features' of generics or that they default to another interpretation of the sentence in the absence of knowledge about all members of a kind.

I have shown that taking truth value judgments at face value would lead us to propose complex and variable truth-conditions for universal generalisations. Since this is a route that I suspect few working in semantics or cognitive science would like to take, I take it as good evidence that ordinary speakers' truth value judgments can be mediated by heuristics which can introduce inconsistency into our judgments.

5.2.4 *Heuristics or meaning?*

A defender of PTG's status as a theory of the meaning of generics might object to the argument that our intuitions about the truth values of generics are mediated by heuristics by responding that they have an equally good explanation of why our intuitions about the truth values of generics are unsystematic. This is the Core Claim of PTG: the generalisations that generics express correspond to the judgments of the default mode of generalisation (Leslie, 2007: 381, 394), so the truth-conditions of generics correspond to the conditions under which our default mode of generalisation tells us to accept a generalisation (leaving space for empirical investigation into precisely what those conditions are). Since this mode of generalisation is unsystematic, our intuitions about the truth values of generics are too (Leslie, 2007: 397-398). This claim is independent of the specific clauses Leslie gives for the truth-conditions of different types of generics that I quoted above, so let's focus just on the Core Claim.[13]

As a view about what the world has to be like in order for a generic to be true, the Core Claim of PTG is subject to two damaging objections. The first concerns what happens when two (or more) people's default modes of generalisation lead to different truth value intuitions for a particular generic (in a given context). Imagine

13 I am grateful to an anonymous reviewer from *Philosophical Studies* for raising this objection and suggesting this manner of presentation.

that we are able to hold fixed the conditions against which those people evaluate the sentence. Which of these intuitions is 'correct'? If there is no way of deciding between these intuitions, then there seems to be no unique answer to the question of what the truth-conditions of that generic are. That means that if we pursue the cognitive explanation of the meaning of generics proposed by PTG, we have to choose between two unpalatable options: either each of our default modes of generalisation functions in the same way (i.e. with the same inputs, they produce the same outputs), or the truth of generics is relativised to each distinct default mode of generalisation (i.e. each person). Even if the truth-conditions of generics are indeterminate or underdetermined in some way, they are presumably not different for each observer.

Second, even assuming unanimity in truth value judgments, with the clauses regarding the truth-conditions of generics removed, PTG would predict that any intuition about the truth value of a generic informed by the default mode of generalisation is 'correct'. Reduced to the Core Claim, then, PTG makes no falsifiable predictions, since if our default mode of generalisation leads us to endorse a generic, it is true. That result would be especially damaging for Leslie's project, as she wants to be able to maintain that striking property generics, like 'Mosquitoes carry the West Nile virus', are true, while superficially similar striking property generics about social groups, like 'Muslims are terrorists', are false (Leslie, 2017: 403). This illustrates the problem with this mode of overfitting the data: if PTG makes no falsifiable claims about the meaning of generics, it is not a good theory of the meaning of generics.

I suspect Leslie is aware of that flaw, which explains why she rightly appeals to evidence from the psychological literature so that she is able to make more systematic claims about our judgments of the truth of generics (Leslie, 2008: 16). Without these claims about the patterns in our judgments, PTG is vacuous. With the addition of the clauses stating the truth-conditions of different types of generics, the theory is at least falsifiable, but this does not protect PTG from the first criticism. By tying the truth-conditions of generics so closely to our judgments about the truth of generics, PTG leaves no room for disagreement and hence errors in our evaluations of generics.

I therefore conclude that the Core Claim of PTG does not constitute a viable basis for a theory of the meaning of generics.

5.3 EVALUATION HEURISTICS

In what follows, I will argue that the hypothesis that fallible heuristics mediate truth value judgments helps us to make sense of the variability in our judgments about the truth values of generics, and that PTG is fruitfully interpreted as a theory of these heuristics of judgment, not of the meaning of generics.

Heuristics of judgment are methods of making decisions under conditions of uncertainty or limited resources: "A heuristic is a strategy that ignores part of the information, with the goal of making decisions more quickly, frugally, and/or accurately than more complex methods" (Gigerenzer & Gaissmaier, 2011: 454). In practice, most decisions are made with limited cognitive resources or incomplete knowledge, so we frequently employ heuristics to help us make judgments.

5.3.1 *Why heuristics?*

First, one might ask why we would need to use heuristics to judge the truth of generics. The most convincing answer is that it is hard for us to know for certain whether a generic is true (Prasada, 2000: 67). We frequently need to make decisions in cases where there is either insufficient information or where there is no optimal solution to a problem and it is in these cases that the use of heuristics is appropriate (and perhaps even required). Judging the truth value of a generic may be just like this, since it is clear that even on reflection we (both ordinary speakers and semanticists) do not know what its truth-conditions are, let alone whether they obtain.

The truth-conditions could, for example, include some modal element (Asher & Morreau, 1995; Nickel, 2016; Thakral, 2018; Kirkpatrick, 2023), require the existence of some sort of pattern which must be recognised over time (Carlson, 2008; Cohen,

1999b),[14] rest on a notion of kind-predication where unobservable kinds inherit some of the properties of their members (Liebesman, 2011; Teichman, 2023), or they may even be context-sensitive (Sterken, 2015a; Lee & Nguyen, 2022), semantically incomplete (Nguyen, 2020), or unspecific with respect to the generalisations they express (Bosse, 2021). If these suggestions are along the right lines, it will be extremely difficult for ordinary speakers to know with any certainty that a generic is true, even given a body of evidence. But even if generics are just unbounded universals concerning all or almost all actual instances of a property, in practice we would not come to endorse or reject them by considering each instance in turn. Just as Hampton suggested in the passages quoted above, we do not in practice evaluate universal generalisations on the basis of the observation of every instance, but instead evaluate them on the basis of "the importance of their association with the content of a concept" (2012: 20).

To see some evidence for this claim, consider the 'modifier effect', which has been observed in several studies (e.g. Connolly et al., 2007; Hampton et al., 2011; Jönsson & Hampton, 2012). Participants were asked to score the following sentences for likelihood to be true (Connolly et al., 2007):

(10) a) Candles are made of wax. *unmodified*

b) Scented candles are made of wax. *prototypical modifier*

c) Purple candles are made of wax. *non-prototypical modifier*

d) Expensive purple candles are made of wax. *additional non-prototypical modifier*

The mean ranking across a set of target sentences decreased from (10a) to (10d), i.e. participants judged each sentence to be less likely to be true than the last. The introduction of a non-prototypical modifier decreased the truth value rating relative to sentences with a prototypical modifier, which suggests that easy availability of a prototypical representation of a concept modulates truth value judgments.

14 Van Rooij & Schulz (2020) explicitly draw a connection between Cohen's probability theory of generics and Kahneman's Heuristics and Biases research programme. The basis of this connection is the fact that people often mistake typicality for probability, i.e. if my representation of a typical F has the property G, I will consider it probable that an arbitrary F is G.

In an unpublished study reported in Hampton (2012), Hampton and Lam investigated the interactions between the modifier effect and the inverse conjunction fallacy. They found that the addition of 'all' into a bare plural sentence led, as one would expect, to a lower mean average in the judgments of the likelihood of the sentence being true. The modifier effect was significant for both the 'all'-sentences and the bare plurals, although the effect was larger for bare plurals. Plausibly, then, the same heuristics are being used to evaluate both universal generalisations and generics, with the difference that for universal generalisations, people are less tolerant (although not completely intolerant) of counterinstances.

This evidence supports the conclusion that we use heuristics to evaluate both generics and universal generalisations. These heuristics may lead us into inconsistency and do not necessarily reflect the truth-conditions of either generics or universal generalisations.

Acknowledging the impossibility of its justification on the basis of deductive inference, Frank Ramsey described the practice of induction as a "habit of the mind" without which we would be "very much worse off" (1990: 92-93). We should think of the heuristics that guide our judgments of the truth of generics in much the same way; without these heuristics, we would be unable to judge whether or not generics are true and, given the important role that they play in our conceptual knowledge (Prasada & Dillingham, 2006) and early education about the world (Cimpian & Markman, 2009), there is a clear utility in being able to make these judgments, even if they are fallible. Perhaps this is the best that we can do.

5.3.2 *What is preserved?*

Once we view PTG as a theory of how heuristics affect our judgements of the truth values of generics rather than of their truth-conditions, we are better equipped to understand some of their puzzling features.

5.3.2.1 *Content-sensitivity*

One of the challenges for giving a semantics for generics is that our truth value judgments appear to be sensitive to the wording and content of generics, which is what motivates Leslie's complex and variable truth specifications. Since heuristics can be sensitive to quite fine-grained distinctions in language (such as word order or choice), their use would explain why our truth value judgments appear very sensitive to the content of generics. Making things more concrete, we can modify Leslie's truth-conditions to suggest heuristics for generics.

Judge a generic, 'Ks are F' true when the following conditions hold:

(11) a) The counterinstances are negative.

b) If F is a property of a prototypical K then some Ks are F.

c) If K is an artifact or social kind then F is the function or purpose of the kind K.

d) If F is striking or dangerous then some Ks are F and the other Ks are disposed to be F.

e) If none of the antecedents of the three preceding conditions are true, then almost all Ks are F.

I do not claim that these heuristics correctly predict our truth value judgments for generics, merely that PTG can be restated as a theory of truth value judgments rather than worldly truth specifications.

The explanation for these heuristic rules will presumably relate to Leslie's default mode of generalisation; people make judgments in this way because of features of human cognition. The difference is now that this mode of generalisation plays no role in the *semantics* of generics.

Condition (11d) calls for special consideration though, as we might ask why striking or dangerous properties are singled out in this way, since prototypical Ks are typically counterinstances to such generics. I suggest that the explanation for this heuristic would be purely pragmatic, in the broadest sense of the word: it is advantageous to

treat Ks as if they were F when F is a property that poses some kind of threat (e.g. carrying the West Nile Virus) (compare with Leslie, 2017: 396). So, on the basis that it is not generally possible to identify which mosquitoes are carriers of a disease without carrying out some kind of impractical testing, I might avoid all mosquitoes on the basis of the belief that some mosquitoes carry the West Nile Virus, and would utter or endorse the generic, 'mosquitoes carry the West Nile Virus'. As Restall puts it, in these cases "it is much better to err on the side of *false positives* than *false negatives*" (2018: 2).

5.3.2.2 *Inferential asymmetry*

The inferential asymmetry of generics, the tendency to infer more from a generic than we require to believe that it is true (Cimpian et al., 2010), plays an important role in Leslie's more recent work on the harmful effects of social generics (e.g. Leslie, 2017). Acknowledging the role of heuristics in our judgments of generics, we can now see that the inferential asymmetry of generics might be caused by the fact that we habitually accept false generics as true and then use these false generics to draw inferences, which are unsound because they have a false premise. If the conditions under which we accept generics as true bear no direct relation to prevalence, but instead to representativeness, as suggested by the model of bounded rationality presented in Tversky and Kahneman (1982), it is no surprise that there is a mismatch between the conditions under which we accept them as true and the extent of the inferences we use them to license.

This raises the question, why should the asymmetry hold in the direction that it does, i.e. why do we accept generics on the basis of little evidence but use them to license strong inferences, rather than the other way around? To give a satisfactory answer to this question would take me too far into the realms of cognitive psychology, but it may be the case people habitually overestimate the extent to which the properties of a prototypical instance of a category are prevalent among other instances of the category (cf. n. 14). This would be, as Cimpian et al. suggest (2010: 1475-1476), an exemplification of the tension between intensional reasoning (reasoning about a concept) and extensional reasoning (reasoning about objects which fall under a

concept), since people may take a property to be typical of a concept even if it is not possessed by a high proportion of objects which fall under the concept. In the long run, it may be better to overgeneralise on the basis of a few instances than to undergeneralise. We are able, then, to make more and more confident predictions about new instances of categories that we encounter; these predictions may later be revised in the face of counterevidence. But this does not, however, force us to posit any inconsistency in the semantics of generics. The inconsistency is in our heads.

5.3.2.3 *Stereotyping and prejudice*

As we have seen, the world contains too much information for us to process it all in detail, so our cognitive system must be selective in what information it chooses to process. Heuristics help us to navigate the complexity of the world, but at the same time they can lead us to make or endorse morally questionable generalisations when they concern social kinds, such as ethnic, religious or gender groups, as Leslie has argued (2014, 2017). This coheres with the view in social psychology that stereotyping is closely related to the use of heuristics of judgment (e.g. Rothbart et al., 1978; Bodenhausen, 1990);[15] if we make judgments about generics using heuristics, it is no surprise that generics are associated with stereotyping since the same mechanism is responsible for our endorsement of both generics and stereotypes. This suggests that there is still mileage in the increasingly popular research programme concerning the links between the use of generics and prejudice.[16]

In their landmark discussion of the semantics of generics, Krifka et al. rejected the theory that generics express stereotypes, which they attribute to Declerck (1986) and Geurts (1985),[17] on the basis that this ties the truth-conditions of generics too closely to subjective issues such as cultural norms (Krifka et al., 1995: 49); for example, the

15 See also Allport's description, in his seminal text, of prejudice as the result of the "normal and natural tendency to form generalizations, concepts, categories whose content represents an oversimplification of [our] world of experience" (1954: 27).

16 On this issue, see Haslanger (2012), Langton et al. (2012), Leslie (2014, 2017), Wodak et al. (2015), Wodak and Leslie (2017), Lemeire (2021), McKeever and Sterken (2021) and Bosse (2024).

17 Recently, Bosse (2024) has reversed the explanatory direction and argued that generics express stereotypes precisely because the content of stereotypical beliefs is generic.

following sentence is false, regardless of whether it is held to be an item of cultural knowledge within a given community:

(12) Snakes are slimy.

Their conclusion is surely correct; generics are about the world and we can be mistaken about their truth values if we are mistaken about the nature of world. But even if stereotypes play no role in their semantics, is it just a coincidence that generics are closely associated with stereotyping? By acknowledging the role of heuristics in the evaluation of generics, we can see how the same psychological processes that are involved in stereotype formation are involved in the judgment of the truth values of generics.

Since on the model I have argued for here, universal generalisations are also judged using heuristics, one might ask why they are not equally associated with stereotyping. The answer is simple: universal generalisations are less tolerant of counterexamples than generics are, and so are more easily falsified (Lemeire, 2021). As such, it is much easier to show someone categorically that their belief in a universal generalisation is false than it is for a generic, so it is easier to hold on to a belief in a generic in the face of counterevidence.

Reconceptualising PTG as a theory of heuristics can help us see how what is usually an effective and efficient way of thinking can lead us astray in certain cases, with potentially harmful effects (see also van Rooij and Schulz, 2020 for a similar suggestion). But, given that the problem is not located in the semantics of generics, this view encourages us to ask whether the focus on *generics* in particular is justified, since it points towards a more general psychological explanation, rather than one associated with a particular form of words.

5.4 CONCLUSION

I have argued that we should reconceptualise Leslie's theory of the meaning of generics as a theory of how heuristics affect our judgements of the truth values of generics.

Viewing the theory in this way sheds light on the link between generics and cognitive biases, and opens up the possibility of drawing connections between our disposition to judge certain generics as true and other cases of bias resulting from the employment heuristics of judgment. As PTG is compatible with a range of theories of the semantics of generics, it offers an opportunity to explain away some of the truth value intuitions which diverge from the predictions of our theories, making the daunting project of giving a semantics for generics just a little easier.

In uncovering and probing these heuristics, Leslie makes a valuable contribution to our understanding of generics. But it is important that philosophers, linguists and cognitive scientists keep in mind that PTG does not provide an account of the meaning of generics so that they do not overestimate the potential to change how humans think by changing how they speak.

6

CONCLUSION

Across the five chapters of this thesis, my objective has been to promote a version of a plural reference theory of generics. I have presented a range of linguistic evidence supporting the hypothesis that generic bare plurals are referential and refer to more than one thing. I gave a compositional semantic theory for a fragment of English and showed that this theory can explain a range of puzzles associated with the meaning of generics, including collective and distributive ambiguities, the felicity of generic co-predications, the availability of both existential and functional readings of bare plurals and the nomicity of characterising sentences. I have also argued that the phonologically null generic operator, 'GEN', is not required to explain linguistic data that have traditionally been claimed to support its existence.

I gave two closely related versions of the temporal plural reference theory, one in Chapter 3 and a slightly more complex version in Chapter 4. I view the truthmaker semantics of Chapter 4 as a way of interpreting the language of regimentation of Chapter 3. The semantics of Chapter 3 state the truth-conditions of generics; the truthmaker semantics of Chapter 4 enable us to make distinctions between the different ways that these conditions can be satisfied. Strictly speaking, the version of the theory given in Chapter 3 can stand independently of the truthmaker semantics of Chapter 4, since the latter makes finer distinctions between states of affairs than the former.

A central theme in this thesis, especially in Chapters 3 and 4, has been the relationship between generics and time. This relationship is often ignored by those who adopt kind-predication or modal theories of generics. That is because, having avoided the obvious embarrassment of giving a completely extensional theory for generics, they fail to attend to the subtle linguistic evidence of generics that are explicitly restricted to some period of time. As I argued in Chapter 4, this is especially important when

considering the nomicity of generics, which is often taken as an important motivation in favour of a modal theory. Acknowledging that generics concern not only the present, but the past and future too, allows us to avoid the embarrassments of an extensional theory without taking our eye off the actual world.

I have also strongly rejected the notion of a kind which frequently appears in discussions of the meaning of generics. The reason for this is that it is clear to me that kinds play no important role in the compositional semantics of generics. Moreover, there is little clarity on what different theorists mean by a kind; I have argued that many of the natural ways of explicating the notion fail to do the work needed by a semantic theory for generics.

One source of confusion here is that the notion of a kind is used by philosophers, linguists and psychologists in subtly different ways. Yet it is invoked by all three groups in discussions of the meaning of generics. It is unhelpful for different groups working on a related topic to use one term in such different ways, which is why I have tried to avoid using the term at all in my own analysis of generics.

The notion of a kind in contemporary analytic philosophy, usually attributed to Mill, is a grouping of objects which stand in some relation of similarity to each other. These kinds may have some important role to play in our philosophical theories (e.g. of the justification of induction), but I have argued that the notion of a kind that would be required to play the role of the objects of reference of generic bare plurals is too liberal to play a useful role in philosophical theories.

When psychologists say that "generics are about kinds" or that "generics designate kinds", I believe that this is consistent with my claim that kinds play no role in the compositional semantics of generics. That is because their claims relate to the way in which we represent information to ourselves; kinds, in this context, are ways of organising information internally. These kinds do play a crucial role in our inductive and explanatory practices, but they can't play the justificatory role that philosophers need them to because they are merely internal representations. In giving a truth-conditional semantics for a fragment of a natural language, we are concerned with relations between words and the world, not only the ways that our minds organise

information. So, once again, the kinds needed to play the role of the objects of reference in a semantic theory cannot be the same as the kinds which psychologists study.

Another theme of this thesis has been an attack on the idea that generics are linguistically 'special'. An inference that is commonly drawn from this idea is that a semantic theory for generics will need to invoke some new semantic resources to explain how they mean what they mean. I have attempted to show that many of the puzzles associated with generics are not as distinctive as they first appear and that recognising parallels with other linguistic expressions is the key to understanding the semantics of generics. By borrowing resources from existing semantic theories for other expressions of English, we can produce a more unified theory and make fewer contentious assumptions (or perhaps: the contentious assumptions will not be made only in the service of this one narrow theory). This attitude is most clearly expressed in the arguments in Chapter 2 that the quantificational distributivity operator commonly invoked to explain the availability of distributive readings of non-generic plurals can play the variable-binding function often believed to necessitate a special generic operator or predicate modifier.

The plurality of bare plural generics has also played a central role in my analysis. One disadvantage of this approach is that it might obscure some of the parallels between bare plurals and singular linguistic forms capable of expressing generic meanings, such as definite or indefinite singulars. Some may feel that this approach to the semantics of generics 'misses a generalisation' by failing to assign a closely related semantics to these other expressions.

My reply to this is that we must do semantics from the bottom up, as I suggested in Chapter 5. What we want is a compositional semantic theory for each expression which explains both the similarities and differences between the expressions. I argued in Chapter 1 that genericity is a psychological phenomenon which arises when we believe that a sentence holds in virtue of a stable state of affairs. Viewing genericity in this way – rather than as tied to a particular expression – means that it is far easier to see how the independent projects of giving semantic theories for the full range of linguistic forms capable of expressing generic meanings should go. That this

class includes bare plurals, definite singulars, indefinite singulars as well as, more controversially, definite plurals and sentences with overt quantifiers such as 'all' or 'most' suggests that the bottom up methodology will be especially valuable.

Although I touched on the empirical questions concerning the links between generics and psychological essentialism in Chapter 4 and Leslie's discussions of the dangers of the default mode of generalisation in Chapter 5, I have not engaged with the wide range of normative questions surrounding generics. To do justice to this project would have required another whole thesis.

However, what I have had to say about the semantics of generics has some bearing on questions concerning the difficulty of falsifying generics. A practical application of this concerns fact-checking by services such as Full Fact or BBC Verify. For these services, it is important to be able to say whether a statement is true, false or perhaps misleading, but this of course presupposes that they know the truth-conditions of the statements that they assess. If generics and other plural sentences have the truth-conditions that I have claimed, then the task of fact-checking these sentences will be very difficult, because whether it is correct to say of some things that they have a property will depend on the purposes of conversation. That leaves ample space for unscrupulous (or careless) speakers to use generics in ways that leaves them plenty of leeway to deny that what they have said is false. It also suggests that the lack of specificity built into the semantics of generics really does make them distinct – at least in this regard – from explicitly quantified sentences, whose truth-conditions are (for the most part) precise.

The subject of generics continues to represent a fertile area for interdisciplinary research. So, although this thesis has focused on a relatively narrow question concerning the compositional semantics of a kind of English sentence, I hope that the arguments made here can help us to better understand many puzzles in philosophy, linguistics and psychology.

BIBLIOGRAPHY

Acton, E. K. (2019). "Pragmatics and the Social Life of the English Definite Article". In: *Language* 95.1, pp. 37–65. DOI: 10.1353/lan.2019.0010.

Agolli, E. (2023). "The Complex Lives of Proper Names". In: *Linguistics and Philosophy* 46.6, pp. 1393–1439. DOI: 10.1007/s10988-023-09389-y.

Alexiadou, A. (2021). "Definite Plural Generics in English: Evidence from De-adjectival Nominalization". In: *Determiners and Quantifiers*. Ed. by C. Gianollo, K. von Heusinger and M. Napoli. Vol. 44. Brill, pp. 29–54. DOI: 10.1163/9789004473324_003.

Allport, G. W. (1954). *The Nature of Prejudice*. Addison-Wesley.

Amiraz, O. (2020). "Non-Maximality and Homogeneity: Parallels between Collective Predicates and Absolute Adjectives". In: *Proceedings of Sinn und Bedeutung* 24.1, pp. 63–82. DOI: 10.3765/salt.v1i0.5344.

Aparicio, H., M. Xiang and C. Kennedy (2015). "Processing Gradable Adjectives in Context: A Visual World Study". In: *Semantics and Linguistic Theory* 25, pp. 413–432. DOI: 10.3765/salt.v25i0.3128.

Armstrong, D. M. (1983). *What Is a Law of Nature?* Cambridge University Press.

Asher, N. and M. Morreau (1995). "What Some Generic Sentences Mean". In: *The Generic Book*. Ed. by G. N. Carlson and F. J. Pelletier. University of Chicago Press, pp. 300–338.

Asher, N. and F. J. Pelletier (2013). "More Truths about Generic Truth". In: *Genericity*. Ed. by A. Mari, C. Beyssade and F. Del Prete. Oxford University Press, pp. 312–333.

Augurzky, P., M. Bonnet, R. Breheny, A. Cremers, C. Ebert, C. Mayr, J. Romoli, M. Steinbach and Y. Sudo (2023). "Putting Plural Definites into Context". In: *Proceedings of Sinn und Bedeutung* 27, pp. 19–32. DOI: 10.18148/sub/2023.v27.1050.

Bach, K. (1994). "Conversational Impliciture". In: *Mind & Language* 9.2, pp. 124–162. DOI: 10.1111/j.1468-0017.1994.tb00220.x.

Bach, K. (2002). "Seemingly Semantic Intuitions". In: *Meaning and Truth: Investigations in Philosophical Semantics*. Ed. by J. K. Campbell, M. O'Rourke and D. Shier. Seven Bridges Press, pp. 21–33.

Bar-Lev, M. E. (2021). "An Implicature Account of Homogeneity and Non-maximality". In: *Linguistics and Philosophy* 44.5, pp. 1045–1097. DOI: 10.1007/s10988-020-09308-5.

Barker, C. (1992). "Group Terms in English: Representing Groups as Atoms". In: *Journal of Semantics* 9.1, pp. 69–93. DOI: 10.1093/jos/9.1.69.

Barwise, J. and R. Cooper (1981). "Generalized Quantifiers and Natural Language". In: *Linguistics and Philosophy* 4.2, pp. 159–219. DOI: 10.1007/BF00350139.

Barwise, J. and J. Perry (1983). *Situations and Attitudes*. MIT Press.

Bastian, B. and N. Haslam (2006). "Psychological Essentialism and Stereotype Endorsement". In: *Journal of Experimental Social Psychology* 42.2, pp. 228–235. DOI: 10.1016/j.jesp.2005.03.003.

Behrens, L. (2005). "Genericity from a Cross-Linguistic Perspective". In: *Linguistics* 43.2, pp. 275–344. DOI: 10.1515/ling.2005.43.2.275.

Benitez, J., E. Foster-Hanson and M. Rhodes (2025). "Generic References to Gender Predict Essentialism and Stereotyping Even When They Express Counter-Stereotypic Ideas". In: *Journal of Experimental Psychology: General*. DOI: 10.1037/xge0001871.

Benitez, J., R. A. Leshin and M. Rhodes (2022). "The Influence of Linguistic Form and Causal Explanations on the Development of Social Essentialism". In: *Cognition* 229, p. 105246. DOI: 10.1016/j.cognition.2022.105246.

Bodenhausen, G. V. (1990). "Stereotypes as Judgmental Heuristics: Evidence of Circadian Variations in Discrimination". In: *Psychological Science* 1.5, pp. 319–322. DOI: 10.1111/j.1467-9280.1990.tb00226.x.

Boneh, N. (2019). "Dispositions and Characterizing Sentences". In: *Glossa: a journal of general linguistics* 4.1. DOI: 10.5334/gjgl.540.

Boneh, N. and E. Doron (2013). "Hab and Gen in the Expression of Habituality". In: *Genericity*. Ed. by A. Mari, C. Beyssade and F. Del Prete. Oxford University Press, pp. 176–191. DOI: 10.1093/acprof:oso/9780199691807.003.0006.

Boolos, G. (1984). "To Be Is to Be a Value of a Variable (or to Be Some Values of Some Variables)". In: *The Journal of philosophy* 81.8, pp. 430–449. DOI: 10.2307/2026308.

Borik, O. and M. T. Espinal (2015). "Reference to Kinds and to Other Generic Expressions in Spanish: Definiteness and Number". In: *The Linguistic Review* 32.2, pp. 167–225. DOI: 10.1515/tlr-2014-0023.

Bosse, A. (2021). "Generics: Some (Non) Specifics". In: *Synthese* 199.5, pp. 14383–14401. DOI: 10.1007/s11229-021-03426-8.

Bosse, A. (2024). "Stereotyping and Generics". In: *Inquiry* 67.10, pp. 3876–3892. DOI: 10.1080/0020174X.2022.2074879.

Breheny, R. (2005). "Exhaustivity, Homogeneity and Definiteness". In: *Proceedings of the Fifteenth Amsterdam Colloquium*. Amsterdam Colloquium. Ed. by P. Dekker and M. Franke, pp. 59–64.

Brisson, C. (1998). "Distributivity, Maximality, and Floating Quantifiers". PhD thesis. Rutgers University.

Brisson, C. (2003). "Plurals, 'all', and the Nonuniformity of Collective Predication". In: *Linguistics and Philosophy* 26.2, pp. 129–184. DOI: 10.1023/A:1022771705575.

Buchanan, R. (2010). "A Puzzle about Meaning and Communication". In: *Noûs* 44.2, pp. 340–371. DOI: 10.1111/j.1468-0068.2010.00743.x.

Burge, T. (1973). "Reference and Proper Names". In: *The Journal of Philosophy* 70.14, pp. 425–439. DOI: 10.2307/2025107.

Büring, D. (2016). *Intonation and Meaning*. Oxford University Press. DOI: 10.1093/acprof:oso/9780199226269.001.0001.

Burnett, H. (2016). *Gradability in Natural Language: Logical and Grammatical Foundations*. Oxford University Press. DOI: 10.1093/acprof:oso/9780198724797.001.0001.

Burton-Roberts, N. (1977). "Generic Sentences and Analyticity". In: *Studies in Language* 1.2, pp. 155–196. DOI: 10.1075/sl.1.2.02bur.

Carlson, G. N. (1977). "Reference to Kinds in English". PhD thesis. University of Massachusetts Amherst.

Carlson, G. N. (1982). "Generic Terms and Generic Sentences". In: *Journal of Philosophical Logic* 11.2, pp. 145–181.

Carlson, G. N. (1989). "On the Semantic Composition of English Generic Sentences". In: *Properties, Types and Meaning*. Ed. by G. Chierchia, B. H. Partee and R. Turner. Kluwer Academic, pp. 167–192.

Carlson, G. N. (1995). "Truth Conditions of Generic Sentences: Two Contrasting Views". In: *The Generic Book*. Ed. by G. N. Carlson and F. J. Pelletier. University of Chicago Press, pp. 224–237.

Carlson, G. N. (1999). "Evaluating Generics". In: *Studies in the linguistic sciences* 29.1, pp. 13–24.

Carlson, G. N. (2008). "Patterns in the Semantics of Generic Sentences". In: *Time and Modality*. Ed. by J. Guéron and J. Lecarme. Springer Netherlands, pp. 17–38.

Champollion, L. (2015). "Stratified Reference: The Common Core of Distributivity, Aspect, and Measurement". In: *Theoretical Linguistics* 41.3–4, pp. 109–149. DOI: 10.1515/tl-2015-0008.

Champollion, L. (2017). *Parts of a Whole: Distributivity as a Bridge between Aspect and Measurement*. Oxford University Press. DOI: 10.1093/oso/9780198755128.001.0001.

Champollion, L. (2020). "Distributivity, Collectivity, and Cumulativity". In: *The Wiley Blackwell Companion to Semantics*. John Wiley & Sons, Ltd, pp. 1–38. DOI: 10.1002/9781118788516.sem021.

Champollion, L., D. Bumford and R. Henderson (2019). "Donkeys under Discussion". In: *Semantics and Pragmatics* 12.1, pp. 1–50. DOI: 10.3765/sp.12.1.

Cheng, K.-Y. (2011). "A New Look at the Problem of Rule-Following: A Generic Perspective". In: *Philosophical Studies* 155.1, pp. 1–21. DOI: 10.1007/s11098-010-9561-9.

Chierchia, G. (1995). "Individual-Level Predicates as Inherent Generics". In: *The Generic Book*. Ed. by G. N. Carlson and F. J. Pelletier. University of Chicago Press, pp. 176–223.

Chierchia, G. (1998). "Reference to Kinds across Languages". In: *Natural Language Semantics* 6.4, pp. 339–405. DOI: 10.1023/A:1008324218506.

Chierchia, G. (2022). "'People Are Fed up; Don't Mess with Them.' Non-quantificational Arguments and Polarity Reversals". In: *Journal of Semantics* 39.3, pp. 475–521. DOI: 10.1093/jos/ffac006.

Cimpian, A., A. C. Brandone and S. A. Gelman (2010). "Generic Statements Require Little Evidence for Acceptance but Have Powerful Implications". In: *Cognitive Science* 34.8, pp. 1452–1482. DOI: 10.1111/j.1551-6709.2010.01126.x.

Cimpian, A. and E. M. Markman (2009). "Information Learned from Generic Language Becomes Central to Children's Biological Concepts: Evidence from Their Open-Ended Explanations". In: *Cognition* 113.1, pp. 14–25. DOI: 10.1016/j.cognition.2009.07.004.

Cimpian, A. and E. M. Markman (2011). "The Generic/Nongeneric Distinction Influences How Children Interpret New Information About Social Others". In: *Child Development* 82.2, pp. 471–492. DOI: 10.1111/j.1467-8624.2010.01525.x.

Cimpian, A. and E. Salomon (2014). "The Inherence Heuristic: An Intuitive Means of Making Sense of the World, and a Potential Precursor to Psychological Essentialism". In: *Behavioral and Brain Sciences* 37.5, pp. 461–480. DOI: 10.1017/S0140525X13002197.

Claveau, F. and J. Girard (2019). "Generic Generalizations in Science". In: *Erkenntnis* 84.4 (4), pp. 839–859. DOI: 10.1007/s10670-018-9983-x.

Cohen, A. (1999a). "Generics, Frequency Adverbs, and Probability". In: *Linguistics and Philosophy* 22.3, pp. 221–253. DOI: 10.1023/A:1005497727784.

Cohen, A. (1999b). *Think Generic!: The Meaning and Use of Generic Sentences*. CSLI Publications.

Cohen, A. (2001). "On the Generic Use of Indefinite Singulars". In: *Journal of Semantics* 18.3, pp. 183–209. DOI: 10.1093/jos/18.3.183.

Cohen, A. (2004). "Existential Generics". In: *Linguistics and Philosophy* 27.2, pp. 137–168. DOI: 10.1023/B:LING.0000016441.89129.3d.

Cohen, A. (2020). *Something out of Nothing: The Semantics and Pragmatics of Implicit Quantification*. Brill.

Cohen, A. (2022). "Genericity". In: *Oxford Research Encyclopedia of Linguistics*. Ed. by M. Aronoff. Oxford University Press. DOI: 10.1093/acrefore/9780199384655.013.326.

Cohen, A. and N. Erteschik-Shir (2002). "Topic, Focus, and the Interpretation of Bare Plurals". In: *Natural Language Semantics* 10.2, pp. 125–165. DOI: 10.1023/A:1016576614139.

Collins, J. (2015). "Genericity as a Unitary Psychological Phenomenon: An Argument from Linguistic Diversity". In: *Ratio* 28.4, pp. 369–394. DOI: 10.1111/rati.12102.

Collins, J. (2018). "Genericity sans Gen". In: *Mind & Language* 33.1, pp. 34–64. DOI: 10.1111/mila.12168.

Comrie, B. (1976). *Aspect: An Introduction to the Study of Verbal Aspect and Related Problems*. Cambridge University Press.

Condoravdi, C. (1992). "Strong and Weak Novelty and Familiarity". In: *Semantics and Linguistic Theory*, pp. 17–38. DOI: 10.3765/salt.v2i0.3038.

Condoravdi, C. (1994). "Descriptions in Context". PhD thesis. Yale University.

Connolly, A. C., J. A. Fodor, L. R. Gleitman and H. Gleitman (2007). "Why Stereotypes Don't Even Make Good Defaults". In: *Cognition* 103.1, pp. 1–22. DOI: 10.1016/j.cognition.2006.02.005.

Cooper, R. (1996). "The Role of Situations in Generalized Quantifiers". In: *The Handbook of Contemporary Semantic Theory*. Ed. by S. Lappin. Basil Blackwell, pp. 65–86.

Corblin, F. (2013). "The Roots of Genericity: Indefinite Singulars vs Definite Plurals". In: *Genericity*. Ed. by A. Mari, C. Beyssade and F. del Prete. University Press, pp. 352–371.

Dahl, Ö. (1975). "On Generics". In: *Formal Semantics of Natural Language*. Ed. by E. L. Keenan. Cambridge University Press, pp. 99–111. DOI: 10.1017/CBO9780511897696.009.

Dalrymple, M., M. Kanazawa, S. Mchombo and S. Peters (1994). "What Do Reciprocals Mean?" In: *Semantics and Linguistic Theory*, pp. 61–78. DOI: 10.3765/salt.v4i0.2466.

Dayal, V. (2004). "Number Marking and (in)Definiteness in Kind Terms". In: *Linguistics and Philosophy* 27.4, pp. 393–450. DOI: 10.1023/B:LING.0000024420.80324.67.

Dayal, V. (2006). "Semantic Variation and Pleonastic Determiners: The Case of the Plural Definite Generic". In: *Proceedings of V Asian GLOW*. Citeseer.

Dayal, V. (2013). "On the Existential Force of Bare Plurals across Languages". In: *From Grammar to Meaning: The Spontaneous Logicality of Language*. Ed. by C. Cecchetto and I. Caponigro. Cambridge University Press, pp. 49–80. DOI: 10.1017/CBO9781139519328.006.

Declerck, R. (1986). "The Manifold Interpretations of Generic Sentences". In: *Lingua* 68.2–3, pp. 149–188. DOI: 10.1016/0024-3841(86)90002-1.

Declerck, R. (1991). "The Origins of Genericity". In: *Linguistics* 29.1, pp. 79–102. DOI: 10.1515/ling.1991.29.1.79.

DeJesus, J. M., M. A. Callanan, G. Solis and S. A. Gelman (2019). "Generic Language in Scientific Communication". In: *Proceedings of the National Academy of Sciences* 116.37, pp. 18370–18377. DOI: 10.1073/pnas.1817706116.

Delfitto, D. (2006). "Bare Plurals". In: *The Blackwell Companion to Syntax*. John Wiley & Sons, Ltd, pp. 214–259. DOI: 10.1002/9780470996591.ch8.

Delgrande, J. P. (1987). "A First-Order Conditional Logic for Prototypical Properties". In: *Artificial Intelligence* 33.1, pp. 105–130. DOI: 10.1016/0004-3702(87)90053-1.

Diesing, M. (1988). "Bare Plural Subjects and the Stage/Individual Contrast". In: *Genericity in Natural Language*, pp. 107–154.

Diesing, M. (1992). "Bare Plural Subjects and the Derivation of Logical Representations". In: *Linguistic Inquiry* 23.3, pp. 353–380.

Dobrovie-Sorin, C. (2001). "Adverbs of Quantification and Genericity". In: *Empirical Issues in Formal Syntax and Semantics 4*. Colloque de Syntaxe et Sémantique à Paris. Presses Universitaires de la Sorbonne.

Dobrovie-Sorin, C. and I. Giurgea (2021). *Majority Quantification and Quantity Superlatives: A Crosslinguistic Analysis of Most*. Oxford University Press. DOI: 10.1093/oso/9780198791249.001.0001.

Dowty, D. R. (1987). "Collective Predicates, Distributive Predicates and 'All'". In: *Proceedings of the Third Eastern States Conference on Linguistics* 3. Ed. by F. Marshall, A. Miller and Z.-s. Zhang, pp. 97–117.

Drewery, A. (2005). "The Logical Form of Universal Generalizations". In: *Australasian Journal of Philosophy* 83.3, pp. 373–393. DOI: 10.1080/00048400500191966.

Dutilh Noaves, C. (2008). "Tarski's Hidden Theory of Meaning: Sentences Say Exactly One Thing". In: *Unity, Truth and the Liar : The Modern Relevance of Medieval Solutions to the Liar Paradox*. Ed. by S. Rahman, T. Tulenheimo and E. Genot. Springer, pp. 41–63.

Ebert, K. H. (1971a). *Referenz, Sprechsituation und die bestimmten Artikel in einem nordfriesischen Dialekt (Fering)*. Nordfriisk Inst.

Ebert, K. H. (1971b). "Zwei Formen des bestimmten Artikels". In: *Probleme und Fortschritte der Transformationsgrammatik*. Ed. by D. Wunderlich. Max Hueber, pp. 159–174.

Eckardt, R. (1999). "Normal Objects, Normal Worlds and the Meaning of Generic Sentences". In: *Journal of Semantics* 16.3, pp. 237–278. DOI: 10.1093/jos/16.3.237.

Elbourne, P. D. (2005). *Situations and Individuals*. MIT Press.

Elbourne, P. D. (2008). "Ellipsis Sites as Definite Descriptions". In: *Linguistic Inquiry* 39.2, pp. 191–220. DOI: 10.1162/ling.2008.39.2.191.

Elbourne, P. D. (2013). *Definite Descriptions*. Oxford University Press.

Enç, M. (1986). "Towards a Referential Analysis of Temporal Expressions". In: *Linguistics and Philosophy* 9.4, pp. 405–426. DOI: 10.1007/BF00603217.

Erteschik-Shir, N. (1998). *The Dynamics of Focus Structure*. Cambridge University Press. DOI: 10.1017/CBO9780511519949.

Fara, D. G. (2015). "Names Are Predicates". In: *The Philosophical Review* 124.1, pp. 59–117. DOI: 10.1215/00318108-2812660.

Farkas, D. and H. de Swart (2007). "Article Choice in Plural Generics". In: *Lingua* 117.9, pp. 1657–1676. DOI: 10.1016/j.lingua.2006.06.011.

Ferreira, M. (2005). "Bare Habituals and Plural Definite Descriptions". In: *Proceedings of Sinn und Bedeutung* 9, pp. 102–115. DOI: 10.18148/sub/2005.v9i0.721.

Ferreira, M. (2016). "The Semantic Ingredients of Imperfectivity in Progressives, Habituals, and Counterfactuals". In: *Natural Language Semantics* 24.4, pp. 353–397. DOI: 10.1007/s11050-016-9127-2.

Fiengo, R. and R. May (1994). *Indices and Identity*. MIT Press.

Fine, K. (2017). "Truthmaker Semantics". In: *A Companion to the Philosophy of Language*. John Wiley & Sons, Ltd, pp. 556–577. DOI: 10.1002/9781118972090.ch22.

von Fintel, K. (1994). "Restrictions on Quantifier Domains". PhD thesis. University of Massachusetts Amherst.

von Fintel, K. (1997). "Bare Plurals, Bare Conditionals, and *Only*". In: *Journal of Semantics* 14.1, pp. 1–56. DOI: 10.1093/jos/14.1.1.

von Fintel, K. and I. Heim (2011). *Intensional Semantics*. URL: https://github.com/fintelkai/fintel-heim-intensional-notes/blob/main/IntensionalSemantics.pdf.

Florio, S. and Ø. Linnebo (2021). *The Many and the One: A Philosophical Study of Plural Logic*. Oxford University Press. DOI: 10.1093/oso/9780198791522.001.0001.

Fodor, J. D. (1970). "The Linguistic Description of Opaque Contexts". PhD thesis. Massachusetts Institute of Technology.

Frazier, L. and C. Clifton (2006). "Ellipsis and Discourse Coherence". In: *Linguistics and Philosophy* 29.3, pp. 315–346. DOI: 10.1007/s10988-006-0002-3.

Frege, G. (1960). "Begriffsschrift". In: *Translations from the Philosophical Writings of Gottlob Frege*. Trans. by P. T. Geach and M. Black. 2nd ed. Basil Blackwell, pp. 1–20.

Gajewski, J. R. (2005). "Neg-Raising: Polarity and Presupposition". PhD thesis. Massachusetts Institute of Technology.

Gelman, S. A., E. A. Ware and F. Kleinberg (2010). "Effects of Generic Language on Category Content and Structure". In: *Cognitive Psychology* 61.3, pp. 273–301. DOI: 10.1016/j.cogpsych.2010.06.001.

Gerstner-Link, C. and M. Krifka (1993). "Genericity". In: *Halbband*. De Gruyter Mouton, pp. 966–978. DOI: 10.1515/9783110095869.1.14.966.

Geurts, B. (1985). "Generics". In: *Journal of Semantics* 4.3, pp. 247–255. DOI: 10.1093/jos/4.3.247.

Gigerenzer, G. and W. Gaissmaier (2011). "Heuristic Decision Making". In: *Annual Review of Psychology* 62.1, pp. 451–482. DOI: 10.1146/annurev-psych-120709-145346.

Gillon, B. S. (1987). "The Readings of Plural Noun Phrases in English". In: *Linguistics and Philosophy* 10.2, pp. 199–219. DOI: 10.1007/BF00584318.

Goodman, N. (1983). *Fact, Fiction, & Forecast*. 4th ed. Athlone Press.

Greenberg, Y. (2003). *Manifestations of Genericity*. Taylor & Francis Group.

Greenberg, Y. (2007). "Exceptions to Generics: Where Vagueness, Context Dependence and Modality Interact". In: *Journal of Semantics* 24.2, pp. 131–167. DOI: 10.1093/jos/ffm002.

Grice, H. P. (1975). "Logic and Conversation". In: *Syntax and Semantics*. Ed. by P. Cole and J. Morgan. Academic Press, pp. 41–58.

Grimau, B. (2021). "Structured Plurality Reconsidered". In: *Journal of Semantics* 38.1, pp. 145–193. DOI: 10.1093/jos/ffaa012.

Guerrini, J. (2024). *Revisiting Kind Predication*. URL: https://lingbuzz.net/lingbuzz/008172 (visited on 26/06/2024). Pre-published.

Hale, B. (2020). "What Makes True Universal Statements True?" In: *Essence and Existence*. Ed. by B. Hale, J. Leech and K. Fine. Oxford University Press, pp. 104–122. DOI: 10.1093/oso/9780198854296.003.0007.

Hampton, J. A. (2012). "Generics as Reflecting Conceptual Knowledge". In: *Recherches linguistiques de Vincennes* 41 (41), pp. 9–24. DOI: 10.4000/rlv.2036.

Hampton, J. A., A. Passanisi and M. L. Jönsson (2011). "The Modifier Effect and Property Mutability". In: *Journal of Memory and Language* 64.3, pp. 233–248. DOI: 10.1016/j.jml.2010.12.001.

Haslanger, S. (2012). "Ideology, Generics, and Common Ground". In: *Resisting Reality*. Oxford University Press, pp. 446–478. DOI: 10.1093/acprof:oso/9780199892631.003.0017.

Hawkins, J. A. (1978). *Definiteness and Indefiniteness: A Study in Reference and Grammaticality Prediction*. Croom Helm.

Heim, I. (1982). "The Semantics of Definite and Indefinite Noun Phrases". PhD thesis. University of Massachusetts Amherst.

Heim, I. (2011). "Definiteness and Indefiniteness". In: *Semantics: An International Handbook of Natural Language Meaning*. Vol. 2. De Gruyter Mouton, pp. 996–1025. DOI: 10.1515/9783110255072.996.

Heim, I. and A. Kratzer (1998). *Semantics in Generative Grammar*. Blackwell.

Higginbotham, J. and B. Schein (1989). "Plurals". In: *North East Linguistics Society* 19.1.

Hoeltje, M. (2017). "Generics and Ways of Being Normal". In: *Linguistics and Philosophy* 40.2, pp. 101–118. DOI: 10.1007/s10988-016-9203-6.

Hoicka, E., J. Saul, E. Prouten, L. Whitehead and R. Sterken (2021). "Language Signaling High Proportions and Generics Lead to Generalizing, but Not Essentializing, for Novel Social Kinds". In: *Cognitive Science* 45.11, e13051. DOI: 10.1111/cogs.13051.

Horn, L. R. (2021). "Ceteris Paribusiness: On the Power of Salient Exceptions". In: *Inquiries in Philosophical Pragmatics: Issues in Linguistics*. Ed. by F. Macagno and A. Capone. Springer, pp. 7–31. DOI: 10.1007/978-3-030-56696-8_2.

Hussak, L. J. and A. Cimpian (2018). "Memory Accessibility Shapes Explanation: Testing Key Claims of the Inherence Heuristic Account". In: *Memory & Cognition* 46.1 (1), pp. 68–88. DOI: 10.3758/s13421-017-0746-8.

Izumi, Y. (2012). "The Semantics of Proper Names and Other Bare Nominals". PhD thesis. University of Maryland, College Park.

Jönsson, M. L. and J. A. Hampton (2006). "The Inverse Conjunction Fallacy". In: *Journal of Memory and Language* 55.3, pp. 317–334. DOI: 10.1016/j.jml.2006.06.005.

Jönsson, M. L. and J. A. Hampton (2012). "The Modifier Effect in Within-Category Induction: Default Inheritance in Complex Noun Phrases". In: *Language and Cognitive Processes* 27.1, pp. 90–116. DOI: 10.1080/01690965.2010.544107.

Kahneman, D. (2003). "A Perspective on Judgment and Choice: Mapping Bounded Rationality". In: *American Psychologist* 58, pp. 697–720. DOI: 10.1037/0003-066X.58.9.697.

Keenan, E. L. (1971). "Names, Quantifiers, and the Sloppy Identity Problem". In: *Papers in Linguistics* 4.2, pp. 211–232. DOI: 10.1080/08351817109370257.

Keller, J. (2005). "In Genes We Trust: The Biological Component of Psychological Essentialism and Its Relationship to Mechanisms of Motivated Social Cognition". In: *Journal of Personality and Social Psychology* 88.4, pp. 686–702. DOI: 10.1037/0022-3514.88.4.686.

Kim, C. and J. T. Runner (2009). "Strict Identity, Coherence, and Parallelism in VP Ellipsis". In: *Semantics and Linguistic Theory*, pp. 275–287. DOI: 10.3765/salt.v19i0.2535.

Kirkpatrick, J. R. (2023). "Generic Conjunctivitis". In: *Linguistics and Philosophy* 46, pp. 379–428. DOI: 10.1007/s10988-022-09371-0.

Kirkpatrick, J. R. (2024a). "Are Generics Quantificational?" In: *Synthese* 204.1, p. 17. DOI: 10.1007/s11229-024-04653-5.

Kirkpatrick, J. R. (2024b). "The Acquisition of Generics". In: *Mind & Language* 39.4, pp. 492–517. DOI: 10.1111/mila.12499.

Kirkpatrick, J. R. (2025). "Generic Excluded Middle". In: *Philosophers' Imprint* 25.0. DOI: 10.3998/phimp.4085.

Kiss, K. É. (1998). "On Generic and Existential Bare Plurals and the Classification of Predicates". In: *Events and Grammar*. Ed. by S. Rothstein. Springer Netherlands, pp. 145–162. DOI: 10.1007/978-94-011-3969-4_7.

Kleinschmidt, S. and J. Ross (2013). "Repeatable Artwork Sentences and Generics". In: *Art and Abstract Objects*. Ed. by C. Mag Uidhir. Oxford University Press, pp. 125–157. DOI: 10.1093/acprof:oso/9780199691494.003.0007.

Koslicki, K. (1999). "Genericity and Logical Form". In: *Mind & Language* 14.4, pp. 441–467. DOI: 10.1111/1468-0017.00120.

Kratzer, A. (1995). "Stage-Level and Individual-Level Predicates". In: *The Generic Book*. Ed. by G. N. Carlson and F. J. Pelletier. University of Chicago Press, pp. 125–175.

Kratzer, A. (2007). *Situations in Natural Language Semantics*. In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta and U. Nodelman. Fall 2007. Metaphysics Research Lab, Stanford University.

Krifka, M. (1987). "An Outline of Genericity". In: *SNS-Bericht 87-25*. Seminar Fur Naturlich-Sprachliche Systeme, pp. 1–34.

Krifka, M. (1995). "Focus and the Interpretations of Generic Sentences". In: *The Generic Book*. Ed. by G. N. Carlson and F. J. Pelletier. University of Chicago Press, pp. 238–264.

Krifka, M. (1996). "Pragmatic Strengthening in Plural Predications and Donkey Sentences". In: *Semantics and Linguistic Theory* 6.0, pp. 136–153. DOI: 10.3765/salt.v6i0.2769.

Krifka, M. and R. Musan (2012). "Information Structure: Overview and Linguistic Issues". In: *The Expression of Information Structure*. De Gruyter Mouton, pp. 1–44. DOI: 10.1515/9783110261608.1.

Krifka, M., F. J. Pelletier, G. N. Carlson, A. ter Meulen, G. Chierchia and G. Link (1995). "Genericity: An Introduction". In: *The Generic Book*. Ed. by G. N. Carlson and F. J. Pelletier. University of Chicago Press, pp. 1–124.

Kripke, S. (1980). *Naming and Necessity*. Harvard University Press.

Križ, M. (2015). "Aspects of Homogeneity in the Semantics of Natural Language". PhD thesis. University of Vienna.

Križ, M. (2016). "Homogeneity, Non-Maximality, and 'All'". In: *Journal of Semantics* 33.3, pp. 493–539. DOI: 10.1093/jos/ffv006.

Križ, M. and E. Chemla (2015). "Two Methods to Find Truth-Value Gaps and Their Application to the Projection Problem of Homogeneity". In: *Natural Language Semantics* 23.3, pp. 205–248. DOI: 10.1007/s11050-015-9114-z.

Križ, M. and B. Spector (2021). "Interpreting Plural Predication: Homogeneity and Non-Maximality". In: *Linguistics and Philosophy* 44.5, pp. 1131–1178. DOI: 10.1007/s10988-020-09311-w.

Kroch, A. S. (1974). "The Semantics of Scope in English". PhD thesis. Massachusetts Institute of Technology.

Kuhn, J. (2020). "Gather/Numerous as a Mass/Count Opposition". In: *Natural Language Semantics* 28.3, pp. 225–253. DOI: 10.1007/s11050-020-09163-x.

Langton, R., S. Haslanger and L. Anderson (2012). "Language and Race". In: *Routledge Companion to Philosophy of Language*. Ed. by G. Russell and D. G. Fara. Taylor & Francis Group.

Lasersohn, P. (1995). *Plurality, Conjunction and Events*. Kluwer Academic.

Lasersohn, P. (1998). "Generalized Distributivity Operators". In: *Linguistics and Philosophy* 21.1, pp. 83–93. DOI: 10.1023/A:1005317815339.

Lasersohn, P. (1999). "Pragmatic Halos". In: *Language* 75.3, pp. 522–551.

Lawler, J. M. (1973). "Studies in English Generics". PhD thesis. University of Michigan.

Lazaridou-Chatzigoga, D. and L. Stockall (2013). "Genericity, Exceptions and Domain Restriction: Experimental Evidence from Comparison with Universals". In: *Proceedings of Sinn und Bedeutung* 17, pp. 325–343.

Lazaridou-Chatzigoga, D., L. Stockall and N. Katsos (2019). "Contextualising Generic and Universal Generalisations: Quantifier Domain Restriction and the Generic Overgeneralisation Effect". In: *Journal of Semantics* 36.4, pp. 617–664. DOI: 10.1093/jos/ffz009.

Le Bruyn, B., H. de Swart and J. Zwarts (2017). "Bare Nominals". In: *Oxford Research Encyclopedia of Linguistics*. DOI: 10.1093/acrefore/9780199384655.013.399.

Lee, J. and A. Nguyen (2022). "What's Positive and Negative about Generics: A Constrained Indexical Approach". In: *Philosophical Studies* 179.5, pp. 1739–1761. DOI: 10.1007/s11098-021-01727-9.

Lemeire, O. (2021). "Falsifying Generic Stereotypes". In: *Philosophical Studies* 178.7, pp. 2293–2312. DOI: 10.1007/s11098-020-01555-3.

Lerner, A. and S.-J. Leslie (2016). "Generics and Experimental Philosophy". In: *A Companion to Experimental Philosophy*. Ed. by W. Buckwalter and J. Sytsma. Oxford: Wiley-Blackwell, pp. 404–416.

Leshin, R. A., S.-J. Leslie and M. Rhodes (2021). "Does It Matter How We Speak About Social Kinds? A Large, Preregistered, Online Experimental Study of How Language Shapes the Development of Essentialist Beliefs". In: *Child Development* 92.4, e531–e547. DOI: 10.1111/cdev.13527.

Leslie, S.-J. (2007). "Generics and the Structure of the Mind". In: *Philosophical Perspectives* 21, pp. 375–403.

Leslie, S.-J. (2008). "Generics: Cognition and Acquisition". In: *The Philosophical Review* 117.1, pp. 1–47. DOI: 10.1215/00318108-2007-023.

Leslie, S.-J. (2014). "Carving Up the Social World with Generics". In: *Oxford Studies in Experimental Philosophy*. Ed. by J. Knobe, T. Lombrozo and S. Nichols. Oxford University Press, pp. 208–230. DOI: 10.1093/acprof:oso/9780198718765.003.0009.

Leslie, S.-J. (2015). "Generics Oversimplified". In: *Noûs* 49.1, pp. 28–54. DOI: 10.1111/nous.12039.

Leslie, S.-J. (2017). "The Original Sin of Cognition: Fear, Prejudice, and Generalization". In: *The Journal of Philosophy* 114.8, pp. 393–421. DOI: 10.5840/jphil2017114828.

Leslie, S.-J., S. Khemlani and S. Glucksberg (2011). "Do All Ducks Lay Eggs? The Generic Overgeneralization Effect". In: *Journal of Memory and Language* 65.1, pp. 15–31. DOI: 10.1016/j.jml.2010.12.005.

Leslie, S.-J. and A. Lerner (2022). *Generic Generalizations*. In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta and U. Nodelman. Fall 2022. Metaphysics Research Lab, Stanford University.

Lewis, D. (1975a). "Adverbs of Quantification". In: *Formal Semantics of Natural Language*. Ed. by E. L. Keenan. Cambridge University Press, pp. 3–15. DOI: 10.1017/CBO9780511897696.003.

Lewis, D. (1975b). "Languages and Language". In: *Minnesota Studies in the Philosophy of Science*. Ed. by K. Gunderson. Vol. 7. University of Minnesota Press, pp. 3–35.

Lewis, D. (1979). "Scorekeeping in a Language Game". In: *Journal of Philosophical Logic* 8.1, pp. 339–359. DOI: 10.1007/BF00258436.

Lewis, D. K. (1991). *Parts of Classes*. Basil Blackwell.

Liebesman, D. (2011). "Simple Generics". In: *Noûs* 45.3, pp. 409–442. DOI: 10.1111/j.1468-0068.2010.00774.x.

Liebesman, D. and O. Magidor (2025). *Property Versatility and Copredication*. Oxford University Press. DOI: 10.1093/9780191915291.001.0001.

Liebesman, D. and R. K. Sterken (2021). "Generics and the Metaphysics of Kinds". In: *Philosophy Compass* 16.7, e12754. DOI: 10.1111/phc3.12754.

Link, G. (1983). "The Logical Analysis of Plurals and Mass Terms: A Lattice-theoretical Approach". In: *Meaning, Use, and Interpretation of Language*. Ed. by R. Bäuerle, C. Schwarze and A. von Stechow. De Gruyter, pp. 302–323. DOI: 10.1515/9783110852820.302.

Link, G. (1984). "Hydras: On the Logic of Relative Constructions with Multiple Heads". In: *Varieties of Formal Semantics: Proceedings of the 4th Amsterdam Colloquium*. Ed. by F. Landman and F. Veltman. Foris.

Link, G. (1987). "Generalized Quantifiers and Plurals". In: *Generalized Quantifiers: Linguistic and Logical Approaches*. Ed. by P. Gärdenfors. D. Reidel Pub. Co, pp. 151–180.

Linnebo, Ø. (2022). "Generality Explained". In: *The Journal of Philosophy* 119.7, pp. 349–379. DOI: 10.5840/jphil2022119725.

Löbner, S. (2000). "Polarity in Natural Language: Predication, Quantification and Negation in Particular and Characterizing Sentences". In: *Linguistics and Philosophy* 23.3, pp. 213–308. DOI: 10.1023/a:1005571202592.

Longobardi, G. (1994). "Reference and Proper Names: A Theory of N-Movement in Syntax and Logical Form". In: *Linguistic Inquiry* 25.4, pp. 609–665.

Magri, G. (2012). "No Need for a Dedicated Theory of the Distribution of Readings of English Bare Plurals". In: *Semantics and Linguistic Theory* 22.0, pp. 383–402. DOI: 10.3765/salt.v22i0.2644.

Magri, G. (2014). "An Account for the Homogeneity Effect Triggered by Plural Definites and Conjunction Based on Double Strengthening". In: *Pragmatics, Semantics and the Case of Scalar Implicatures*. Ed. by S. P. Reda. Palgrave Macmillan UK, pp. 99–145. DOI: 10.1057/9781137333285_5.

Maienborn, C. (2005). "On the Limits of the Davidsonian Approach: The Case of Copula Sentences". In: *Theoretical Linguistics* 31.3, pp. 275–316. DOI: 10.1515/thli.2005.31.3.275.

Malamud, S. A. (2012). "The Meaning of Plural Definites: A Decision-Theoretic Approach". In: *Semantics and Pragmatics* 5, pp. 1–58. DOI: 10.3765/sp.5.3.

Mari, A., C. Beyssade and F. Del Prete (2013). "Introduction". In: *Genericity*. Ed. by A. Mari, C. Beyssade and F. Del Prete. Oxford University Press, pp. 1–92. DOI: 10.1093/acprof:oso/9780199691807.003.0001.

Matthewson, L. (2001). "Quantification and the Nature of Crosslinguistic Variation". In: *Natural Language Semantics* 9.2, pp. 145–189. DOI: 10.1023/A:1012492911285.

McConnell-Ginet, S. (2012). "Generic Predicates and Interest-Relativity". In: *Canadian Journal of Linguistics/Revue canadienne de linguistique* 57.2, pp. 261–287. DOI: 10.1017/S0008413100004771.

McKeever, M. and R. K. Sterken (2021). "Social and Political Aspects of Generic Language and Speech". In: *The Routledge Handbook of Social and Political Philosophy of Language*. Ed. by R. K. Sterken and J. Khoo. Routledge, pp. 259–280. DOI: 10.4324/9781003164869-20.

McKillen, A. (2016). "On the Interpretation of Reflexive Pronouns". PhD thesis. McGill University.

Milsark, G. L. (1974). "Existential Sentences in English". PhD thesis. Massachusetts Institute of Technology.

Moltmann, F. (1997). "Intensional Verbs and Quantifiers". In: *Natural Language Semantics* 5.1, pp. 1–52. DOI: 10.1023/A:1008245409172.

Moltmann, F. (2013). "On the Distinction between Abstract States, Concrete States, and Tropes". In: *Genericity*. Ed. by A. Mari, C. Beyssade and F. Del Prete. Oxford University Press, pp. 293–311. DOI: 10.1093/acprof:oso/9780199691807.003.0011.

Moltmann, F. (2015). "Quantification with Intentional and with Intensional Verbs". In: *Quantifiers, Quantifiers, and Quantifiers: Themes in Logic, Metaphysics, and Language*. Ed. by A. Torza. Springer International Publishing, pp. 141–168. DOI: 10.1007/978-3-319-18362-6_8.

Moltmann, F. (2016). "Plural Reference and Reference to a Plurality: Linguistic Facts and Semantic Analyses". In: *Unity and Plurality: Logic, Philosophy, and Linguistics*. Ed. by M. Carrara, A. Arapinis and F. Moltmann. Oxford University Press, pp. 93–120. DOI: 10.1093/acprof:oso/9780198716327.003.0006.

Musan, R. (1997). *On the Temporal Interpretation of Noun Phrases*. Garland Science.

Musan, R. (1999). "Temporal Interpretation and Information-Status of Noun Phrases". In: *Linguistics and Philosophy* 22.6, pp. 621–661. DOI: 10.1023/A:1005435423259.

Neale, S. (1990). *Descriptions*. MIT Press.

Neufeld, E. (2022). "Psychological Essentialism and the Structure of Concepts". In: *Philosophy Compass* 17.5, e12823. DOI: 10.1111/phc3.12823.

Neufeld, E., A. Bosse, G. Del Pinal and R. Sterken (2025). "Giving Generic Language Another Thought". In: *WIREs Cognitive Science* 16.1, e70000. DOI: 10.1002/wcs.70000.

Nguyen, A. (2020). "The Radical Account of Bare Plural Generics". In: *Philosophical Studies* 177.5, pp. 1303–1331. DOI: 10.1007/s11098-019-01254-8.

Nickel, B. (2008). "Generics and the Ways of Normality". In: *Linguistics and Philosophy* 31.6, pp. 629–648. DOI: 10.1007/s10988-008-9049-7.

Nickel, B. (2016). *Between Logic and the World: An Integrated Theory of Generics*. Oxford University Press. DOI: 10.1093/acprof:oso/9780199640003.001.0001.

Nickel, B. (2017). "Generics". In: *A Companion to the Philosophy of Language*. Ed. by B. Hale, C. Wright and A. Miller. John Wiley & Sons, Ltd, pp. 437–462. DOI: 10.1002/9781118972090.ch18.

Nicolas, D. and J. D. Payton (2025). "Superplurals Analyzed Away". In: *Inquiry*, pp. 1–34. DOI: 10.1080/0020174X.2024.2440765.

Noyes, A. and F. C. Keil (2019). "Generics Designate Kinds but Not Always Essences". In: *Proceedings of the National Academy of Sciences of the United States of America* 116.41, pp. 20354–20359.

Noyes, A. and F. C. Keil (2020). "There Is No Privileged Link between Kinds and Essences Early in Development". In: *Proceedings of the National Academy of Sciences* 117.20, pp. 10633–10635. DOI: 10.1073/pnas.2003627117.

Oliver, A. and T. Smiley (2016). *Plural Logic*. 2nd ed. Oxford University Press. DOI: 10.1093/acprof:oso/9780198744382.001.0001.

Ong, M. and A. Brasoveanu (2014). "Strict and Sloppy Reflexives in VP Ellipsis". In: *Empirical Issues in Syntax and Semantics*. Ed. by C. Pińon. Vol. 10, pp. 251–268.

Parent, T. (2021). "A Dilemma about Kinds and Kind Terms". In: *Synthese* 198.12, pp. 2987–3006. DOI: 10.1007/s11229-019-02278-7.

Partee, B. H. (1991). "Topic, Focus and Quantification". In: *Semantics and Linguistic Theory* 1.0, pp. 159–188. DOI: 10.3765/salt.v1i0.2918.

Partee, B. H. (1995). "Quantificational Structures and Compositionality". In: *Quantification in Natural Languages*. Ed. by E. Bach, E. Jelinek, A. Kratzer and B. H. Partee. Springer Netherlands, pp. 541–601. DOI: 10.1007/978-94-017-2817-1_17.

Pelletier, F. J. and N. Asher (1997). "Generics and Defaults". In: *Handbook of Logic and Language*. Ed. by J. van Benthem and A. ter Meulen. North Holland, pp. 1125–1177. DOI: 10.1016/B978-044481714-3/50025-4.

Peters, U. (2021). "Science Communication and the Problematic Impact of Descriptive Norms". In: *The British Journal for the Philosophy of Science*. DOI: 10.1086/715001.

Pietroski, P., J. Lidz, T. Hunter and J. Halberda (2009). "The Meaning of 'Most': Semantics, Numerosity and Psychology". In: *Mind & Language* 24.5, pp. 554–585. DOI: 10.1111/j.1468-0017.2009.01374.x.

Prasada, S. (2000). "Acquiring Generic Knowledge". In: *Trends in Cognitive Sciences* 4.2, pp. 66–72. DOI: 10.1016/S1364-6613(99)01429-1.

Prasada, S. and E. M. Dillingham (2006). "Principled and Statistical Connections in Common Sense Conception". In: *Cognition* 99.1, pp. 73–112. DOI: 10.1016/j.cognition.2005.01.003.

Prasada, S., S. Khemlani, S.-J. Leslie and S. Glucksberg (2013). "Conceptual Distinctions amongst Generics". In: *Cognition* 126.3, pp. 405–422. DOI: 10.1016/j.cognition.2012.11.010.

Ralston, T. (2024a). "Generic Explanations for Generics". In: *Proceedings of the 24th Amsterdam Colloquium*, pp. 291–297.

Ralston, T. (2024b). "Reconceptualising the Psychological Theory of Generics". In: *Philosophical Studies* 181.11, pp. 2973–2995. DOI: 10.1007/s11098-024-02242-3.

Ramsey, F. P. (1990). "Truth and Probability". In: *Philosophical Papers*. Ed. by D. H. Mellor. Cambridge University Press, pp. 52–109.

Rayo, A. (2002). "Word and Objects". In: *Noûs* 36.3, pp. 436–464. DOI: 10.1111/1468-0068.00379.

Rayo, A. (2007). "Beyond Plurals". In: *Absolute Generality*. Ed. by G. Uzquiano and A. Rayo. Oxford University Press, pp. 220–254.

Read, S. (2009). "Plural Signification and the Liar Paradox". In: *Philosophical Studies* 145.3, pp. 363–375. DOI: 10.1007/s11098-008-9236-y.

Reinhart, T. (1981). "Pragmatics and Linguistics: an analysis of Sentence Topics". In: *Philosophica* 27.0. DOI: 10.21825/philosophica.82606.

Restall, G. (2018). *Accommodation, Inference, Generics and Pejoratives*. URL: https://consequently.org/handouts/accommodation-unimelb-handout-2018.pdf (visited on 20/12/2022).

Rhodes, M., S. A. Gelman and S.-J. Leslie (2025). "How Generic Language Shapes the Development of Social Thought". In: *Trends in Cognitive Sciences* 29.2, pp. 122–132. DOI: 10.1016/j.tics.2024.09.012.

Rhodes, M., S.-J. Leslie, K. Saunders, Y. Dunham and A. Cimpian (2018). "How Does Social Essentialism Affect the Development of Inter-Group Relations?" In: *Developmental Science* 21.1, e12509. DOI: 10.1111/desc.12509.

Rhodes, M., S.-J. Leslie and C. M. Tworek (2012). "Cultural Transmission of Social Essentialism". In: *Proceedings of the National Academy of Sciences* 109.34, pp. 13526–13531. DOI: 10.1073/pnas.1208951109.

Rhodes, M. and K. Moty (2020). "What Is Social Essentialism and How Does It Develop?" In: *The Development of Social Essentialism*. Ed. by M. Rhodes. Academic Press, pp. 1–30. DOI: 10.1016/bs.acdb.2020.05.001.

Rieppel, M. (2015). "Pluralities and Plural Logic". In: *Analysis* 75.3, pp. 504–514. DOI: 10.1093/analys/anv010.

Rimmel, L. (2004). "Habituals and Generic Quantification". In: *WCCFL 23: Proceedings of the 23rd West Coast Conference on Formal Linguistics*. Ed. by V. Chand, A. Kelleher, A. J. Rodríguez and B. Schmeiser. Cascadilla Press.

Ritchie, K. (2017). "Plural and Collective Noun Phrases". In: *The Routledge Handbook of Collective Intentionality*. Ed. by M. Jankovic and K. Ludwig. Routledge, pp. 464–475.

Ritchie, K. and J. Knobe (2020). "Kindhood and Essentialism: Evidence from Language". In: *The Development of Social Essentialism*. Ed. by M. Rhodes. Academic Press, pp. 133–164. DOI: 10.1016/bs.acdb.2020.05.005.

Roberts, C. (1987). "Modal Subordination, Anaphora, and Distributivity". PhD thesis. University of Massachusetts Amherst.

Roberts, C. (1996). "Information Structure in Discourse: Towards an Integrated Formal Theory of Pragmatics". In: *Semantics and Pragmatics* 5, pp. 1–69.

Robinson, H. M. (2005). "Unexpected (in)Definiteness: Plural Generic Expressions in Romance". PhD thesis. Rutgers, The State University of New Jersey.

van Rooij, R. and K. Schulz (2020). "Generics and Typicality: A Bounded Rationality Approach". In: *Linguistics and Philosophy* 43.1, pp. 83–117. DOI: 10.1007/s10988-019-09265-8.

Rooth, M. (1992). "A Theory of Focus Interpretation". In: *Natural Language Semantics* 1.1, pp. 75–116.

Rosola, M. and F. Cella (2020). "Generics and Epistemic Injustice". In: *Ethical Theory and Moral Practice* 23.5, pp. 739–754. DOI: 10.1007/s10677-020-10095-y.

Rothbart, M., S. Fulero, C. Jensen, J. Howard and P. Birrell (1978). "From Individual to Group Impressions: Availability Heuristics in Stereotype Formation". In: *Journal of Experimental Social Psychology* 14.3, pp. 237–255. DOI: 10.1016/0022-1031(78)90013-6.

Rumfitt, I. (2005). "Plural Terms: Another Variety of Reference?" In: *Thought, Reference, and Experience: Themes from the Philosophy of Gareth Evans*. Ed. by J. L. Bermúdez. Oxford University Press, pp. 84–123. DOI: 10.1093/acprof:oso/9780199248964.003.0004.

Sag, I. A. (1976). "Deletion and Logical Form". PhD thesis. Massachusetts Institute of Technology.

Salomon, E. and A. Cimpian (2014). "The Inherence Heuristic as a Source of Essentialist Thought". In: *Personality and Social Psychology Bulletin* 40.10, pp. 1297–1315. DOI: 10.1177/0146167214541659.

Sauerland, U. and P. Elbourne (2002). "Total Reconstruction, PF Movement, and Derivational Order". In: *Linguistic Inquiry* 33.2, pp. 283–319. DOI: 10.1162/002438902317406722.

Saul, J. (2023). "Are Generics Especially Pernicious?" In: *Inquiry* 66.9, pp. 1682–1699.

Schiller, H. I. (2023). "Genericity and Inductive Inference". In: *Philosophy of Science* 90.3, pp. 539–557. DOI: 10.1017/psa.2022.94.

Schubert, L. K. and F. J. Pelletier (1987). "Problems in the Representation of the Logical Form of Generics, Plurals and Mass Nouns". In: *New Directions in Semantics*. Ed. by E. Lepore, pp. 385–451.

Schwarz, F. (2009). "Two Types of Definites in Natural Language". PhD thesis. University of Massachusetts Amherst.

Schwarz, F. (2012). "Situation Pronouns in Determiner Phrases". In: *Natural Language Semantics* 20.4, pp. 431–475. DOI: 10.1007/s11050-012-9086-1.

Schwarz, F. (2013a). "Maximality and Definite Plurals - Experimental Evidence". In: *Proceedings of Sinn und Bedeutung* 17, pp. 509–526.

Schwarz, F. (2013b). "Two Kinds of Definites Cross-linguistically". In: *Language and Linguistics Compass* 7.10, pp. 534–559. DOI: 10.1111/lnc3.12048.

Schwarzschild, R. (1994). "Plurals, Presuppositions and the Sources of Distributivity". In: *Natural Language Semantics* 2.3, pp. 201–248. DOI: 10.1007/BF01256743.

Schwarzschild, R. (1996). *Pluralities*. Kluwer Academic.

Schwarzschild, R. (1999). "Givenness, AvoidF and Other Constraints on the Placement of Accent". In: *Natural Language Semantics* 7.2, pp. 141–177. DOI: 10.1023/A:1008370902407.

Sharvy, R. (1980). "A More General Theory of Definite Descriptions". In: *The Philosophical Review* 89.4, pp. 607–624. DOI: 10.2307/2184738.

Smith, C. S. (1991). *The Parameter of Aspect*. Kluwer Academic.

Soames, S. (1986). "Incomplete Definite Descriptions". In: *Notre Dame Journal of Formal Logic* 27.3, pp. 349–375. DOI: 10.1305/ndjfl/1093636680.

Stanley, J. and Z. G. Szabó (2000). "On Quantifier Domain Restriction". In: *Mind & Language* 15.2–3, pp. 219–261. DOI: 10.1111/1468-0017.00130.

Sterken, R. K. (2015a). "Generics in Context". In: *Philosopher's Imprint* 15.21, pp. 1–30.

Sterken, R. K. (2015b). "Generics, Content and Cognitive Bias". In: *Analytic Philosophy* 56.1, pp. 75–93. DOI: 10.1111/phib.12056.

Sterken, R. K. (2015c). "Leslie on Generics". In: *Philosophical Studies* 172.9, pp. 2493–2512. DOI: 10.1007/s11098-014-0429-2.

Sterken, R. K. (2016). "Generics, Covert Structure and Logical Form". In: *Mind & Language* 31.5, pp. 503–529. DOI: 10.1111/mila.12118.

Sterken, R. K. (2017). "The Meaning of Generics". In: *Philosophy Compass* 12.8, e12431. DOI: 10.1111/phc3.12431.

de Swart, H. (1996). "(In)Definites and Genericity". In: *Quantifiers, Deduction, and Context*. Ed. by M. Kanazawa, C. J. Piñón and H. de Swart. CSLI Publications, pp. 171–194.

de Swart, H. (2002). "Topichood and the Stage/Individual Distinction". In: *Context-Dependence in the Analysis of Linguistic Meaning*. Brill, pp. 489–499. DOI: 10.1163/9789004487222_035.

Syrett, K., E. Bradley, C. Kennedy and J. Lidz (2006). "Shifting Standards: Children's Understanding of Gradable Adjectives". In: *Proceedings of the Inaugural Conference on Generative Approaches to Language Acquisition — North America*. Vol. 2, pp. 353–364. DOI: 10.7282/T39W0HHS.

Syrett, K., C. Kennedy and J. Lidz (2010). "Meaning and Context in Children's Understanding of Gradable Adjectives". In: *Journal of Semantics* 27.1, pp. 1–35. DOI: 10.1093/jos/ffp011.

Tanaka, H. (2011). "Syntactic Identity and Ellipsis". In: *The Linguistic Review* 28.1, pp. 79–110. DOI: 10.1515/tlir.2011.003.

Taub, A. (1989). "Collective Predicates, Aktionsarten and 'All'". In: *Papers on Quantification*. Ed. by E. Bach, A. Kratzer and B. H. Partee, pp. 1–32.

Teichman, M. (2023). "The Sophisticated Kind Theory". In: *Inquiry* 66.9, pp. 1613–1654. DOI: 10.1080/0020174X.2016.1267407.

Tessler, M. H. and N. D. Goodman (2019). "The Language of Generalization". In: *Psychological Review* 126.3, pp. 395–436. DOI: 10.1037/rev0000142.

Thakral, R. (2018). "Generics and Weak Necessity". In: *Inquiry*, pp. 1–28. DOI: 10.1080/0020174X.2018.1426683.

Tonhauser, J. (2020). "Temporal Properties of Noun Phrases". In: *The Wiley Blackwell Companion to Semantics*. John Wiley & Sons, Ltd, pp. 1–25. DOI: 10.1002/9781118788516.sem126.

Tversky, A. and D. Kahneman (1982). "Availability: A Heuristic for Judging Frequency and Probability". In: *Judgment under Uncertainty: Heuristics and Biases*. Ed. by A. Tversky, D. Kahneman and P. Slovic. Cambridge University Press, pp. 163–178. DOI: 10.1017/CBO9780511809477.012.

Uzquiano, G. (2004). "The Supreme Court and the Supreme Court Justices: A Metaphysical Puzzle". In: *Noûs* 38.1, pp. 135–153. DOI: 10.1111/j.1468-0068.2004.00465.x.

Vasilyeva, N., A. Gopnik and T. Lombrozo (2020). "When Generic Language Does Not Promote Psychological Essentialism". In: *Proceedings of the 42nd Annual Conference of the Cognitive Science Society*. Annual Meeting of the Cognitive Science Society.

Vasilyeva, N., A. Gopnik and T. Lombrozo (2018). "The Development of Structural Thinking about Social Categories". In: *Developmental Psychology* 54, pp. 1735–1744. DOI: 10.1037/dev0000555.

Vasilyeva, N. and T. Lombrozo (2020). "Structural Thinking about Social Categories: Evidence from Formal Explanations, Generics, and Generalization". In: *Cognition* 204, p. 104383. DOI: 10.1016/j.cognition.2020.104383.

Vendler, Z. (1957). "Verbs and Times". In: *The Philosophical Review* 66.2, pp. 143–160. DOI: 10.2307/2182371.

Vlach, F. (1993). "Temporal Adverbials, Tenses and the Perfect". In: *Linguistics and Philosophy* 16.3, pp. 231–283. DOI: 10.1007/BF00985970.

de Vries, H. (2017). "Two Kinds of Distributivity". In: *Natural Language Semantics* 25.2, pp. 173–197. DOI: 10.1007/s11050-017-9133-z.

Weicker, M. and P. Schulz (2020). "Not All Gradable Adjectives Are Vague – Experimental Evidence from Adults and Children". In: *Proceedings of Sinn und Bedeutung* 24.2, pp. 406–422. DOI: 10.18148/sub/2020.v24i2.905.

Westerståhl, D. (2015). "Generalized Quantifiers in Natural Language Semantics". In: *The Handbook of Contemporary Semantic Theory*. John Wiley & Sons, Ltd, pp. 7–39. DOI: 10.1002/9781118882139.ch1.

Wilkinson, K. J. (1991). "Studies in the Semantics of Generic Noun Phrases". PhD thesis. University of Massachusetts Amherst.

Williams, E. S. (1977). "Discourse and Logical Form". In: *Linguistic Inquiry* 8.1, pp. 101–139.

Williamson, T. (2003). "Everything". In: *Philosophical Perspectives* 17, pp. 415–465.

Williamson, T. (2013). *Modal Logic as Metaphysics*. Oxford University Press. DOI: 10.1093/acprof:oso/9780199552078.001.0001.

Williamson, T. (2020). *Suppose and Tell: The Semantics and Heuristics of Conditionals*. Oxford University Press. DOI: 10.1093/oso/9780198860662.001.0001.

Winter, Y. (1997). "Choice Functions and the Scopal Semantics of Indefinites". In: *Linguistics and Philosophy* 20.4, pp. 399–467. DOI: 10.1023/A:1005354323136.

Winter, Y. (2001). *Flexibility Principles in Boolean Semantics: The Interpretation of Coordination, Plurality, and Scope in Natural Language*. MIT Press. DOI: 10.7551/mitpress/3034.001.0001.

Winter, Y. and R. Scha (2015). "Plurals". In: *The Handbook of Contemporary Semantic Theory*. John Wiley & Sons, Ltd, pp. 77–113. DOI: 10.1002/9781118882139.ch3.

Wodak, D. and S.-J. Leslie (2017). "The Mark of the Plural: Generic Generalizations and Race". In: *The Routledge Companion to Philosophy of Race*. Ed. by P. C. Taylor, L. M. Alcoff and L. Anderson. Routledge, pp. 277–289.

Wodak, D., S.-J. Leslie and M. Rhodes (2015). "What a Loaded Generalization: Generics and Social Cognition". In: *Philosophy Compass* 10.9, pp. 625–635. DOI: 10.1111/phc3.12250.

Yi, B.-U. (2005). "The Logic and Meaning of Plurals. Part I". In: *Journal of Philosophical Logic* 34.5, pp. 459–506. DOI: 10.1007/s10992-005-0560-9.

Yoon, Y. (1996). "Total and Partial Predicates and the Weak and Strong Interpretations". In: *Natural Language Semantics* 4.3, pp. 217–236.

Zamparelli, R. (2002). "Definite and Bare Kind-Denoting Noun Phrases". In: *Romance Languages and Linguistic Theory 2000: Selected Papers from 'Going Romance' 2000, Utrecht, 30 November–2 December*. Ed. by C. Beyssade, R. Bok-Bennema, F. Drijkoningen and P. Monachesi. John Benjamins Publishing Company.

Zwicky, A. M. and J. M. Sadock (1975). "Ambiguity Tests and How to Fail Them". In: *Syntax and Semantics*. Vol. 4. Brill, pp. 1–36. DOI: 10.1163/9789004368828_002.

DECLARATION

I declare that all work in this thesis is my own and is the result of research carried out while enrolled for the degree of DPhil at the University of Oxford.

Oxford, May 2025

Tom Ralston