Title
Normality Judgments: Structure and Content

Author
Lavin, Andrew Shawn

Publication Date
2019

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Los Angeles

Normality Judgments:

Content and Structure

A dissertation submitted in partial satisfaction of the

requirements for the degree

Doctor of Philosophy in Philosophy

by

Andrew Lavin

2019

ABSTRACT OF THE DISSERTATION

Normality Judgments: Content and Structure

by

Andrew Lavin

Doctor of Philosophy in Philosophy

University of California, Los Angeles, 2019

Professor Gavin Lawrence, Co-Chair

Professor Calvin G. Normore, Co-Chair

A normality judgment is a type of generalization that ascribes a *normal* feature or property to a kind. The explanatory link account of normality—an account of what it is to be a normal feature or property and therefore an account of the content of a normality judgment—is defended in this dissertation. According to the explanatory link account, a feature is normal for a kind if and only if an individual's possession of that feature is explained by its being a member of that kind. The account is inspired by the Aristotelian notion that being a carpenter explains one's house-building activities whereas being a grammarian wouldn't explain the same. The second chapter defends this form of explanation—called herein "kind explanation." The dissertation closes with a study of the extensional relationship between normality and statistical regularity, and finally with a study of the inferential profile of normality judgments.

The dissertation of Andrew Lavin is approved

Sheldon Russell Smith

Joshua David Armstrong

Gavin Lawrence, Committee Co-Chair

Calvin G. Normore, Committee Co-Chair

University of California, Los Angeles

2019

For Kellie Marie

TABLE OF CONTENTS

ACKNOWLEDGEMENTS

No one deserves more credit than Kellie Marie, who has supported my graduate studies for the past 8 years. On top of working full-time throughout most of our marriage and pursuing her own goals in addition, she has cared for an ailing, depressed, anxious, and often fatigued spouse—abnormal conditions if ever there were any. This dissertation is, therefore, her accomplishment as well. I'm so grateful for her partnership through this journey.

I want to thank my advisors, Calvin Normore and Gavin Lawrence, who gave me the freedom to explore the questions I wanted to explore in the way I wanted to explore them. I am a sometimes fiercely independent person, and I appreciate that their style as advisors made room for me to craft a project that matched my eccentricities. The result is a dissertation that feels like a natural fit with my philosophical personality, but one that is informed by their desire that I approach simple questions with great care and circumspection. I moreover appreciate the hours they spent with me conversing, the time they spent reading my unmastered prose, and the energy they put into having me on their to-do lists. A conversation with Calvin or Gavin is not an opportunity to be missed. Their reputations throughout the philosophy world as sharp, but amiable and collegial academics are wholly deserved. They are masters of the craft and I am grateful to have learned from them and studied under them.

I owe a great deal of gratitude, also, to the other members of my committee (official and unofficial). I thank Sheldon Smith for being a supportive, but critical voice. His work in particular on ceteris paribus laws has been particularly instructive and his generosity with his time have been a boon to my philosophical training and my sense of mastery over my subject. I thank Josh Armstrong for being a source of great practical advice in how to shape my research around the goal of going on the job market; and I'm grateful to have had someone who is not wholly unsympathetic to my conclusions, but whose approach is importantly different and so whose perspective is

invaluable in helping me get a sense of what my work looks like from the outside. I thank Katie Elliott for offering the perspective of a well-trained metaphysician, who approaches my project with the analytic clarity that I sometimes lose sight of as I get lost among the weeds of my own thinking.

I am grateful to the UCLA philosophy department as a whole, which provided me with a generous amount of funding and support, and most importantly facilitated an intellectual training program for me over these past five years that has shaped my mind in qualitative and significant ways. Virtually everyone on faculty has contributed in some way to my growth as a philosopher and I am grateful to each of them.

I am grateful in particular to John Carriero, whose mentorship proved invaluable over the first couple of years of the program. His careful and systematic approach to philosophy influenced me greatly and I am glad to have had the time to explore early modern and medieval thought with such a well-tuned intellect. His intuitive and clear writing style has also influenced my writing in a quite positive way. When I'm tempted to write in a dense and analytic style, I'm reminded that he is able to do careful and precise philosophical work in prose that virtually anyone could understand on the first pass (though so remembering does not always produce that effect in my own writing). Furthermore, his care for the graduate and undergraduate students around him is readily evident and I have learned much from him as a philosopher and a teacher.

I am grateful to members of the Albritton Society at UCLA for their helpful commentary during presentations of parts of the dissertation, particularly at the early stages of development. Thank you to the graduate student community more broadly, who have given me informal philosophical feedback on my ideas throughout the years.

I presented parts of the dissertation at the Boulder Medieval Philosophy Conference hosted by the philosophy department at the University of Colorado, Boulder; at the Berkeley-Stanford-Davis graduate conference; at the Eastern Division Meeting of the American Philosophical Association; at

a meeting of the Chico Triad, and in other venues. I'm grateful in particular to the organizers, but also to the attendees of these sessions.

I would like to thank the UCLA Graduate Division for their financial support in the form of travel stipends and the Graduate Research Mentorship and Graduate Summer Research Mentorship.

Thank you to my friends and family for your support, encouragement, and love. I have an amazing community standing behind me and I am ever grateful. A special thanks to Cooie Grey-Lavin, who always believed in, cheered for, and supported me; and who always let me be my own person—whether I was being normal or not.

VITA

Education

2018	UCLA: C. Phil in Philosophy
2017	UCLA: Graduate Certificate in Writing Pedagogy
2015	UCLA: MA in Philosophy
2013	San Jose State University: MA in Philosophy
2010	CSU, Chico: BA in Comparative Religion, minors in Philosophy and Ethics
2008	Butte College: AA in Social and Behavioral Sciences

Normality Judgments: Content and Structure

Introduction

Judgments are mental actions. As a class, they comprise comparisons, attributions, descriptions, classifications, characterizations, and, among many others, generalizations. A generalization is an attribution of a property or feature to a kind. The project here is to understand a particular sort of generalization called a 'normality judgment.' What can one say about the family of generalizations that—at least intuitively—centrally involve the concept of normality or normalcy?

Since the project is about a particular sort of mental action, the bearing it has on the empirical world is limited. The goal is to offer an analysis of *what it means* to say that some feature is normal for a kind of thing. With that in mind, I present no grand picture of what is and is not normal. This is no treatise on how to find out what is normal or what sorts of cases should count as normal, nor indeed does any argument hinge on the truth of a particular normality judgment. The intent is to offer an account of *what we are arguing about* if we have a disagreement about what is normal, or *what we are saying* when we say that some feature is normal for a kind. No claims are made about contentious normalities of deep social and moral importance: human sexuality, human neural structure, human gender, human disabilities, or human subspecies groupings more generally. One must take great care not to divorce such discussions from the moral and ethical dimensions inherent in deciding whether a sort of human being is normal or abnormal. Furthermore, one must take great care to ensure that all are in agreement that no way of treating another human being can immediately be justified by appeal to the abnormality of their condition or status. Oppression is never justifiable, whether the oppressed turn out to be normal or abnormal members of the human species in this or that respect. Of course, I have views about each of these "first order" cases, but the analysis on offer here doesn't decide such issues, nor is it meant to.

Normality has played a variety of roles in philosophical theorizing to date. A number of philosophers use normality in the development of a semantics for generic sentences—for instance, Nickel (2008, 2016), Eckardt (1999), Asher, Morreau, and Pelletier (Asher and Morreau 1991; Pelletier and Asher 1997). Epistemologists have employed some notion akin to normality or explicitly called normalcy or normality in the development of their theories. For example, see Ball (2013), Burge (2003), Leplin (2009), Martin Smith (2007, 2010, 2014), and Sosa (2007, 2009, 2010). Philosophers of language—most prominently Millikan (1984)—have argued that normality in a special evolutionary sense plays an important role in grounding the meaning of thoughts and words. Leslie (2012a) has argued that generic judgments—which appear to share almost all of the same features as normality judgments—articulate a form of generalization that is the default mode of generalization for the human mind. Schurz (2001a, 2001b, 2009a, 2009b, 2014) has argued that normality plays a central role in our understanding of *ceteris paribus* laws and more broadly defeasible generalizations about dynamical systems. Gardiner (2015, 2018) has explored the importance of the distinction between probabilistic judgments and normality judgments in the epistemology of law and in philosophical judgments themselves. Normality has played a perennial role in theories of dispositions—cf. in particular Haukioja (2007), Hauska (2008), and Pettit (1999). Yalcin (2014) has argued that there is no properly epistemic reading of deontic modals like *ought* and *should* and that instead the non-deontic reading should be understood in terms of normality. Veltman (1996) and Bastiaanse (Bastiaanse and Veltman 2016) have argued that default semantics can elucidate the inferential structure of normality because the default *is* the normal.

The normality judgment itself—though conceptualized slightly differently and always under different names—has been the subject of a host of particularly developmental psychological studies exploring the structure and nature of human cognition in relation to kinds and their individual members: for a sampling, see (Ahn et al. 2001; Bear and Knobe 2017; Cimpian et al. 2010; Gelman

and Bloom 2007; Gelman and Markman 1986; Khemlani et al. 2007, 2012; Leslie 2008; Leslie et al. 2009; Prasada et al. 2013; Prasada 2000; Prasada and Dillingham 2006, 2009; Rhodes and Gelman 2009). Kahneman and Miller (1986) have explored the hypothesis that the mind is fundamentally concerned with comparing perceived situations to an internalized norm or sense of what is normal. Boorse (1975, 1977, 2014) has argued that normality is central to an understanding of health and has offered a statistical account of normality. Among others, Kingma (2007, 2009, 2016) has critiqued this view on the grounds that the reference classes for a statistical account cannot be specified without tacit appeal to normality or function or other normative concepts. The concept of normality as it relates to clinical psychology has also received a variety of treatments, including in Abraham Kaplan (1967), Canguilhem (1978), and Murphy (1966).

This survey is incomplete, but it suffices to establish the wide range of applications of the concept. Normality is thus an unquestionably important if not central concept to philosophy and psychology. Nevertheless, few have undertaken to give an account of what it is to be normal. Strößner (2014, 2015) offers an account based on majority; Schurz (2001a) offers an account based on the nature of an open dynamical system; Millikan (1984) offers at least a series of comments out of which an account based on historical reproduction and selection can be reconstructed; Wachbroit (1994) offers the beginnings of an account of biological normality; Smith (2007, 2010, 2014) offers some preliminary reflections on what it means to be normal and particularly on the role of normality in epistemology, but falls short of a complete account; Nickel (2008, 2016) offers the most direct account, in which he appeals to the explicability of the presence of a property in a population to ground the normality of that property for that kind. None of these accounts, though, discusses what I take to be a central feature of normality: that the kind for which a property is normal explains a normal individual's having that property. One has a normal property *because* one is the type of thing for which it is normal to have that property.

Aside from normality judgments, there are at least two other sorts of generalizations. Some generalizations have to do with *every* individual in a set of things. *All* natural numbers are either even or odd. *No* prime numbers are divisible by 4. *Every* dog is a mammal. *Each* student must complete the comprehensive exam before being awarded a master's degree. *Only* those who pass the state Bar are allowed to practice law. Other generalizations have to do with some more or less precise proportion of the population of members of a set or kind of thing. *Most* species of locust hatch in prime intervals. *A small percentage* or mosquitoes carry West Nile Virus. *Some* mammals lay eggs. The first chapter will be occupied with the task of setting up the distinction between nomic or universal judgments, statistical or frequency judgments, and normality or normic judgments.

Chapters two, three, and four form the core argumentative component of the dissertation. In the second chapter, the sufficiency of a particular form of explanation is defended. When a question like "why does that thing have stripes like that?" receives the answer, "because it's a zebra," there are two things that might be going on. First, it may be a mere pragmatic phenomenon that only *appears* to be an explanation.[1] Second, it may be an explanation in itself. If it is the latter, then again we are faced with the question: is it a sufficient explanation all by itself or is it only the beginnings of an explanation? This chapter seeks to show that this sort of explanation—'kind explanation'—is indeed a form of explanation proper and further that it is a *sufficient* explanation on its own. It won't answer every question, but there is a particular sort of question this sort of explanation indeed answers and better yet: answers fully.

In the third chapter—the flagship of the dissertation—the novel account of normality (and therefore of the content of normality judgments) is presented and defended in comparison with the closest extant account available in the literature: that of Bernhard Nickel. Nickel supposes that normal properties—the sort of normality that is somewhat independent of statistical prevalence—

[1] This may be so even if explanation is itself a mere pragmatic phenomenon.

are properties the presence of which in the population of members of a kind is explicable by appeal to the right sorts of explanatory practices. At least this is Nickel's theoretical account of normality posited in service of giving an account of the semantics of generic sentences. I think, to the contrary, that a normal property is one the possession of which by an individual is explained by that individual's being a member of the kind for which the property is normal. Nickel's account, I worry, invites a sort of circularity into the analysis in its appeal to deference to experts (like biologists, for instance) in specifying what the "right sorts of explanatory strategies" are.

The fourth chapter rounds out the positive account in offering responses to a number of pressing worries for the explanatory link account of normality. In addition, the first part of the chapter reflects on the pragmatic elements of normality and the role of normality and normality judgments in explanatory strategies. Normality plays an important role in determining explanatory priority, but also plays a central role in the formulation of explanatory accounts and narratives.

The fifth and penultimate chapter explores Thomas Aquinas' position—partially inherited from Aristotle—on the extensional relation between normality and regularity. Is what is normal what one should also expect to see *regularly*? Sure, the vast majority of dogs have four legs and it is normal for dogs to have four legs, but how general is this extensional link? After drawing Aquinas' considered view out of his corpus, the chapter addresses two problematics: 1) there appear to be a variety of counterexamples to the view that the normal is always regular and the regular is always normal, and 2) it's unclear how *Aquinas in particular* could think this, given his commitment elsewhere to the idea that the world is full of rival processes that interfere with one another and introduce abnormalities throughout nature.

Finally, the sixth chapter addresses the "structure" in the dissertation's subtitle "Content and Structure." Having settled on a view about the content of normality judgments, an exploration of their inferential profile is in order: what sorts of inferences license normality judgments and what

sorts of inferences do they in turn license? The basic inferential structure of normality judgments is presented, then the chapter discusses a few controversial inferences defended by some in the literature, and finally the chapter offers a critique of various mistakes that researchers have made in the literature.

The result is a study of normality judgments as situated within the conceptual space of other generalizations and judgments more broadly; the content of normality judgments and the sense of the concept of normality; the extensional relationship between normality and statistical regularity; and finally of the inferential structure of normality judgments. Though there is much more to say about normality judgments, the basic groundwork will be in place at the close of the project.

Before beginning, it is worth clarifying that an *outcome* is a feature, an activity or behavior, a relation, a condition or set of conditions, the result of a process, or something similar. 'Property' would have been the common word to use here, but it has connotations with contingent possession or exemplification and with a sort of Platonism or at minimum reism about features, activities, and the like; all of which make an Aristotelian such as myself nervous. I have therefore for the most part avoided talking in terms of properties. Throughout the dissertation, 'outcome' and 'feature' are often used interchangeably as stand-ins for the list contained in this paragraph.

Chapter One: Regularity, Necessity, Normality

Humans are in the business of making judgments. We judge that it is too soon to take the roast out of the oven; that it is warm outside today; that penguins and puffins are more similar than are robins and parrots; that busses are regularly behind schedule; that nothing moves faster than the speed of light; and that computers turn on when we depress their 'on' buttons. It is the latter three sorts of judgment that will occupy us in this dissertation. The judgment about busses is a judgment about what frequently, usually, or ordinarily occurs—it is a statistical or frequency judgment. The judgment about the universal speed limit, alternatively, is about what *must* occur (physically speaking) or in this case more properly what *cannot* occur—it is a nomic judgment or a universal law. Our final judgment is a judgment about what is *supposed* to happen or what will happen absent *failure* or *interference*—it is what I'm calling a normality judgment. Each of these is a class of generalization: a sub-species of judgments wherein one predicates some property, say, of a population or kind. The aim of this dissertation is simply to understand the last of these three kinds of generalization more fully. In this chapter, the goal is to distinguish it from other kinds of judgments and attempt to characterize what ties this *prima facie* family of judgment together.

To structure this study, we'll focus on two central problems. If the goal is to present normality judgments as a distinct form of judgment, then we need to see how they are a) different from statistical and nomic judgments, and b) the same in the sense that all nomic and all statistical judgments are respectively the same. With that in mind, there are two worries to overcome:

> (Problem of Reduction) It is possible that the content and inferential profile of any normality judgment might be captured with either a statistical judgment or a nomic judgment.
>
> (Problem of Pluralism) It is possible that there is no coherence to be found among the diverse array of normality judgments and therefore that there is no one family of judgments we might call 'normality judgments.'

The task here, is therefore, to attempt to show either that each is false, or at minimum that each is implausible. We'll address each in time. Sections 1 and 2 will lay the conceptual groundwork for the discussion to follow, introducing statistical and nomic judgments in section 2. Section 3 will introduce normality judgments and attempt to handle the problem of reduction. Section 4 will deal with the problem of pluralism, and section 5 will address the question of whether an account of normality judgments writ large can be developed from extant accounts of generic sentences or *ceteris paribus* laws.

1. The Conceptual Space

Some preliminaries are in order. Judgments are predications—i.e. they are expressed by propositions or sentences. Generalizations are a sub-class of judgment that link a kind (or over an individual through time) to an *Outcome*—a feature, activity, relation, or condition. So, any normality judgment will comprise a general subject and an outcome that is supposed to be normal for that general subject or members thereof. There are at least two other kinds of generalization: nomic or law-like generalizations, and statistical or frequency generalizations.

On the world side, a 'regularity' is the object of a statistical generalization, a 'necessity' the object of a nomic generalization, and a 'normality' the object of a normality generalization. These, as I conceive of it are the truthmakers of these kinds of judgments. They are relations between kinds (or individuals through time) and features or properties.

Additionally, there is an important, but orthogonal, distinction among judgments: loose and strict. Loose judgments simply aren't meant *seriously* or *literally*, even though they are meant to be true enough to support coordination with other agents, practical ends, and pragmatic inferences. They aren't meant to be *false* or immune to truth valuation, but they also aren't meant to be literally, specifically, or exhaustively true. Challenging the truth of such judgments is likely to be met with "C'mon, you knew what I meant." These are judgments that satisfice. What I'm calling 'strict

judgments' are the opposite: they are meant to be taken literally or specifically. They are attempts at accurately capturing some fact about the world or—apropos to our purposes here—some relation in the world. While loose judgments can take the form of any of the kinds of judgments in my taxonomy, they often stand in for strict judgments of the same kind. When I say 50% of the population are males, I might really mean that "something close to 50% of the population are males." I'm signaling that the relevant strictly true judgment is nearby, but likely not exactly, what the strict reading of my statement claims. Now that this distinction is on the table, I can specify that I'm most interested in characterizing strict judgments, though there may yet be implications for loose judgments.

2. Statistical and Nomic

The class of judgments I'm calling generalizations—judgments about a whole kind of thing or about an individual through time, usually on the basis of evidence about a subset of the kind or temporally extended individual in question—might not be exhausted by the three categories I've introduced. I'm tempted to think that they are, but I'm open to the possibility that there are non-nomic, non-statistical, and non-normality generalizations. Let us assume for the sake of this study that we've got generalizations covered by these three categories. What I hope to argue for, instead, is the *exclusivity* of the three categories. With that in mind, we need to distinguish these three kinds from one another at this general level. Along the way, we must address the Problem of Reduction. There are many kinds of each type of generalization, so the goal now is to find out what makes all of the nomic judgments nomic, the statistical judgments statistical, and the normality judgments "normic."

At first pass, nomic judgments are simply those judgments which are intolerant of exceptions,[1] which means that every apparent exception must be dealt with by proviso or revision, or else the nomic judgment in question must be abandoned. The following are clear cases of nomic judgments in that they are all in some sense universal judgments:

(1) The total entropy in a closed system can *never* decrease.

(2) *Each* prime number is divisible only by 1 and itself.

(3) For *all* massive bodies, $f = ma$.

In each case, a single counterexample[2] would render the claim false. The mark of a nomic judgment is an inferential profile which demands revision in the face of true exceptions. In fact, there is a sense in which *only* nomic judgments have true counterexamples in that part of what it means to be a true counterexample is that the judgment for which it is a counterexample is refuted by it. This, I take it, is in part due to the fact that nomic judgments are perhaps centrally claims about what *cannot* occur rather than centrally predictions about what will occur—though they are that as well. Nomic judgments, finally, are virtually always[3] relative to some frame or domain or set of boundary conditions, and so amount to the claim that *all entities in the domain* show a particular outcome.[4]

---

[1] Woodward (2002, p. 27) "Whatever else a law is, it must at least describe an exceptionless regularity. In particular all laws have the 'All As are Bs' form of (U) universally quantified conditionals in which the condition in the antecedent of the law is "nomically sufficient" for the condition in its consequent."

[2] N.B. that there are real and apparent exceptions to nomic generalizations—any particular case where $f \neq ma$ (indeed, in the real world *f* almost *never* equals *ma* is not a real counterexample to the physical law $f=ma$ because of the way physical laws interact with one another before producing empirical generalizations or predictions. Cf. Smith (2002). A real exception to a physical law, for example, would be an instantiation of an empirical state of affairs ruled out by the law. A body traveling faster than the speed of light would presumably be a real exception to the law that no body can travel faster than the speed of light.

[3] Is the claim "all things are self-identical" relative to a particular domain? It may be that it has a universal domain and so it would be awkward to say that it is true only relative to a particular domain.

[4] Cohen (1999) supposes that lawlike means extensible—"its truth value would remain the same if the number of elements in its domain were greater than it actually is" (223)—and so explicitly quantified statements aren't lawlike. If this is right, and I'm tempted to think that it is, universal quantification will then have two readings: one a lawlike extensible reading and one a statistical reading meaning "100% of the current population…" An evocative example of the statistical reading is "every Supreme Court justice has a prime social security number."

Statistical judgments, again at first pass, are those judgments which make some claim about the prevalence of an outcome in a population, often on the basis of the prevalence of the outcome in some sample or set of samples. The following are clear cases of statistical judgments:

(4) 47% of frogs are males.

(5) About 100 million Americans currently approve of the President's performance.

(6) Any given Ukranian is about 20% likely to be a separatist.[5]

(7) Fatima regularly clocks in about 5 minutes late to work.

These are judgments about the frequency of a particular outcome in a population and so a change in frequency in the population is incompatible with the continued truth of statistical judgments.[6] Statistical judgments are centrally concerned with regularity, frequency, and prevalence. They are therefore essentially linked to or *sensitive* to a particular frequency, regularity, or prevalence—they become falsified if the actual frequency, regularity, or prevalence should change.[7]

3. Normality

There is a different form of generalization humans make, however. Centrally, this kind of judgment involves function ascriptions like "the heart pumps blood," trajectory claims like "walnuts grow into walnut trees," and characterization claims like "dogs have four legs." At the final analysis, the hope is that this kind of generalization will encompass habituals; disposition ascriptions; cause ascriptions;

[5] Note that (6) can only be read clearly as a statistical judgment insofar as it is shorthand for the claim that "20% of Ukrianians are separatists." If it's a probability judgment, then it fast becomes unclear how to classify it. Probability judgments are complex and varied, but I suspect that on final analysis they will be normality judgments; excepting cases like (6) (which stand in for frequency claims).

[6] Schurz (2009b) points out that statistical claims are gradually disconfirmed in the way that they are gradually confirmed and thus exhibit a "symmetry of non-strict confirmation and disconfirmation" whereas strict (nomic) laws are asymmetrical in that a single counterinstance refutes them but positive instances either don't confirm at all or only quite weakly confirm them.

[7] Statistical judgments have their own unique insensitivities as well. Most prominently, they are by their nature insensitive to the non-accidentality of the outcome being predicated of the kind. This is not the sort of insensitivity that interests us in this study.

different forms of *ceteris paribus* claims;[8] reliability claims; competence, capacity, and power ascriptions; intention ascriptions; and perhaps many more. For now, though, the focus will be on these central sorts of normality judgment. That these are all *generalizations* is evidenced by the fact that each makes a claim about either a kind or an individual through time which has implications for (respectively) an individual member of that kind or that individual at any particular time. Smoking, for instance, is something an individual does at a time, but "Alice smokes" can be a true generalization about Alice as a diachronic individual. "*Ceteris paribus* the price will equal the intersection of the supply and demand curves," similarly, is a generalization about all prices, but the property ascribed will be true of a particular price with its attendant supply and demand curves. So these are all *generalizations*. That these are all members of the same family of judgment, however, remains to be seen.

For now, the goal is simply to distinguish this class of generalization from the others and try to alleviate any lingering suspicions that normality generalizations might simply be kinds of nomic generalizations or might be statistical claims of a particular sort. Given our core cases, we can say that all normality judgments are tolerant of apparent exceptions—exceptions to the corresponding universal claim—and are insensitive to diversity in the prevalence of the properties they predicate in the populations of the kinds to which they predicate those properties. So, at this early stage, we can identify two core features of normality judgments:

> (Tolerance) The generalization remains true in spite of some (potentially large) number of exceptions to the corresponding universal generalization.
>
> (Insensitivity) The generalization remains true through (most) changes of the prevalence in the population of the predicated property.

[8] Why not *all* cp claims? Well, certain cp claims (or certain interpretations of the meaning of all cp claims) are specifically meant to have the inferential structure of laws: they're supposed to admit of no true exceptions. Claims that have a proviso which is unspecified and so can be made to rule out *any and all* apparent counterexamples are one such class of cp claims. See Strevens (2012) for a taxonomy.

If doctors should stop performing surgery in such large numbers—perhaps due to advances in non-invasive interventions—it will still remain true for quite some time that *doctors perform surgeries*. After all, Doctors are still the things that perform surgeries when surgeries are performed. This judgment is therefore at least *relatively* insensitive to the prevalence of the practice in the population of doctors. At a certain point in this scenario, the claim, "doctors don't perform surgeries anymore" will become true. Alternatively, if we should come across a white raven, this won't upend our conception of ravens such that "ravens are black" is true. This judgment is tolerant of particular kinds of counterexamples. With just these two features, we can already make clear how these generalizations differ from statistical and nomic generalizations.

First, nomic generalizations are *defined by* their intolerance. At the appearance of even one counterexample, they either need to be revised with a proviso or need to be abandoned. Any apparent counterexample—that is, any exception to the universal judgment that corresponds to the nomic generalization in question—is a true exception and therefore refutes the generalizations. This core feature of their inferential structure thus distinguishes them from the whole class of tolerant generalizations, normality judgments included.[9]

Second, statistical generalizations are intimately tied to specific prevalences in the population over which they generalize. In fact, they are or are perhaps reducible to claims about a specific prevalence in the population, so they are inherently sensitive to changes in that prevalence. Should the prevalence of females among the population of dogs change, then a strict reading of "49% of dogs are female" will be rendered false. Of course, when we say "50% of x's are y" we often don't mean *exactly* 50%, but statistical generalizations are still going to be *relatively* sensitive when compared

[9] Although it actually might turn out that all tolerant judgments are normality judgements. Statistical judgments don't have a clear relationship with the property of tolerance.

with normality generalizations. Again, an essential fact about the inferential structure of statistical claims distinguishes them from the class of insensitive generalizations.

The class of insensitive and tolerant judgments, then, just is the class of non-nomic and non-statistical judgments. This partition thus exhausts the conceptual space, but without being able to decide whether or not there are more than the three categories of nomic, statistical, and normality judgments; we won't be able to determine whether all of the insensitive and tolerant judgments are normality judgments. What we can say at this point is that all normality judgments are insensitive and tolerant. Whether it is *only* normality judgments that are insensitive and tolerant is a question for another day.

3.1 Normativity

We can say a bit more about the inferential profile of normality generalizations. For one, they seem to participate in one set of inferences surrounding explanatory practice. It seems like, if it's true that hearts pump blood, when a heart stops pumping blood or fails to do it well, something will have gone *wrong* for it as a heart or at minimum we'll need to look for the *interfering* factor or *suboptimal* conditions. So, though normality generalizations allow for apparent exceptions, this doesn't appear to be the whole story. At the discovery of an apparent exception, an explanatory task is presented: what went wrong or what interfered? More generally: what's the explanation of why this particular instance didn't present the normal outcome?

This is the *normative* part of the inferential profile of certain normality judgments. A special explanation is required, but by default that explanation will be to identify what has gone wrong or was defective.[10] Minimally, something must have *interfered* in the normal goings on and interference is itself importantly normative in that there is a standard process against which the interfered process is judged.

[10] Strößner (2015, p. 795 n2) and Nickel (2014), for instance, explicitly reject normativity.

To get onto the type of normativity in question here, we have to do a bit of work. The kind of normativity often discussed in contemporary analytic philosophy—in particular philosophy of action, epistemology, and metaethics—is one that is intimately tied to rational agents and the reasons they have for acting or updating beliefs. If it doesn't make direct reference to reasons for a rational agent, many think, then it can't be normative. One can understand the motivation here, but the dualism inherent in the intuition is deeply problematic. Obviously, reasons for rational agents are a central case of normativity, and it's perhaps true that the type of normativity that adheres to reasons for rational agents is a paradigm for other forms of normativity, but this is simply not the only form of normativity. There are norms for things that have functions, for instance, quite apart from what reasons the functions of things provide rational agents. More controversially, a thing, insofar as it is a participant in a particular process or system, *ought* to have all of the features constitutive of the role it is playing in that process or system. If it doesn't, then it has—as a conceptual matter—failed or has suffered defects *qua* participant in that role. If I want to use a fork to eat my spaghetti, but the fork I reach for is made of paper and can't be used for the task, then *qua* eating utensil, the paper fork has failed or is at minimum not apt or well-suited to the role. Of course, I *qua agent* made a mistake and I have therefore failed in my own way, but again it is wrong to say that it is *exclusively* the agent that has failed. The paper fork wasn't a good eating utensil and so *failed* to be an eating utensil as I intended it to be.[11]

With this in mind, we can see how this role-dependent normativity might attach to the inferential structure of normality claims. It appears that, as far as we can tell at this point in the analysis, each normality judgment will carry with it the inferential disposition such that at the

[11] It's tempting to think that *failure* attaches exclusively to *trying*. Under that analysis, it feels odd to say that a paper fork, which wasn't *trying* to be a real fork, nevertheless has *failed*. I think this analysis is problematic, particularly for the reason that things with functions seem clearly to be able to fail—by my lights at least—and yet aren't clearly *trying* in the full sense of the word. An engine, bolt, weld, and heart can all fail. Were they trying? It's not so clear and the mere fact that it isn't clear speaks against a tight conceptual relationship between failing and trying.

appearance of an apparent exception, the inference that something has gone *wrong* is warranted. Let's run through some clear examples, and then look at a few more difficult examples.

In our well-worn example "dogs have four legs," it seems fairly clear that in the case of any dog without four legs, something will have gone *wrong* somewhere. Either the dog is deprived of having four legs by some environmental event, or the dog suffers a defect by having a genetic mutation, etc. Because of the role that four-leggedness plays in dog life, lacking four legs is a bad thing for a dog relative to the ends of the canine form of life.[12] Similarly, for "nail guns shoot nails when the trigger is pulled," which is an artifactual claim, we will think that something has gone *contrary to design* or *contrary to intention* if a nail doesn't shoot at the pulling of the trigger. It will be contrary to design if the gun has broken down, and it will be contrary to intention if the nails run out without my knowing. *I designed it to work this way, but it is not. I expected and therefore intended it to shoot a nail just then, but nothing happened.* I take it that these are very clear examples of the normative inferential structure of normality claims.

Here are some less clearly amenable examples:

(8) Glass breaks when struck

(9) Alice smokes on her lunch break

It's harder to make the case that something will have gone *wrong* in any case where there's an apparent exception to these generalizations since they both on their face are dispositions or habits to do something with a negative valence. To get onto the right intuition here, since I'm contextualizing my normative claims, we must situate these examples in contexts where the relevant outcomes are

[12] Let us presume that there is some, perhaps large contingent of professional biologists who take issue with these claims. Does that matter for our purposes here? What if, even stronger, biology as a practice makes no room for such judgments owing to its methodological constraints? To suppose that this makes a difference to whether or not we should so describe a species, it seems to me, is to suppose that the world is such that there is only one or very few true descriptions of a thing. A species is a concept useful to biologists of every stripe, but also to philosophers, ethicists, and persons in general. The methodological monism of biology, it seems to me, is no mark against a pluralism that invites many theoretically fruitful characterizations of a thing.

good relative to those contexts. Regarding (8), if we're shooting a chase scene in a film and we want our hero to jump through a glass window onto the street below, it will indeed be *inconvenient* for the glass to fail to break. We'll need an explanation so that we can remedy the situation, but we'll also be tempted to think that something will have gone wrong insofar as the glass was a prop in our movie. For (9), if I'm an assassin hired to off Alice (philosophers love assassins) and I'm relying on her smoking habit to bring her outside and into my line of fire, then her deciding to skip that day or begin a nicotine patch treatment will be inconvenient indeed. Something has gone wrong in my plan—Alice isn't doing as she's supposed to with respect to my plan. So, we can see how the relativization of the normativity in question can get us out of some initially counterintuitive examples. The hope is that a similar analysis can be given to any similarly counterintuitive example.[13]

The normative inferential structure—conceived of as dealing in context-relative normativity—won't extend to statistical and nomic generalizations. A statistical claim will often have the inference to an individual probability claim available. For example, to say that roughly 80% of animals are non-domesticated allows for the inference that any given animal has a roughly 80% probability of being wild. If an individual animal is domesticated, though, nothing will have clearly gone *wrong*. A nomic claim is always responsible to the class of actual individuals over which it generalizes. An exception to a nomic judgment, therefore, will require a change in the law being posited, perhaps by the addition of a proviso or disjunction. It won't allow us to claim that something has gone wrong in the exceptional case.

[13] One might wonder about exceptionally *good* specimens. What about the Flash, who can move abnormally fast for a human? What about less fantastical cases? It seems to me that in many of these cases, something will have gone *wrong*. Superheroes famously often meet some otherwise ill fate that just so happens to grant them superpowers. Michael Phelps, though, arguably has abnormally proportioned limbs and other organs, which make him an incredibly fast swimmer. Is he merely a rare specimen or an abnormal specimen? It isn't at all clear what should be said and so one can at least set aside the case and those of other elite athletes as potential counterexamples. More work needs to be done on the abnormally good.

3.2 Do Normality Judgments Reduce to Nomic Judgments?

Are normality claims constrained or loosened forms of nomic or statistical generalizations? Do they reduce to one or the other of the forms in question? Consider the most straightforward way onto the claim that normality judgments are distinct from both statistical and nomic judgments. I'll take my examples from Prasada (2000):

(10) Dogs are four-legged.

(11) All dogs are four-legged.

(12) Fido has three legs.

(13) X% of dogs are four-legged.

Note that the generic claim (10) does not entail the nomic or law-like generalization (11). That dogs are four-legged in no straightforward way entails any claim about all dogs. Generic claims famously "tolerate exceptions" and nomic or quantified claims do not. There is no inconsistency between (10) and (12), which is perhaps the most interesting feature of generic claims. There is, however, straightforward logical inconsistency between (11) and (12) given the auxiliary claim that Fido is a dog. Therefore (10) does not entail (11).

It's also apparently true that (10) does not entail any specification of (13). Nor does it seem to entail any weaker prevalence claim involving "most" or "a great deal" or "many." Many have attempted to find such a link for generics or certain classes of generics, and it seems that all have failed in one respect or another.[14] I take it that similar approaches to normality claims will fail for similar reasons. There do not appear to be any principles governing what sorts of percentages or frequencies of an outcome in a population is consistent with a given normality judgment. Very few dogs could have four legs and yet four-leggedness would still be normal for dogs. Taken historically, we might think "dogs have four legs" would be false if most dogs throughout the history of the

[14] Cohen (1999) is probably the most promising approach available in the literature.

species had missing legs. Yet we can posit worlds in which dogs are subject *en masse* to a disease which causes limb loss. In such a world, I can't imagine it would be false to say, "dogs have four legs." Of course, there might be contexts in which it will be right to say, "dogs don't have four legs for long" or, "dogs have 3 legs," but one can imagine responding "no, dogs have four legs, and then the disease removes one of them." We need not travel to other worlds, though, to find now well-known cases of dramatically low prevalence still being consistent with the truth of a generic judgment. For example, sea turtles have long lifespans, even though only around 10% of them survive into adulthood. It appears that normality judgments don't entail, nor are they reducible to statistical or frequency claims. I don't take myself to have shown this at this point, but few seem to dispute the claim anymore, so I won't spend more time on the matter.

For normality judgments to reduce to one or the other form of judgment would be for it to be possible to capture the content and inferential structure of such judgments by appeal to resources available to the reduction base. For instance, normality judgments might turn out to be reducible to nomic judgments in that nomic judgments remain nomic even with radically constrained domains of application. It might be the case that "dogs have four legs" is reducible to "all dogs that have suffered neither loss of a limb nor genetic mutation or developmental deprivation affecting limb count will have four legs." This is a tempting thought. If it turned out that we could capture the content and inferential structure of normality judgments by simply restricting the domains of nomic judgments, then we'd have no need for a separate account of normality judgments.

One way this strategy shows up in the literature is in discussions of circumscription as a semantic tool.[15] Using a proviso like "unless abnormal" allows one to make a nomic statement out of a normality claim. Alternatively, there is a camp we might call the "explicit proviso" camp in the literature on ceteris paribus claims. They claim, as I mentioned in the previous paragraph, that the

[15] Cf. McCarthy (1980, 1986), Bastiaanse and Veltman (2016).

proviso summarized by *ceteris paribus* is an explicit set of exceptions that get disjoined in an "unless" clause. Something like: glass objects break when struck unless they are tempered or not struck very hard or padded or have dampening mechanisms on them or… Either strategy would, it seems, effectively turn a normality claim into a nomic claim by rendering it intolerant of counterexamples.

It also seems that we could potentially capture the normative inferential structure because the proviso specifying the domain of the judgment includes explicit explanations for any apparent counterexample, or includes a blanket classification of any apparent counterexample as abnormal. The normativity that we find in certain normality judgments would, on this strategy, be contained in the proviso.

The core problem with such a strategy is that it simply can't capture the intuitive content of most normality judgments. When I say, "nail guns shoot nails," I don't intend to be making a quantificational generalization about all of the individual (actual and possible, perhaps) nail guns. Instead, I'm attempting to characterize the kind of thing quite apart from what individual nail guns do or don't do. Conversely, it seems that the most natural reading of the main clause of "nail guns shoot nails, unless they're broken" is "any given nail gun will shoot nails." By including the proviso, the content of the claim shifts to the quantificational or individual level, but this is not the content of a normality judgment.

Boutilier (1994, p. 90) offers an intriguing, though perhaps not entirely convincing critique of circumscription and other explicit proviso interpretations of normality judgments (though normality judgments aren't his true target). Antecedent strengthening, where one adds conjuncts to the antecedent of a conditional, is a valid inference pattern. Given that it is true that if I go swimming, I will become wet. It is also true that if I go swimming and kayaking, I will become wet. It's even true that if I go swimming and then dry off, I will still have become wet at some point in the process. The problem, Boutilier thinks, is that we can strengthen the antecedent with anything at all

according to the logic of circumscription, which is built on classical logic. So we can change "if it's a bird, then it flies" to "If it's a dead bird, then it flies." This inference is clearly not truth preservative and there's nothing in the logic which prevents it. The logic of circumscription, therefore, allows an invalid inference pattern. Though this counterexample isn't terribly convincing, perhaps similar inference patterns concluding in "a wet match will light when struck" or "a sick cheetah can outrun a human" make clear that this is a general problem for the logic of circumscription. Conclusion: circumscription does not appear to be a workable strategy for reducing normality judgments to nomic judgments.

Instead of a simple "not abnormal" circumscription, why not include an explicit proviso delineating all of the conditions that must hold for the normal outcome to be guaranteed? This would make what appears to be an "unless something interferes" normality judgment into a universal or nomic judgment. Here, I'll simply advert to the intuitive sense that many have reported[16] that this doesn't seem possible in the slightest. There will always be some new condition that hasn't been included in one's proviso and so one cannot—so the argument goes—create a true universal judgment out of a true normality judgment.

Another way that normality judgments might reduce to nomic claims is by being loose versions thereof. It might be that when you say, "zebras have stripes," you express a nomic judgment like "all zebras have stripes" but with the understanding that I don't mean this strictly to be true, it's just a useful heuristic for forming expectations.

To accept such an account, we have quite a few pills to swallow. For one, most normality claims seem on their face to be uncontroversially true and strictly so. "Glass shatters when struck" doesn't seem to be tolerant of counterexamples because it is a loose generalization, but because each

[16] A few have made the claim that explicitly listing defeators for a defeasible proposition like a normality judgment or ceteris paribus law is unrealistic if not in principle impossible. Cf. Hempel (1988), Rescher (1994, p. 14), Pietroski and Rey (1995), and Strevens (2012).

counterexample can be explained in a satisfying way. Similarly for "dogs have four legs" and "nail guns shoot nails": these aren't nomic claims that aren't meant to be literally true. They are instead (people widely accept) literally true claims that are simply tolerant of apparent counterexamples.

Finally, we should consider the possibility that normality generalizations are universal generalizations (over actual and possible cases) over domains restricted by context. Virtually every universal claim has a restricted domain. Why not simply think that normality judgments are judgments with especially restrictive domains?[17] For instance, why not think that "dogs have four legs" isn't a claim to the effect of "all dogs that haven't been amputated or dismembered or had a necrotic disease or a genetic disorder causing them to have fewer or more than four legs have four legs"?

One principle shortcoming of such an account is that generics are (now famously) insensitive to context. If I am at the zoo near the lion cage[18] and I lean down to my daughter and say "all the lions are hungry", I obviously don't mean all of the lions in the world, and I certainly don't mean all actual and possible lions. Instead, I simply mean "each lion in this cage is hungry". If, alternatively, I say "lions are hungry", I've said something the content of which is simply false: lions aren't a hungry kind of thing, it's not clearly normal for lions to be hungry, etc. The generic sentence doesn't appear to be sensitive to context in the way that the explicitly quantified sentence is.

Furthermore, as Nickel (2010a) notes, restricted quantification accounts cannot handle conjunctive strengthening, which appears to be endemic to normality judgments. Conjunctive strengthening is an inference pattern in which a disjunctive generic sentence is given a conjunctive semantic interpretation. For example

(14) Elephants live in Africa or Asia

[17] A few have discussed such a strategy in the semantics for generics: Schubert and Pelletier (1987), Declerck (1991), and Chierchia (1995).

[18] Similar examples are discussed in Greenberg (2007, 2012).

Implies both of the following;

(15) Elephants live in Africa

(16) Elephants live in Asia

Whereas it might be natural to treat a disjunction as a disjunction, there are countless cases—both generic sentences and others—where an apparent disjunction is in fact semantically conjunctive.

According to Nickell, an account of generic sentences that treats them as domain-restricted universally quantified statements cannot account for conjunctive strengthening and so appears to fail on a ubiquitous phenomenon of generics. The force of his argument comes from a study of the inferential structure of different grammatical forms. Sentences like "most of the passengers on the train had trouble breathing or felt nauseous" fail to support the inference to claims like "most of the passengers on the train had trouble breathing and felt nauseous." Domain-restricted universal statements do not have the same inferential structure as the sort of sentences for which they were supposed to serve as the semantic structure. It follows that generic sentences—which almost always express normality judgments—do not have the semantics of domain-restricted universal quantifications.

3.3 Do Normality Judgments Reduce to Statistical Judgments?

According to Schurz (2009b), Hempel (1965 ch. 12.3) was the first to suggest that normality statements could be replaced by "numerical-statistical laws."[19] Cohen's (1999) account is perhaps the most compelling version of an attempt to reduce generic predications—and therefore at least potentially all normality judgments—to claims of .51 or greater probability given some important constraints. I'll therefore focus on Cohen's account as the exemplar of an attempt to reduce normality judgments to special cases of statistical judgments.

[19] Schurz offers three challenges to this view in (Schurz 2009b).

According to Cohen, a generic is true if and only if the probability of any member of its domain having the predicated property is greater than .5 given that we calculate that probability by reference to a homogeneous reference class. For a reference class *R* to be homogeneous with respect to partition *x* and property *p* is for the cells produced by *x* to have no different probability of *p* than *R*. Cohen requires that the partitions be effected by means of all of the salient partitions for the speaker of the generic sentence being evaluated. Put in more ordinary language, dividing up the domain in different salient ways shouldn't produce subsets with different probabilities of having the relevant property than the probability attaching to the domain as a whole. If they do, then we haven't achieved homogeneity and we must partition until we have before we calculate the probability in question.

For instance, dogs have four legs if and only if the class of all dogs is homogeneous with respect to four leggedness when partitioned in various salient ways. If we divide dogs up into terriers, hounds, retrievers, etc. no group is more or less likely to have four legs. If we again divide dogs into wild dogs and domestic dogs, we've still got more or less the same probability (though perhaps it's slightly skewed in this case). Again, divide dogs by continent of residence, and you'll still find roughly the same probability. Furthermore, that probability will be greater than .5 and so the generic "dogs have four legs" is true.[20]

There is much more to say about the sophistication of Cohen's account, but this will suffice for our purposes. If Cohen's account is successful, then we have the beginnings of a potentially successful account of normality judgments in terms of a frequentist notion like statistical probability.

[20] An interesting case reveals the built-in relativity in Cohen's account which is necessary to account for divergences in judgments of the truth of generics between people. People disagree about the truth of "birds fly" because, according to Cohen, different partitions are salient to different people. A lot of people will think that species are a salient partition and so will treat the generic as false (penguins have a probability 0 of flying), but others might think that location, size, color, etc. are salient and so will treat the generic as true.

This would mean we could reduce normality judgments to statistical judgments, potentially without remainder.

According to Nickel's critique of Cohen's account, the account "systematically makes the wrong predictions in cases in which there is more than one way of being normal in a given respect."[21] The example Nickel focuses on is the case of the dual generics:

(17) (15) Elephants live in Africa.

(18) (16) Elephants live in Asia.

As Nickel demonstrates, Cohen's truth conditions for generics will have the undesired implication that some individual elephants live on both continents, perhaps by migrating. If his analysis is correct, then Cohen's account has a serious problem when it comes to cases where there are more than one ways of being normal.

Cohen's appears to be the most sophisticated attempt at a reduction to statistics available in the literature, and it doesn't appear to be able to capture the content of an important class of normality judgments.[22] We might, on the assumption that there isn't likely to be a better account, conclude that normality judgments simply cannot be reduced to statistical judgments. I think this is a reasonable inference, but I'll offer a few more arguments against the possibility of reducing normality judgments to statistical judgments in an attempt to more tightly seal the coffin.

First, it is difficult to imagine how one might capture the normativity of normality judgments in a statistical judgment. Statistical judgments are, after all, paradigmatically true or false independently of the accidentality or non-accidentality of the outcome being predicated. It's not impossible that

[21] Nickel (2016, p. 97).

[22] Cohen's account is about probability rather than statistics, which are importantly different. Upon reflection, it seems that most probability claims are themselves a species of normality judgment rather than a species of statistical judgment. A frequency interpretation of probability treats them more like statistical claims; but on most accounts of probability, they end up having the inferential profile of a normality judgment rather than a statistical judgment. Nevertheless, Cohen still offers the most sophisticated attempt to do anything in the neighborhood of reducing normality judgments to statistical judgments.

there might be tools up the sleeve of the champion of the statistical judgment that might allow them to deal with the problem of at minimum tracking non-accidentality, but we have good reason to be pessimistic that it can be done, simply given the track record of such efforts.[23] Beyond tracking non-accidentality by statistical means, though, is the perhaps deeper problem of capturing any sense in which something has gone wrong when it is rare—even if one utilizes a specialized conception of the "rare".

On a related note, it's unclear how one might fashion a statistical claim to cover cases where the actual frequency and even the conditional probability of a normal outcome is extremely low in spite of the normality of that outcome. As Cartwright notes (1999, p. 45), countless uncontroversially true *ceteris paribus* generalizations hold only in ideal conditions, which are quite rare, and so one is blocked from restricting the domain of a statistical generalization to cases where conditions are normal in the statistical sense. Furthermore, certain cases such as:

(19) Mary handles the mail from Antarctica

Plausibly have no actual instances even though the *conditional probability* might be quite high. The conceptual resources that probability makes available are quite powerful, but they are not at all clearly made available by statistical prevalence.

Before moving on, it is important to distinguish between two claims:

(Statistical Consequence) Normality judgments entail or at minimum informally license the inference to corresponding statistical judgments.

(Statistical Reduction) Normality judgments can be reduced to statistical judgments.

Schurz (2009b) and Strößner (2015, p. 801) defend Statistical Consequence, at least when it comes to certain normality judgments. They think this entailment is necessary if normality judgments are to

[23] Most notable, is one of Cohen's chief influences, Wesley Salmon (1971, 1977), who eventually found that he needed non-statistical conceptual resources to track explanatory and causal—i.e. non-accidental—relations.

have empirical content.[24] I'm not so sure this is correct, but I'll pick up the discussion somewhat in the final chapter. No one, according to a no doubt incomplete survey of the literature, seems to defend Statistical Reduction. Cohen is the closest to this position, though, again, his appeal is to probability rather than statistical frequency.

4. The Problem of Pluralism

Given this assortment of generalizations, it might seem difficult to get a handle on what they have in common as a group, and it might be tempting to think that there is no such thing—that non-nomic and non-statistical generalizations are a grab bag of disparate judgments with no unity to be found among them. First, let's quickly sketch some key differences in inferential profile, and then I'll make a *prima facie* case that these are a proper family of generalizations rather than a mere grab bag.

The core distinction among normality judgments appears to be the distinction between *Complete* and *Incomplete* normality judgments. Complete normality judgments track normal outcomes which will obtain unless there is interference or defect or the like. Incomplete normality judgments track normal outcomes that are merely one of a set of contrary possibilities. A computer will turn on when its power button is pressed. If the power cord is unplugged or if the internal power supply has a short circuit, then it will not. That is, if there is interference or something goes wrong for it *qua* computer, it might not turn on. A cat, though, might be black, grey, white, orange tabby, or many other colors without anything being *abnormal*.

A related distinction, which we'll discuss in a bit more detail in the final chapter, is that between what one might call *Exclusive* vs *Inclusive* normality judgments. An exclusive normality judgment is one that predicates an outcome of a kind where there is no other normal outcome for that kind.

[24] Strößner (2015, p. 801): "a lack of testability is by any means undesirable. It means that you can believe a normality statement and no statistical evidence whatsoever can lower this degree of believe [sic]."

Though it is empirically suspect, it is often thought that 'Ravens are black' is such a judgment.[25] An inclusive normality judgment is one that attributes an outcome to a kind without there being any implication that this outcome is exclusively normal. 'Dogs have fur' is a true normality judgment in spite of the fact that 'dogs sometimes have hair' is also true. The distinction between exclusive and inclusive normality judgments seems to track the distinction between complete and incomplete normality judgments quite closely. Nevertheless, these distinctions are not identical. Where a complete normality judgment covers what will happen if we eliminate all relevant abnormalities, an exclusive normality judgment need not. It is normal for a President in the United States to have won the popular election, but as we have seen, they need not in order to become president. It seems that there might only be one normal way to be elected President even though it won't necessarily happen absent abnormal interference. A complex and shaky example, to be sure, but it does illustrate the idea. As I'll note in the final chapter, the distinction between exclusive and inclusive normality judgments doesn't seem to track something that is going on in the inferential structure of the judgments themselves, but instead seems to track cases where our background beliefs shape the way we think about the inferential structure.

Grammatical differences seem to loosely track differences in inferential profile.[26] For instance, to say "it is normal for piano players to be able to play *Für Elise*" seems importantly different from "piano players can normally play *Für Elise*." The latter seems to bring with it the implication that it is only an abnormal piano player who cannot play *Für Elise*. The former, alternatively, doesn't seem to carry such an implication. There are at least four basic grammatical forms used to express normality judgments:

(A) It is normal for *X*'s to *φ*.

[25] Nickel (2014). Leucistic ravens are an important and compelling challenge to such a claim. This is discussed in more detail in chapter 6.

[26] Strößner places a lot of weight on grammatical categories. I am hesitant to do the same because normality judgments as I have defined them are individuated in terms of content rather than along grammatical lines.

(B) Normal *X*'s *φ*.
(C) *X*'s normally *φ*.
(D) *X*'s *φ*.

Especially D, but the others as well, are meant to be stand-ins for grammatical categories. Generic sentences, for instance, have at least three grammatical forms: bare plural, indefinite, and definite.[27] D is meant to stand in for all generic grammatical forms. These four basic forms (A-D) will likely elicit different intuitions about the inferential profile of example sentences and so we mustn't suppose that all normality judgments will have the *same* inferential profile.

At long last, what reason have we for thinking that there isn't a mere grab bag of judgments at play here? To motivate the pluralist intuition, here is a sampling of various statements that I take to express normality judgments:

> Wide receivers are normally smaller and faster than linebackers.
>
> Computers turn on when you push their power buttons (provided they are off when you push the button).
>
> Springs return to more or less one length when they are released after being stretched or compressed.
>
> A NAND transistor receiving only one positive charge outputs a positive charge.
>
> Jamal takes his smoking break at around 10:15am.
>
> Lions have manes.
>
> Ducks lay eggs.
>
> Safira plans to attend Dartmouth.
>
> Francis can hit a bullseye in archery from 50m away.
>
> An eagle can see the movement of a small rodent from 100 meters up.

[27] Bare plural sentences are sentences of the same form as "dogs have four legs." Indefinites have the same form as "a dog has four legs." Definites have the same form as "the dog has four legs."

Represented here are statements comparing state properties, dispositionals, functionals, intention ascriptions, habituals, characteristic claims, capacity claims, and competence or mastery claims. Why think they have anything in common? They seem to be *prima facie* disconnected sorts of claims.

Nevertheless, we might think that an account can be given of what one sort of connection these kinds of judgments are tracking in the world. To get onto the intuition in play, consider the following set of scenarios:

(20) You consider going to the ballet dressed in your workout clothes, but decide against it because that's not what one does, or because it's not the done thing or because it would be disrespectful or rude.

(21) A frog develops only its hind legs in its maturation from a tadpole into an adolescent frog and is eaten before it lives a long frog life.

(22) You turn on the light switch and the light does not turn on.

(23) When choosing a table, you reject a candidate because it is too wobbly.

(24) You decide to forgo your usual morning cup of coffee or tea.

(25) Mutated ants are born which emit a foul odor unpleasant to predators, they survive to reproduce in great numbers while their kin slowly die off.

(26) We'd expect orchid pollen consumers, given the structure of the orchid flower and the rigidity of its stem, to be either heavier hoverers (like hummingbirds) or light flyers (like butterflies). Lo and behold, they are either heavier hoverers or light flyers.

In each case, there is some constraint on things that isn't always satisfied by means of a mechanical cause or set of mechanical causes, but which nevertheless plays a role in the explanation of the outcome. The frog is eaten, but it might just as well have survived. We'd simply be a bit surprised that it did so well. "Frogs have four legs" is a true generic judgment about frogs because of the relation frogs have to the constraints in their environment. Their ability to maneuver quickly is very important to their ability to avoid predation. If a frog without such ability survives, we might think it serendipitous, unexpected, or even accidental. The constraint is real whether or not it is realized in a mechanical cause (i.e. the gnashing of teeth) in every instance. Similarly, one might never catch a wry

glance or a side eye if one wears their workout clothes to the ballet because of the gentility of the other guests, but one would know what the norms are and one would nevertheless feel out of place.

In every case, there is what I call a "Principled Connection" (Prasada 2000, 2012; Prasada et al. 2012; Prasada and Dillingham 2006, 2009) between the subject of the corresponding normality judgment and its predicate. A principled connection is opposed to a nomic connection or a statistical/probabilistic connection, each of which might obtain between a member of a kind and some feature or activity judged to belong to members of that kind in some way. A principled connection is one that attaches to top-down constraints which may or may not be realized in mechanical causes in every case. These are outcomes that "in principle" should obtain even though they might not "in practice" obtain

The basic idea is that there are some types of explanations that center around what would be a means to some end or what "should" obtain given the kind by which an individual is specified. Dispositions are an example of the *explananda* of such explanations in that dispositions are posited to explain reliable or regular occurrences—they tell us what should in principle be the case. If we're trying to explain why a light switch's motion is connected with a light's illumination or ceasing of illumination, then positing a disposition of the light switch to turn on the light would explain this further occurrence. This sort of functional, top-down predication of a kind is part and parcel of the normality judgment, which seems to track the normal or in principle what "should" happen.

Being a member of a kind (in the fullest sense of "kind"), as we'll explore in Chapter 2, is similar to (or perhaps identical with) playing a role in a system or organization. It is therefore to be subject to a particular set of constraints, norms, standards, and goals; and each of these licenses inferences surrounding the individual in question. Most prominently, there are a number of inferences involving what "should"—in some sense of 'should'—be the case for that individual. Shortstops should have fast reflexes, pitchers should have a good fast ball and good aim, ballerinas should have

strong feet and exceptional flexibility and core strength, frogs should have front legs, opera attendees should be dressed to the nines, and so on.

In short, there's not only a common inferential profile to the normality judgments we've explored (albeit a potentially quite limited profile across all normality judgments), but there's also a case to be made that normality judgments all track similar sorts of phenomena in the world: the principled connection between being a kind of thing or admitting of some description and exemplifying certain features. Further proof that we are dealing with a distinct sort of judgment will come in Chapter 3, where I will posit an account of what *content* they have in common.

5. Tools of Analysis

Before closing, we should address certain sorts of judgment about which much ink has been spilled. Why not simply extend a successful account of generics or *ceteris paribus* laws to other forms of normality judgment?

In order to understand normality claims, we cannot solely rely on the lexical semantics of "normal" and "normally." There are two distinct senses of the word and there don't appear to be any syntactic clues as to which use is involved in any given case. It would be convenient if we could simply do lexical semantics, since it is at least a clear path forward. But there doesn't appear to be any set of principles available which allow us to distinguish the group of good linguistic data from misleading. Absent such principles, we ought to be careful not to fall into a simple semantic treatment of a word, instead elucidating a *concept* by appeal to the ways it shows up in various judgments.

5.1 Generic Sentences

Since so much work has gone into understanding the truth conditions for generics, why not simply think that normality claims *are* generic claims? It's a tempting thought: we could explain

sentences like "she normally smokes" and "normal male lions have manes" in terms of the generic "she smokes" and "male lions have manes." We could similarly treat all generic sentences as one form of normality judgment or another. I'm not tempted to go this route.

Leslie (2008, 2012b) offers the following counterexamples, inferring that therefore the semantics of generics cannot simply appeal to normality, but must be more sophisticated.[28]

(27) Mosquitoes carry the West Nile Virus.

(28) Sharks attack swimmers.

Both of these appear to be true generics. For Leslie, they are true because they codify striking and ecologically important information for human psychology. It furthermore seems that each is not a normality claim. It seems wrong to say that it is normal when a mosquito carries West Nile Virus, or that normal mosquitoes do so, or that mosquitoes normally do so. It also seems wrong to claim that sharks normally eat people (indeed, they normally eat sea lions and fish), or that normal sharks eat people (perhaps it is only *abnormal* sharks who are convinced that people are sea lions or fish that eat them). If this analysis is correct, then the semantics of generics is not the semantics of normality claims, even though in large part generics might encode normality claims. (27), for instance, seems to be a perfectly respectable normality claim and also a focal example of genericity. Perhaps, therefore, it's the case that though the linguistic phenomenon fails to perfectly represent normality claims, they are nevertheless closely allied forms of judgment. I'll adopt the view for the sake of this discussion that *certain* generics express normality claims and furthermore that the most straightforward way to express most normality claims is by appeal to generics.

[28] Note that a semantics of generics that appeals to normality (like Nickel's (2008, 2014, 2016)) rules out the *equation* of generic claims with normality claims. So Leslie's argument, aimed as it is against accounts of the semantics of generics which appeal to normality, doesn't clearly translate into our current purposes, but the counterexamples are nevertheless troubling.

I go with the latter claim—that a generic form is often preferable to one that includes the word 'normal' or 'normally'—because of sentences like the following:

(29) Politicians are normally lawyers, though other professions are increasingly being represented.

(30) Flowers normally smell nice, but this one smells like rotting flesh.

(29) and (30) seem plausible enough, but neither seems to be a normality claim nor a generic, especially given the attached clauses. They seem to be claims about most or a great deal of politicians and flowers rather than about the normality proper of the features in question. If this analysis is correct, though, then certain claims which appear on the surface to be normality claims are in fact claims about regularity or perhaps some confusion of the two. Generic sentences are less prone to produce the same result—a surface grammar that appears to be about normality but instead seems to involve a statistical notion instead. Consider the fact that "politicians are lawyers" is simply false even though there is a statistical connection between the two.

I go with the former claim—the claim that certain generics just are or express normality claims—because I'm convinced by the work of Gehrman (ms), Thompson (2008), Lawrence (2006), and others in the neo-Aristotelian tradition that a group of generics we might call "characterizing generics" express the kind of normality judgments I'm centrally interested in. If it can be expressed in the form of a true generic, then heuristically it's a true normality judgment; and if it can't be expressed in the form of a true generic, then heuristically it isn't a true normality judgment. Nevertheless, we shouldn't necessarily expect a semantics of generics to give us an account of normality because generics are a set of syntactic categories and therefore aren't obviously supposed to be unified in terms of semantic content (indeed some accounts of generics simply appeal to normality as a semantic device and so trivially cannot offer us an account of normality).[29]

[29] Nickel (2008, 2016) is a notable exception in that he appeals to a notion of non-statistical normality in giving a semantics for generics, but then goes on to explicate that notion. I'll engage with his work in chapter 3.

5.2 Ceteris paribus

Another class of judgment that has received much attention is the *ceteris paribus* judgment (CP). If normality judgments were all simply *ceteris paribus* claims, then we could turn to the CP literature in search of a successful account of normality judgments. Unfortunately, the relation between CP judgments and normality judgments is perhaps even more tenuous than the connection with generics.[30]

First, CP judgments all express *complete* normality judgments if normality judgments at all, and normality judgments include both complete and incomplete normality judgments. It simply isn't true that '*ceteris paribus* dogs are black.' Second, as Sheldon Smith notes, *ceteris paribus* claims "invite consideration of what needs to be equal" (*Personal Correspondence*), whereas countless normality judgements don't seem to do the same. Faced with the claim "*ceteris paribus* the price of a good equals the intersection of the supply and demand curves for that good," one immediately wants to know what it means for other things to be equal. It seems to be wrapped into the "*ceteris paribus*" clause that *specific* parameters or factors are equal.[31]Alternatively, faced with the claim 'dogs have four legs,' it's not clearly part of the judgment that certain specific conditions will hold. Moreover, Smith notes that there is an important linguistic difference between the following, to use Nickel's oft-repeated example:

(31) (16) Elephants live in Asia

(32) Ceteris paribus, elephants live in Asia.

Whereas (16) appears to be a statement about the *kind* elephant, (32) appears to tell us what to expect of *each individual elephant*. CP clauses, that is, tend to shift away from the kind level to the

[30] Furthermore, the success of various accounts of CP statements is in dispute. An even deeper worry is that, as Schiffer (Schiffer 1991, p. 10) and Woodward (2002, p. 305), point out, actual CP statements aren't widespread in sciences outside of economics. Menno Rol, (2012) has a wider discussion of the role of CP statements in economics.

[31] Of course, what exactly "equal" is, is a philosophical challenge in and of itself. Cf. (Schurz 2009a; S. Smith 2002)

individual level. As another example, Smith discusses the fact that "smoking causes lung cancer" is importantly distinct from the claim that "*Ceteris paribus*, smoking causes lung cancer." The former is the claim that there's a causal link between two kinds and the latter is a claim about each individual case of smoking or at least each individual smoker: they will get cancer barring shifts in certain parameters or factors.

Strößner (2015) resists treatment of normality judgments in terms of CP judgments as well. She notes that some CP statements are true only in rare conditions and so they can't facilitate a default inference of the form "If A, then by default B."[32] For instance, *f* essentially *never* equals *ma* in the empirical world because there are always forces other than gravity acting on bodies. Complete normality judgments, though, it seems can facilitate these default inferences. A stag has horns, so I can default to "it has horns" if I am told "it is a stag."

6. Conclusion

Normality Judgments are a distinct form of generalization from both Statistical and Nomic Judgments. While this tripartite distinction may not be exhaustive of the conceptual space, I hope to have shown that it is indeed—or at any rate is very likely—an *exclusive* divide. Both in terms of content and perhaps more strikingly in terms of inferential structure, normality judgments are fundamentally different from nomic or universal judgments on the one hand and statistical or frequency judgments on the other. Whether or not they are a deeply unified set or a true family of generalizations remains to be seen. That will be the project for the next three chapters.

[32] Strößner, (2015, p. 797), citing Cartwright (1999, p. 45) as an influence in her thinking.

Chapter Two: Kind Explanation

There are many different kinds of explanations. A 'why' question can elicit—in different contexts—a mechanical explanation, an ontogenetic or developmental explanation, a functional, adaptationist, or selectionist explanation, a lineage explanation, an equilibrium explanation, an intentional explanation, a necessary explanation, an accidental explanation (if there is such a thing),[1] and so on. One type of explanation, though, might appear on its face to be particularly mysterious or even dubious—let's call it *Kind Explanation*. Kind explanation is an explanation which simply cites a kind as its *explaniens*—one that simply specifies or characterizes the individual, process, or system involved so as to *ipso facto* provide an explanation of the feature, change, relation, activity, condition, etc. (henceforth simply 'feature') in question—i.e. the *explanandum*. A paradigm example is found in Aristotle: being a carpenter explains one's building a house in a way that being a musician doesn't.[2] As I interpret Aristotle, this is the reason we would claim that the specification of the subject in "the musician built a house on the outskirts of town" is incidental, whereas "the carpenter built a house on the outskirts of town" is principled and explanatory—*per se*. The subject's *being a certain kind of thing* can explain why the subject might *do* certain activities or *exhibit* certain features or *participate in* certain relations or the like. We are wont to reply, "oh, well that explains it!" when we are given the right specification of the individual involved, but is this simply a turn of phrase? Or perhaps we should take the exclamation literally and accept that specification is itself a sufficient form of explanation. That's the central claim on offer here.

---

[1] Some philosophers hold that there is a distinction between description and explanation such that certain accounts of an event that might be called explanations are in fact not explanations; they are instead descriptions. If one tells a story of how a die landed on the 6, it will involve a number of independent factors, thresholds, and other radically contingent parameters that make it such that no principled account can be given of why the die should have landed as it did. One, at best, could offer a description of the events as they actually unfolded since one can't offer an explanation in terms of what in principle *should have* unfolded given certain factors.

[2] Aristotle discusses this and similar examples in Physics I and Metaphysics V. Aquinas also discusses similar cases in *De Principiis Naturae*, Caput 5, §39 and §40; and *In Metaphysica* 5, Lectio 3, §787 and §789; *In Phys I*, lect 13, no 112.

Using resources from the Aristotelian tradition—where kind explanation is commonplace—we can begin to form a defense of specification as a form of explanation. I'll focus my attention on Aquinas as a representative of a scholastic-Aristotelian view of kindhood and formal explanation. First, we'll need to understand why kind explanation might be mysterious and to this end it will be helpful to get into the mind of a kind explanation skeptic in section 1. Next, section 2 will look to the corpus of Thomas Aquinas for lessons on kind explanation and then section 3 will develop a picture of kind explanation? Finally, section 4 will present a positive case for kind explanation and respond to a few worries.

1. The Skeptic

How is it that merely citing a kind of thing could provide an explanation? Let us begin by getting into the intuitive headspace of a skeptic about kind explanation. To be clear, our skeptic doubts that kinds can do *sufficient* explanatory work on their own. They don't doubt that kinds can play any explanatory role whatsoever.

It's helpful to begin with a bald notion of a kind: a kind is just a category that we can use to sort the world around us. On such a conception, it is natural to think it strange that we should be able to merely specify the individual in question and thereby provide an explanation for some feature or activity of the individual. It is inescapably odd to attempt to explain the fact that a particular dog has four legs by appeal to the fact that it belongs to some category if belonging to some category is merely a matter of being characterizable or, worse, sortable by that category. It's almost as if we're asking "why did that M&M candy taste like caramel?" and receiving the answer "because it's a green one."[3] If typing individuals is more arbitrary than principled, then it's difficult to see how typing individuals explains anything at all—let alone explaining *on their own*.

[3] Of course, there are contexts in which such an answer would be perfectly acceptable, but not in a great many perfectly normal contexts.

But what if we make our conception of kinds a bit more robust? A natural kind in particular might be a historically-contiguous set of individuals that are similar on account of a set of reproductive, selective, and otherwise *homeostatic* mechanisms which ensure a great deal of similarity amongst individual members of the kind.[4] So being a member of a kind is in part owing one's possession of features and dispositions and the like to one's ancestors and the reproductive mechanisms involved and also to the other homeostatic mechanisms which ensure that one will be fairly similar to one's conspecifics. Now it starts to feel a bit less mysterious how citing a kind to which one belongs could provide an explanation of one's features in that a kind comes along with a whole host of explanatory resources. Being a dog is more than just being classified a certain way, it's being subject to a host of processes and mechanisms which ensure that one has a high probability of having certain features. And should Fido have such features, it will be *because* Fido is subject to the processes and mechanisms which play a role in defining what it is to be a dog—i.e. because Fido is a dog.

Nevertheless, the skeptic may be unimpressed, even with a far more robust notion of a kind. We don't, they might say, actually *explain* the feature in question by specifying the individual that possesses that feature. Instead, we're merely doing one of two things. We're either a) removing the *need* for an explanation by demystifying the source of confusion or mystery or b) providing a sort of shorthand for the *true* explanation by bringing into the common ground a host of truly explanatory facts one or more of which can serve as a true *explaniens* for the feature in question. The former possibility can be illustrated quite clearly by the following anecdote.

> Jayesh and Anita are at the zoo and Jayesh notices that one of the puffins is unusually tall. He turns to Anita and asks, "why is that puffin so much taller than the others?" Anita responds, "because it's not a puffin, it's a penguin." "Oh," Jayesh retorts, "that explains it!"

[4] Here, I'm explicitly recalling Boyd's Homeostatic Property Cluster theory of natural kinds (1999a, 1999b).

Has Jayesh been given an *explanation* or has the need for an explanation been removed? Even though Jayesh responds by saying "that explains it," it still may be more accurate to say that he no longer *needs* an explanation. Again, imagine you and a friend are in the city and you see a figure that you take to be a street performer. "Wow!" you say, "that street performer is amazing at pretending to be a statue!" "Well," your friend replies, "that's because that *is* a statue." Has your friend, in specifying the figure properly, explained why it should be perfectly still? Or have they merely undercut your reason for being surprised in the first place? In asking how a human being is able to stay so still, you are presupposing that it is a human being in front of you. In respecifying the individual as a statue, then, is your friend explaining how this individual is perfectly stationary or are they merely removing the need for an explanation? The skeptic may answer that in these cases we are merely removing the need for an explanation rather than providing an explanation for the feature in question.

The possibility that citing a kind is merely shorthand for the *true* explanation has perhaps more intuitive pull: what explains why any particular wolf has long canines is either i) that it needs them in order to grasp and tear meat and prey or ii) that its genes code for certain developmental processes which produce long canines and it in fact developed in that way. Merely saying "it's a wolf" seems hardly to provide as much of an explanation as either of these. It in fact might be a patently inappropriate answer to the question, "why does that animal have such long canines?"[5]

Another tool is at the skeptic's disposal in raising worries for kind explanation. They may point out that kind explanation bears a dangerous resemblance to *virtus dormitiva* explanations. The namesake example of dormitive virtue explanations is that depicted by Moliére in *The Imaginary Invalid*, when a character responds (after some careful thought) to the question of why a certain

[5] The skeptic may also find themselves influenced by a particular framework for thinking about explanation, one according to which explanations are relations between events and event-types. On such a conception, it is at best mysterious and at worse incoherent how being a kind of thing—hardly an 'event', one might think—can explain the possession of a feature or a particular behavior.

chemical has a soporific effect with the claim that it has a "dormitive virtue." A *virtus dormitiva* explanation is therefore a vacuous explanation—it is an explanation that accomplishes nothing in that it posits a principle of explanation that has no epistemic independence from the outcome it was posited to explain. So instead of saying *what it is* about the world that brought the *explanandum* about, we simply say *that* there is something about the world that brought the *explanandum* about—we state exactly what was presupposed in the asking of the question.

If we posit dispositions, say, in order to explain a particular diachronic outcome, then we might be accused of failing to explain the outcome. As an answer to "why does that lightbulb emit light every time we supply it with electrical energy?", the response "it's disposed to do that" appears at first to be largely inadequate. At worst, it's vacuous; but at best, it's simply uninformative. We'll address this and the other worries in the final section of the chapter.

It is now clear that we who hope to defend kind explanation are in a bad way. It doesn't seem—at least from a particular perspective—like we are dealing with a form of explanation. Instead, it appears that we are discussing a phenomenon in the pragmatics of explanation where sometimes people are informed about the individual in question so as to be able to formulate the true explanation on their own or other times are given the resources necessary to infer what the true explanation must be. So much for kind explanation. After extracting some resources from Aquinas' discussions of natural philosophy, though, we can begin to see that kind explanation is in fact a form of explanation—kinds can explain certain features of their individual members.

2. Kind Explanation in Aquinas

One way of reading the Aristotelian tradition is as a multi-generational debate, in part, on the nature, necessity, and subject matter of different forms of *explanation*. On such a reading, Aquinas represents an initial Latin attempt at grand synthesis of Greek, Arabic, and Latin discussions of explanation. As such, his corpus is in many ways the result of centuries of discussion, position, and observation

about the topic under analysis in this paper. As we read through his corpus, as anticipated, we find a complex and multifaceted treatment of explanation in general in addition to his treatment thereof with respect to specific domains.

Lacking the space for a full discussion, I'll instead highlight some relevant aspects of Aquinas' view of explanation. The most salient aspect of his account is that explanation is *intensional* in that entities that enter into explanations are characterized with greater or lesser specificity, substantiality, and accidentality. A human might help explain one change *qua* thermodynamic entropy engine and another change *qua* practically rational agent.

If Francis is building a house, we'll naturally wonder exactly *why* this is happening.[6] Is this to be Francis' house and so Francis is a *future homeowner*? Is Francis the employee of a contracting company and so a *carpenter*? Is Francis a *property developer* hoping for profit from the venture? Is Francis the drugged *zombie* of a cruel shaman taskmaster with a workforce of laborers who lack rational intention? Depending on how we characterize Francis, we'll be offering different explanations of Francis' activities. Similarly, one particular agent performing (in a sense) one particular activity can be specified in different ways so as to explain different features or activities. Marisol is *qua potter* shaping clay into a pot, but *qua radiant heat source* is heating the room, and *qua small business owner* is generating stock to sell. What's more, when asked why each of these activities is happening, we can simply say "Marisol is a potter," or "Marisol is a radiant heat source," or "Marisol is the proprietor of a small pottery shop," and we will have thereby explained the activity in question.

Similarly, depending on how we characterize the *explanandum*, we will demand different sorts of entities as an *explaniens*. If we're interested in how it is that a rock came to be formed in such a way, we might offer explanations as diverse as landslides, volcanoes, cosmic accidents, and artists. If,

[6] For discussions of similar points using similar examples, see *De Principiis Naturae*, Caput 5, §40.; *In Metaphysica,* Book 5, Lectio 3, §787 and §789.

however, we are after the explanation of why that statue has such a melancholic visage, then we are restricted to explanations that appeal to artistry of one kind or another.

This explanation *qua* or explanation under a description is a sort of artifact of a deeper fact about Aristotelian explanation: there is a sort of explanation which fundamentally involves the nature of a kind of thing. This sort of explanation is commonly called "formal causation" or "formal explanation." The 'form' of a thing in Aristotelian natural philosophy can be understood in a variety of ways. One might think that the form is the structure that attaches to being a kind of thing, or that it is the set of dispositions that go along with that structure, or alternatively that the form just is the thing considered as a member of a kind rather than as an individual. According to the final possibility, a kind is a sort of lens through which a thing comes into focus in a new way—indeed, sometimes *what the thing is* comes into view. We can see a computer as a chunk of metal and silicon and plastic, or we can see it as a computer. When we see it as a computer, a whole host of explanatory resources are available to us that wouldn't be available if we were looking at it in terms of the dispositions of its material parts. For instance, we know what computers are *designed and made for* and so we have a ready explanation of its activities and of the organization of its parts. Most importantly, perhaps, we'll have resources available for understanding what is *supposed* to happen and so for identifying when a part has failed.

Formal explanation is explanation that appeals to *what a thing is* in explaining why it should be the way it is. Why is the nose on this statue so porcine? Well, that's a statue of Socrates—that is, Socrates' physical shape serves as the model on which this statue is structured. Why does this furniture have this large flat platform? Well, that's a bed—that is, the function of a bed is to support a mattress and that function dictates particular structural facts about any given bed. If a bed is to be a good bed, then *ceteris paribus* it should have a flat platform big enough to support a mattress of the

intended size. Functions set normative constraints on structures[7] and so structures characterize the kinds of things with those functions. Being a member of a kind, therefore, explains why one has certain features.

Much of the explanatory power, therefore, of formal explanation comes from its relation to teleological explanation. Forms are intrinsically tied to their functions: the long sloping contour of the spoon handle is formal and functional in that it's a structure the spoon is *supposed* to have by design *in order that* it might fit the human hand more aptly. Forms and functions, therefore, cannot readily be separated. "Is it possible for something with identical structure or form to have different functions than it now has and to lack all of the functions that it now has?" It's unclear the Aristotelian can make sense of such a question. This isn't owing to a lack of conceptual resources, moreover. An Aristotelian such as Thomas Aquinas, for instance, can't understand the question because "structure" isn't mere arrangement of parts or physical disposition—it instead essentially involves the function for which the thing is structured.

Not all of the explanatory power of formal explanation traces to the teleology governing kindhood, though. There are characteristic features that aren't specifically tied to the functions of a kind. A now-famous example is that some rhinos have single horns and others have double horns.[8] This is thought to be the result of genetic drift—of more-or-less accidental genetic deviation as a result of brute physical separation or some other factor that splits the gene pool into two or more subgroups. An Asian rhino has one horn; whereas an African rhino has two. Why? Not for any

[7] This is especially clean in artifactual cases. In the biological world, species and their environments coevolve to suit one another. An ant not only evolves to occupy a particular niche, but also carves that niche out throughout eons. The niche construction school of thought therefore holds that it's in principle impossible to separate the functions that place normative constraints on a species and the structure of the species that evolved to serve those functions. Both evolve in parallel and so must be understood in relation to one another. Cf. (Odling-Smee et al. 1996, 2003).

[8] Gould and Lewontin (1994) discuss this as an example of non-adaptationist evolutionary change. In this case, it is caused by genetic drift rather than adaptation.

functional reason, apparently.[9] Simply due to the accidental trajectories of their relative histories. If this is true, then the fact of the characteristicness of single horns for, say, Sumatran rhinos isn't dictated by the functions of the species. It is instead dictated by the accidents of history. Formal explanation—in this case explaining the physical structure of a Sumatran rhino by appeal to its being a Sumatran rhino and therefore by appeal to its having the form of a Sumatran rhino—sometimes explains features in a way that is somewhat more divorced from the teleology of the kind. It follows that kind explanation doesn't reduce to and is not shorthand for teleological explanation.

This isn't to say that these sorts of explanations are divorced from teleology full stop. In certain cases, the form *just is* the *telos* of the process of becoming of an individual of the species—regardless of what sorts of *teloi* govern that form. For instance, the form of a rabbit (of a particular species of rabbit, that is) is the *telos* or end goal of the process of gestation for a rabbit fetus. The process is "trying" to make a baby rabbit and then the baby rabbit is "trying" to become an adult.[10] Any deviation from the creation of a well-formed rabbit infant is ipso facto a deviation from the norms governing becoming a rabbit. Forms are therefore tied fundamentally to ends, but their explanatory power doesn't reduce to the explanatory power of ends.

Let us expand a bit on Aquinas' conception of kinds before proceeding to a positive account of kind explanation. For Aquinas, as for Aristotle, what it is to be something of a given kind in the fullest sense is tied to being related to a particular final cause[11] that makes sense of (unifies and makes intelligible) a set of features one has. The final cause is also a standard of evaluation so that if

---

[9] Of course, it is in principle possible that there might be differences in function between these horns, but the received view is that there is none. The fundamental point here is that functional differences don't always explain or even track phenotypical differences: sometimes functionally-equivalent phenotypes come about as a result of random mutation or genetic drift and there is no explanation for the polymorphism in terms of function or adaptation.

[10] Of course there is some analogy in play here between trying in the truest sense and "trying" in this liberalized teleological sense.

[11] For most all Aristotelians, humans are not *only* rational, but are instead rational *in addition to being* sensitive, vegetative, locomotive, etc. For Aristotle, at least, each of these subsidiary final causes is in service of the final cause of the species and so there is in a sense only one final cause even if there are at the same time multiple final causes that serve as standards of evaluation in various respects.

one doesn't in fact have a feature one *should*, the final cause makes intelligible why one *should* have that feature. For incompatible or rival features, one *ought* to have one of the set of rival features, but of course isn't in a bad way for not having *all* of them (which would be impossible). So the final causes, for Aquinas, not only serve to define the kind but also serve as standards of completeness, perfection, or excellence for individual members of the kind. Aquinas believes that being a member of a kind is less like fitting into a group and more like participating in an activity. To participate in an activity is inherently normative: there are standards to fully participating and participating well and these standards are built into the activity rather than imposed from without. Consider the activity of playing the guitar, where there are exemplars and norms which set the standard against which one compares oneself—indeed playing the guitar in a particular style *at all* subjects one to such standards of evaluation. To be a member of a kind, therefore, is constitutively to be subject to a set of norms—to be in an evaluative field such that one is evaluable as a more or less perfect member of that kind with respect to the norms and functions governing the kind.

3. What, then, is Kind Explanation?

Before we give an account of how specification might by itself explain feature possession, it will be helpful to say more about the structure of kinds in light of Aquinas' conception of kinds and explanation. Here, the focus will be on what we might call—to reflect Aquinas' theological leanings—"created kinds" or kinds in the natural world. This would include natural kinds, but it would also include artifactual and social kinds, for instance. It's an open question the extent to which these reflections extend to abstract kinds or eternal kinds or the like. A kind—at least the sort of kind which can ground kind explanation—is a complex of features, relations, conditions, and activities that defines a group of things insofar as they are evaluable with respect to that complex—the components of which serve differentially as standards of evaluation. Most fundamentally,

though, a kind is a *functional role* within is a system or organization and so insofar as a thing plays that role (is a member of that kind), it is evaluable with respect to the standards associated with that role.

More concretely, being a hammer is constitutively to play (or at minimum to be evaluable with respect to) a role: the role of hammering nails and pounding hot metal in a smithy and doing all of the other things that hammers do. If this is what it is to be a hammer, then by specifying an artifact as a hammer, I have thereby provided a set of functions with respect to which we can evaluate it and also with respect to which we can *understand* why it should be the way it is. A hammer is shaped the way it is, composed of the materials of which it is composed, and bears certain of the spatial relations it does *because* it is a hammer.

An archaeologist discovers a stone that has been chipped away at one edge. Is it an accident of geological history? Is it an artifact? If so, is it a hand axe? Each would in some sense *explain* why it should have the shape it does (though some would argue that calling it an accident of geological history isn't truly an explanation). By specifying it as a hand axe, our archaeologist has thereby placed it in an evaluative space with countless dimensions. They now have available to them a whole host of explanatory resources that before were merely background possibilities. The edge is sharp because it is an axe and it is good for axes to be sharp. Our *homo habilis* ancestor shaped it this way because they needed it to perform a certain function and being shaped that way allows it to at least minimally perform that function. Our archaeologist can also think in terms of accidents and purposeful features. The nasty chip in the middle of the rudimentary blade might have come from misuse or accidental breakage or it may have served some more specific purpose than merely chopping or served the function of chopping in a particular way—perhaps in conjunction with a swiping motion the "tooth" formed by the chip allowed the axe to bite the wood more like a saw would.

This seems to work quite well for artifacts, but what about natural objects? Biological cases are the most salient here. Sticking with the example of the archaeologist, now they find a porous chunk of hard white material on a dig. Why is it more or less hollow? Perhaps it is a pigeon bone and pigeon bones are hollow in order to lighten the organism as a whole and therefore to enable flight. By specifying this chunk of material as a section of the femur of a pigeon, we therefore have an explanation of its size, its internal structure, and its material composition.

In live cases, it's not clearly so different. It's helpful to get into the mindset of a scientist classifying a new species for the first time. One will typically take on the methodological assumption that the specimens one have access to are normal specimens with respect to at least certain dimensions[12]—that they don't exhibit many unique features as a result of their idiosyncratic genetic structure or individual history. Next, one will likely make certain assumptions at higher phylogenetic levels: it's an animal, it appears to be a chordate, but appears not to be a mammal, so it must be a… Throughout, one will be assuming that this creature has a backbone *because* it's a chordate and not because it's a radically mutated flatworm or the like. Again, one will assume that it lacks mammary glands *because* it's not a mammal and not because it's a radically mutated mammal. One makes these assumptions in order to complete the project at hand: accurately characterize the new species on the basis of the available specimens. Each judgment along the way might be upset by observations of more specimens.

In other cases, scientists are forced to ask the question of whether a specimen is a normal member of an unfamiliar kind or whether it is an abnormal member of a familiar kind. Does it have that coloring pattern because it's not a tiger tail seahorse like we thought it was, or does it have that

[12] Indeed, if one is already convinced that this is a new species being encountered for the first time, then they *must* make the assumption that at least relative to some respects the individuals they have access to are normal or else they can learn nothing about the species itself. Of course, these are all preliminary and defeasible assumptions, but without them one simply has an individual or a set of individuals that tell them nothing about any kind of thing; there's certainly no reason to suppose that they are members of a new species at all if they aren't exemplars of their species.

coloring pattern because it's an abnormal tiger tail seahorse? Have we discovered a new species—which fact would explain the variation we are observing—or have we simply discovered a variant of a familiar species—where facts about that particular individual rather than a kind of thing will explain why it has the features it does?

In a more evaluative register, we can ask questions about features that appear to be defects or privations considered with respect to one species but perhaps are not defects or privations at all with respect to another species. Is that growth the result of genetic mutation or is it simply a camouflage strategy? Is that monkey lethargic due to anemia or is it in fact a sloth and not a monkey at all? Is that a poorly made finishing hammer with an abnormally rough striking face or is it a perfectly normal framing hammer with a grooved striking face? The evaluations we make of features—evaluations of normality and abnormality, adaptivity and maladaptivity, etc.—are relative to the frames of reference defined by different kinds. One's explanatory ends determine which kinds are most apt for a given individual.

Kind explanation, in light of all of this, is paradigmatically situating an individual within a normative landscape and thereby also a *selection* landscape. We're not only stating what is good for things of this kind, but also what we might expect to happen quite often because there is selection pressure in favor of such features. Chemical bonds are less stable outside of certain environmental and structural thresholds and so the chemicals which rely on those bonds are more likely to be found within those thresholds. Conformity in the biological world is rewarded by a higher likelihood of survival and reproduction. Artifacts are often cast aside when they fail to conform to the standards set by the reproductive practices governing the manufacture and design of that kind of artifact. It's at least tempting to think that there are likely to be more members of a kind exhibiting normal properties than abnormal. The relationship between the features that are explained by being the kind of thing one is and the longevity and propagation of individuals who exhibit those features

makes it such that one would expect there to be such a relationship. A more in-depth discussion of this suggestion can be found in Chapter 5.

More precisely, there are a few explanatory mechanisms that could be powering kind explanation. We'll need to understand more clearly what sort of thing is going on in kind explanation. Here are some (mostly mutually compatible) possibilities:

> (Not Special) A kind explanation is or implies the claim that the individual in question is not special for having the feature in question—i.e. its having that feature requires no special explanation.
>
> (Essence) A kind explanation is or implies the attribution of an essential explanatory principle or nature to the individual, and that is what does the explanatory work.[13]
>
> (Shorthand) A kind explanation is in part attributing a kind which is shorthand for:
>
> > all of the necessary information for forming an explanation of the individual's having the feature in question. Or…
> >
> > a series of explanations one of which explains the individual's having the feature in question.[14] Or…
> >
> > a series of explanatory generalizations, one or more of which ground an explanation for the individual having the feature in question.
>
> (Kind Shift) A kind explanation is in part a pragmatic shift of the focus on inquiry from the individual to the kind.
>
> (Functional Role) A kind explanation is in part attributing a functional role to an individual and so providing the material for a functional explanation of the individual's having the feature in question.
>
> (Generic) A kind explanation is specifying an individual and therefore subsuming that individual under a generalization and implicating or explicitly stating that generalization as a non-accidental or explanation-grounding generalization.

[13] This idea undergirds a host of psychological studies. For some examples, see (Barrett 2001; Cimpian and Salomon 2014; Gelman et al. 2010; Gelman and Bloom 2007; Gelman and Hirschfeld 1999; Prentice and Miller 2007; Strevens 2000).

[14] This suggestion is echoed in the theoretical discussion in (Gelman et al. 2018).

> (Capacity) As Cartwright suggests, a kind explanation is attributing a set of capacities to an individual—capacities which are at least partially definitive of the kind.[15]

Each in its own way may elucidate part of what is going on in kind explanation in conjunction with others, and so we need not necessarily settle on one as the sole account of what is going on in kind explanation. In each case, an aspect of a general picture of kind explanation is elucidated: kind explanation explicates a kind in its relationship to its environment through a functional role such that the demands placed on the kind as a result of this relationship are met by a common set of features had by many members of the kind.

One might be tempted to add Same Reason:

> (Same Reason) A kind explanation is or implies the claim that the individual in question has the feature in question for the same reason that any member of the kind in question has that feature.

However, this would be a mistake because it turns out to be in tension with those in the original set above. Same Reason and Kind Shift are particularly at odds, even if they aren't strictly inconsistent. Is kind explanation a shift to the level of the kind or a "horizontal move" from the explanation of why one chicken has feathers to the explanation for why another one has feathers? Each is a pragmatic move: Same Reason says, "apply your understanding of why other chickens have feathers to this case"; and Kind Shift says, "look to your understanding of the kind chicken rather than looking for an explanation for why *any individual* has a feature."[16] Same Reason is supposed to be the

---

[15] "For I maintain that most general causal claims--like 'aspirins relieve headaches' or 'electromagnetic forces cause motions perpendicular to the line of action'--are best rendered as ascriptions of capacity. For example, aspirins--because of being aspirins--can cure headaches. The troublesome phrase 'because of being aspirins--can cure headaches. The troublesome phrase 'because of being aspirins' is put there to indicate that the claim is meant to express a fact about properties and not about individuals: the property of being an aspirin carries with it the capacity to cure headaches. What the capacities of individuals are is another, very complex, matter. For instance, must the relevant conditions for the exercise of the capacity be at least physically accessible to the individual before we are willing to ascribe the capacity to it? These are questions I will have nothing to say about." Nancy Cartwright (1989, p. 141).

[16] If we interpret Same Reason in such a way as to make it consistent with Kind Shift, then it loses some of its spirit. For instance, we might understand the reason referenced in Same Reason to be "because it is a *k*." If "the same reason as any other chicken" is simply "because it is a chicken," then the letter of Same Reason is followed, but not the spirit.

idea that there is an explanation for why any individual chicken has feathers—it has genetics which code for a particular developmental process, for instance—and so by specifying something as a chicken in response to the question "why does that animal have feathers," I'm merely saying, "because of the same genetics and developmental processes that any chicken has feathers." Kind explanation would therefore be relying on the idea that a kind is a collection of individuals with similar explanations for the features they have,[17] but Kind Shift relies on a different conception of a kind—one according to which the kind itself is characterized by certain features due to the kind's relationship to its environment through its functional role. There is, therefore, some tension between Kind Shift and Same Reason. In light of this tension, and because Kind Shift fits more organically into the broader set, I suggest that we abandon Same Reason.[18]

With that settled, it is worth spending a moment fleshing out our understanding of Kind Shift and so the pragmatic function of kind explanation before moving on.

Kind explanation has a particular pragmatic function: it shifts the focus of the explanatory inquiry from the individual to the kind in question. When I ask, "why does that animal have ultraviolet vision?", I'm asking about an individual. In providing a kind explanation such as, "because it's a bee," my conversant is in part saying, "you shouldn't be asking about this individual, because the answer is found in asking about bees writ large." The conversant is shifting the level of the conversation from the level of the individual bee I was interested in "up" to the kind as a whole. Then another—entirely different—question presents itself: "why do bees have ultraviolet vision?" The individual bee has ultraviolet vision because it's a bee, and bees have ultraviolet vision in order to detect the ultraviolet light patterns reflected by flowers (or for some historical reason).

---

[17] This may very well be the conception one gets on a Homeostatic Property Cluster conception of kinds, though this deserves a study of its own.

[18] Note that this move requires we go in for the conception of kinds suggested by Kind Shift (or offer some functionally equivalent alternative). This conception is an important part of the overall account I sketch in the dissertation, and so this move is justified by appeal to the overall project.

In investigating this pragmatic function, we might come to see that there are two questions one might be asking which both might appear on the surface to be similar to one another. In asking "why is that individual behaving that way?" (for instance), we might intend to ask one of two different questions:

(Individual) why is that individual behaving that way?

(Kind) why do individuals of that kind behave that way?

If the individual is behaving *normally* or in keeping with its kind, then the answer to the first question is in part a pragmatic move: to say that they should be asking the second question. It's also, I contend, an *explanation* of the behavior: it's behaving that way because it belongs to this kind.

4. Defending Kind Explanation

We're now in the position to make a case for kind explanation as a form of *explanation proper* rather than merely the provision of some information which either takes away the need for an explanation in the first place or allows the listener to piece together an explanation on their own. Let's make a positive case before addressing some worries we might anticipate from a skeptic.

4.1 Positive Case

Why think that kind explanation is a sufficient form of explanation? Let's explore a pragmatic feature of kind explanations before moving into a defense. By answering "why does that animal have stripes?" with a simple specification of the animal as a zebra—i.e. "because it's a zebra"— I'm implicating by means of the maxim of relevance that zebras have stripes. By answering the same question with "zebras have stripes," I will have thereby implicated that the individual animal in question is a zebra. Specification and a generic judgment go hand in hand at the level of pragmatics and so we should see them as working in tandem to do the explanatory work that kind explanation is supposed to do.

Kind explanation is simply one instance of a general schema of explanation that is uncontroversially a form of explanation proper. Let us call this the Covering Generalization Schema (CGS):

> For the name of a kind F; a predicate ascribing an activity, relation, condition, or feature φ ; and a name of an individual n:
>
> (G) F's (normally, always, probably, etc.) φ
>
> (I) n is an F
>
> [Explains why] n φ's

A generalization G states some general fact that covers the *explanandum*, an instantiation sentence I states that the individual involved in the *explanandum* belongs to the kind cited in G, and the result is that we have an explanation for why the individual should admit of the predicate cited in G.

The CGS is a liberalized form of the Deductive-Nomological (DN) model of explanation introduced by Hempel and Oppenheim,[19] according to which an explanation is a deduction from a law and an observation statement to the claim that the observed individual *must* act in a certain way given the law in question. If we've identified the law that covers the individual instance and makes it the case that things must have happened in the way they did, then we have *ipso facto* explained the individual instance. If we're in agreement that DN explanations in fact explain (and sufficiently explain) the instances they are supposed to explain, then instantiations of the CGS are *prima facie* formally acceptable explanations. Of course, the CGS is a *weaker* form of the DN schema and so this inference isn't strictly warranted. Nevertheless, accepting that one form of generalization can bridge the explanatory gap between an individual member of a kind and its having a certain feature is the first step toward seeing how other forms of generalization might have the power to do the same.[20]

[19] Hempel & Oppenheim (1948).

[20] One class of generalization might not work in the same way: at least some statistical generalizations can't carry explanatory weight in the same way as other generalizations. The (probably false) generalization that 27% of dogs have

If I claim, plausibly, that fair dice have a roughly 1/6 chance of landing with the "2" face up, then it seems I've provided all of the necessary explanatory information for why a particular fair dice should show a 2 in a standard, purely random toss. Similarly, if I note the seemingly true generalization that massive bodies never travel faster than the speed of light, then I'm in the position, it seems, to explain why any particular massive body can't travel faster than the speed of light. That individual body is not special—it can't travel that fast *because it's a massive body* and not for any special individual reason. If we want to know why the law holds, then we're asking a different question altogether. But we seem to have an explanation for why my bouncy ball can't travel faster than the speed of light: nothing can![21]

So, the claim on offer here is simply that normality claims are explanatory generalizations in the same way that laws are. The information that a) an individual belongs to a kind, and b) that individuals of that kind normally have a certain property serves as the basis for an explanation of why that individual has that property.

A second line of argument posits a particular characterization of nearly *all* of our explanatory practices and attempts to show that kind explanation is no different. Scientists along with the lay folk, in seeking an explanation for an event, are nearly always in the business of identifying and refining the proper grain of specification of the entities involved in that event. If my car won't start, it might be because my distributer cap was stolen. So rather than simply being a car, which would explain the car's starting, mine is a *car without a distributer cap*, which explains its failure to start. Cars without distributor caps simply won't start and they won't start *because they are missing their distributer*

---

spotted fur coloration doesn't clearly explain why any individual dog, Fido, might have spotted fur. We'd be better off knowing that Fido is a dalmatian if we want to know why Fido has spotted fur. Worse than that though, it's not clear we've been given any explanatorily relevant information at all in the statistical generalization.
Another class of generalization is technically a statistical generalization even though it appears to be universal. A famous such example is "all of the members of the US Supreme Court have odd social security numbers." This is an exceptionally modally fragile generalization, as are all accidental generalizations. This judgment won't carry the explanatory weight necessary to fit into the CGS, either.

[21] Katie Elliot has been influential in my thinking here. Her work in manuscript is most relevant.

caps. We're not, that is, merely looking for arbitrary sets of things which happen to all and only have the feature we're interested in. Instead, we're looking for the proper explanatory characterization of the item involved such that we can understand why the *explanandum* occurred.

When we wonder why a sphere is moving in an elliptical pattern, we first think of it as a massive body and find that the resources we get out of the specification "massive body" don't give us what we need to explain elliptical motion. Then we add "spherical" to the specification and find we still come up short. Finally, we add "orbiting around a body more massive by a factor of 100,000,000" and we suddenly have the right level at which the explanation is supposed to happen. The question "why is that body moving in an elliptical pattern?" is answered by "because it's a spherical massive body orbiting around a body more massive by a factor of 100,000,000." If one still does not understand, then it's because one does not know what it is to be that kind of thing. I can explicate that kind, but in performing such an explication, I won't be explaining why *that* particular sphere was moving in *that* way. Instead I'll be specifying that individual under a kind which bears an explanatory relationship with the *explanandum*. The same, I contend, goes for zebras and computers and walnuts.

Sociological research is particularly apt for this description. When looking at corporate hiring data in the interest of determining hiring biases, a sociologist will control for potentially confounding variables. Is it socioeconomic discrimination or racial discrimination? Again, is it racial discrimination or linguistic or cultural discrimination? A sociologist will find out by controlling for socioeconomic status, linguistic differences, and/or cultural differences and checking if there is still a statistically significant divergence in outcomes between racial groups. In carrying out this commonplace explanatory practice, the sociologist is doing in a particularly clear way what in some sense all scientists are doing: identifying which kinds of things are causally or otherwise explanatorily related to which. They are identifying which outcomes one can explain by appeal to a given kind.

The folk category jade, it turns out, consists of two different sorts of mineral: jadeite and nephrite. Geologists confirm this by appeal to different features that jadeite has from nephrite. Instead of there being one mineral that sometimes acts this way and sometimes acts another way, a geologist will claim that there are two different types of minerals that each virtually always acts one way or another. The geologist is identifying kinds which are tied to mineral structures and stopping their analysis when they've found a fit between the kind of mineral and the feature in question. They aren't, however, just interested in fit between otherwise arbitrary categories; they are looking for an explanatory link between categories which can do explanatory work.

So, the positive account has two prongs. First, we noted that a kind explanation is formally similar to the DN account of explanation, where an instantiation sentence situates an individual within the domain of a generalization and the generalization offers a link to the feature being explained. Second, I offered a characterization of a number of explanatory practices in the hopes of making plausible the claim that what we're doing in nearly all of our explanatory practices is identifying explanatory links between categories and therefore identifying the proper grain of specification of the individual causes and effects and individuals in question. Together, they constitute an intuitive argument that we're not dealing with some alien form of explanation in kind explanation. To the contrary, kind explanation is structurally and materially similar to uncontroversial forms of explanation.

4.2 The Skeptic's Worries

Let us now take on a few more pressing objections from the skeptic in the hopes of further clarifying the account.

Removing the Need for an Explanation

One worry we discussed at the top of the chapter was that specifying an individual, rather than providing an explanation for its having a particular feature, instead merely performs a pragmatic

task: removes the need for an explanation in the first place. Consider the exchange from before between Jayesh and Anita:

> Jayesh and Anita are at the zoo and Jayesh notices that one of the puffins is unusually tall. He turns to Anita and asks, "why is that puffin so much taller than the others?" Anita responds, "because it's not a puffin, it's a penguin." "Oh," Jayesh retorts, "that explains it!"

It's not that there is no explanation or that the explanation need not make reference to the individual's being a certain kind of thing. Instead, the idea is that the need for an explanation for an individual's having a particular feature is, as a pragmatic fact, a function of there being no explanation available in the common ground. So, in specifying an individual, perhaps what we're doing is, instead of updating the common ground with a new explanation, merely making salient a set of explanations already in the common ground. If that's all that specification accomplishes, then it doesn't seem like one can responsibly call kind explanation a form of *explanation*. It would instead just be a move in the language game governing the pragmatics of explanation.

In response, it's helpful to note that while there may be an answer to the question "why are penguins taller than puffins?" in the common ground, there clearly *cannot* be an explanation for why that *individual* bird is taller than the surrounding puffins. That is, Jayesh cannot be aware that the bird is in fact a penguin.[22] Otherwise the exchange between Jayesh and Anita above simply couldn't occur in good faith. If this is true, then Anita is in fact introducing an explanation into the common ground in noting that the bird is a penguin rather than a puffin.

Sure, the skeptic will say, but the explanation for why that bird is taller isn't that it's a penguin. Instead, the explanation is that penguins are taller than puffins. The fact that this individual is a penguin—while it helps pull the explanatory weight—is hardly an explanation in itself. This response leads our discussion naturally to the second worry of the skeptic:

[22] Otherwise, Jayesh would have to be aware that the individual is a penguin, but not aware that the other birds were puffins. This is possible, but not what is intended by the scenario.

True Explanation at Kind Level

The previous worry was that specifying an individual removes the need for an explanation rather than explaining, and one motivating thought there was that there was an explanation of the individual's having the feature in question, but that specifying the individual merely *identifies* for us of what that explanation is. A serious worry from the skeptic—one that goes hand in hand with the previous worry—rears its ugly head: why not simply think that in answering a question of the form "why is that individual *φ*ing?", merely saying "because it's a *k*" or the like doesn't seem to *explain* the behavior in question, it instead just says "if you want an explanation for the behavior, you should ask why *k*'s behave that way." We might think that the *real* explanation for any particular feature is found at the level of the kind, and so in specifying an individual we haven't *ipso facto* provided an explanation. We have instead merely told our conversant where to look to find the real explanation.

More concretely, being a hammer is in part having a number of functions and it's not—so the thought goes—one's being a hammer that explains one's having particular features, but the *functions* that go hand in hand with being a hammer that explain them. For instance, the explanation isn't that it's a hammer, but instead that hammers are for pounding, which explains my individual hammer's having a metal head and flat pounding surface. If we ever have the exchange: "why does it have a metal head?" "because it is a hammer," then all the response has done—according to this worry—is provide us with the true explanation: hammers are for pounding. The fact that the individual is a hammer would then play a mere pragmatic function and no proper explanatory function.

Importantly, though, there are *two different explanatory questions* in play and so, I think, two different explanatory answers. Whereas the skeptic thinks that there is only one explanation for why any penguin might sit on its eggs—the explanation for why penguins sit on their eggs—the defender of kind explanation thinks instead that there's one explanation (or a set of explanations) for why an individual penguin might sit on its eggs and another (set) for why penguins sit on their eggs. We

explain why an individual animal sits on its eggs by specifying it as a penguin and then we explain why penguins sit on their eggs by appeal to a functional or historical explanation or the like.

As discussed in the previous section, one form of explanation of a particular event is simply to subsume it under a generalization that can carry explanatory weight. To specify an individual penguin as a penguin can explain normal penguin features, behaviors, etc. because it situates the individual within a host of explanatory generalizations. The most pointed response to this worry, and to some extent to the previous worry, is that there must be some link between the kind and the individual in question. If we ask why an individual has feathers, and the response is that robins have feathers, we haven't completely been given a response unless we also find out that the individual is a robin. As noted earlier, there is an implicature here that must be included in a complete account of the explanatory phenomenon: when I respond with "robins have feathers," I'm implicating that the individual bird in question is a robin; but conversely when I answer with "that's a robin," I've implicated that robins have feathers. Either way, the question "why do robins have feathers?" is a different question than "why does that animal have feathers?"

Virtus Dormitiva

Finally, might kind explanation simply be a form of *Virtus Dormitiva* explanation? That is, might responding to "why is that bird black and white?" with "because it's a penguin" be akin to ascribing a "dormitive virtue" to a sleeping draught? The response to this worry involves taking great care in tracking the questions being asked and individuating answers with respect to those questions. When we ask, "why does that lightbulb emit light *every time* we supply it with electrical energy?" we are positing a *diachronic explanandum* the explanation of which can be of two forms. We might say "it's supposed to do that" in some way or another. Alternatively, we might say "it's an accident." The second sort of answer posits a different explanation for each individual outcome grouped together into a single *explanandum* by the question. "Why do I keep getting picked last in dodgeball?" "Well

last month you had a broken arm, so no one wanted to have you on their team, and then two weeks ago, Shantel and Shamik were team captains and neither of them like you very much, and then…" Each outcome has an explanation of a markedly different kind and so *there is no one answer* to the original why question. The first sort of answer posits a single explanatory principle—and here a disposition is the most salient kind of principle—for the *explanandum*. In short, we have to first accept the 'why' question as having a non-accidental answer before we can ask what the explanatory principle in question is.

In order to flesh out this distinction between the accidental and the non-accidental or principled, it's helpful to note that principled connections in the nature world are *modally robust* in that they are counterfactual supporting. Principled connections are *reliable* in that they tend to remain true in spite of shifting parameters and changing factors.[23] Conversely, accidental connections are by their nature *modally fragile*—we can't rely on them remaining true through a wide range of values of the relevant parameters. Change one fact in our dodgeball example above and suddenly the pattern is broken, but if the explanation is principled—the person is wildly unpopular among almost everyone, or is terrible at dodgeball—then even if Shantel weren't team captain two weeks ago, the subject of the regularity would *still* have been picked last. So in positing a principled explanation for a regularity, we are doing more than offering a vacuous "it has an explanation" sort of explanation: we are instead positing a modally robust explanatory connection between a principle and the regularity in question. We simply haven't yet determined what that principle is and so, by the lights of most, haven't truly given an explanation.

Once we've determined that there is a principled explanation of some regularity, the next question that will naturally arise is "how is it that it is so disposed?" or "why is it so disposed?" This is often confused with the first question—"why does it do that regularly?" That these are different

[23] Reminiscent of Woodward's (2000, 2003) "invariances."

questions can be seen in the natural answers we might give to each. Pragmatically speaking, the only sort of answer that seems acceptable to the 'how' question is to explain the mechanisms involved or the like. To the original 'why' question, alternatively, we can give an answer in terms of mechanisms or the like (because we often overlook the distinction between these two questions in conversation) *or* we can give an answer of the kind rehearsed in the previous paragraph—either "it's an accident" or "it's so disposed."

Another sort of accidental explanation must be ruled out as well. Sometimes the kind itself is accidental and so there can be no principle governing the perceived regularity of a feature in the population of members of that kind; alternatively, sometimes the kind cited in the call for explanation is simply at the wrong level of grain. For example, suppose I ask the question "why are evergreens dying in large numbers in the Sierras?" I've posited a kind that might even be a natural kind, but it's the wrong kind to stand in an explanatory relationship with the epidemic of tree deaths currently occurring in the Sierras due to a bark beetle infestation. It is *pine trees* that are dying in large numbers, not evergreens. Bark beetles, it turns out, do not attack cedars and so do not attack all evergreens; they burrow in pines, more specifically. An appropriate answer is therefore that evergreens aren't dying in large numbers, *pines* are. Similarly, if I ask, "why are all of these dogs with collars making this barking sound?" I've posited an accidental kind—dogs-with-collars—and so perhaps there are no explanatory generalizations we can make about them beyond the pseudo-analytic generalizations we can derive deductively from their being dogs with collars. We must fix onto the correct kind before we are able to ask what explanatory principle is governing the apparent regularity in question.

What is the point of this pedantic distinction between these different questions? In seeing that it is a true distinction, we progress from thinking that disposition ascriptions are non-answers to the relevant 'why' questions, to seeing that a disposition ascription advances our explanatory

understanding by ruling out accidentality. Kind explanation as a whole works in a similar way: in citing the kind as the explanation of the outcome, we are thereby ruling out accidentality—we are positing a principle the content of which can be filled in by means of functional or historical explanations.

The argument thus far hasn't served to vindicate kind explanation from the charge that it is a form of *virtus dormitiva* explanation. To the contrary, we have sought to vindicate *virtus dormitiva* explanations. The problem with such explanations was supposed to be that it posited an explanatory principle that was in no way independent from the *explanandum*. It was positing something the existence of which is all but presupposed in the asking of the 'why' question. "What's the explanation for this?" "There's an explanation for that!" This response hasn't clearly illuminated much of anything at all. What I've been arguing thus far, though, is that there *is* a response to the question that would be both an appropriate response and a contrary one to the *virtus dormitiva* response: an accidental explanation of the form, "well, one time the explanation was *x*, and another time the explanation was *y*, and…"

This likely won't be satisfying. Given what I have said above, though, it should be clear how this fits into the overall picture: an individual's having a particular feature either has an individual historical sort of explanation—i.e. an accidental explanation—or it has a principled explanation that involves specifying that individual by reference to a kind. That's the explanation for that individual's having a normal feature: it is the kind of thing which normally has that feature. The explanation we might be expecting—and therefore the reason that we might be dissatisfied with these responses—is an answer to the question "why do individuals of that kind have that feature?" This is an importantly different question.

Chapter Three: The Explanatory Link Account of Normality

1. Introduction

The words 'normal' and 'normally' are often used in a statistical or frequentist sense. That is, they are often used to mean something like 'usual(ly),' 'frequent(ly),' 'ordinar(il)y,' or 'regular(ly).' They are also, however, often used in a different sense: one with a more *normative* rather than merely statistical or frequentist sense and inferential structure.[1] Normality judgments are judgments which explicitly or implicitly employ this concept of normality. As a group, they tolerate apparent exceptions and changes in the prevalence of the predicated property. But apart from their tolerance and insensitivity to prevalence, what are these judgments? Can we say something informative about this apparent family of generalizations?

Bernhard Nickel (2016), it seems, might agree that we can. While giving an account of bare-plural generic sentences, he offers an account of the sort of normality at play in both discussions. He thinks normality is dependent on the ability of our experts to select the right kinds of explanations of the presence of a property in the population of members of the relevant kind. Those are the explanations which ground the normality of the relevant properties. I take issue with this account.

The account I'll develop is one according to which these claims are true if there is an explanatory link between the "kind" (or the individual through time)[2] over which the claim generalizes and the feature, activity, or conditions posited for the kind. In short, the normal outcomes are those

[1] Martin Smith (M. Smith n.d., 2007, 2010, 2014) has directly discussed this sort of normality at length. Direct discussions also appear in Alexander (1973), Millikan (1984), Pietroski and Rey (1995), Veltman (1996), Schurz (Schurz 2001a, 2001b, 2009b, 2014), Yalcin (2014), and Nickel (2016). Leplin (2009), and Burge (2003, 2005, 2010) and Sosa (1991, 2007, 2009, 2010) have discussions of normal conditions that at least implicitly appeal to non-statistical normality as well.

[2] When I use 'kind' throughout the discussion, it will stand in for "kind or individual through time."

explicable with reference to the kind of thing for which they are normal.[3] The account is based on the *Aristotelian* intuition that it is normal for a carpenter to build a house because being a carpenter *explains* one's behavior in building a house.[4]

I'll offer some preliminaries in section 2, address alternative views in section 3—most prominently Bernhard Nickel's account in section 3.1—and finally begin to develop my own *Explanatory Link* account of normality in section 4.

2. The Set Up

Normality judgments are a motley crew at the level of syntax[5] and so simply doing a lexical semantics of the word 'normal' or 'normally' won't capture the phenomenon as a whole. These judgments are expressed using bare plural, determinate, and indeterminate generics;[6] sentences with the adverbial 'normally'; sentences with what appears to be the quantifier or otherwise domain-restricting modifier 'normal'; the grammatical constructions "it's normal for *k*'s to φ," "a normal *k* φ's," and "*k's* normally φ"; among perhaps many others. Moreover, 'normal' and 'normally' appear to have at least two senses—as noted in the introduction—and so a semantic project is further

[3] In other non-rival formulations, the normal is that which is included in being a kind of thing, or that which doesn't call for an explanation, or doesn't need a *special* explanation.

[4] Echoing Aquinas' *Commentary on Aristotle's Metaphysics* VI.2.1174, where he interprets Aristotle's *Metaphysics E*.2.1026a 33-1027a 28; and also Aquinas' *Summa Contra Gentiles* 3.97.8: For since each thing acts on account of its form, so does it suffer passion and is it moved on account of its matter. (my translation).

[5] Note that in grouping these claims together, we're not claiming that they are transmutable *salva* inferential profile. We might expect that adding or removing a "normally" might alter the inferential profile of a claim. Consider the case of transmuting from "we meet at Joe's on Friday for breakfast", which can be used to set a policy for the future; and "we normally meet at Joe's on Friday for breakfast", which it seems cannot. Ströβner (2015) supposes in her discussion of normality that there is a distinct logic to normality laws of the form "*k*'s normally φ" when compared with normality claims of the form "it is normal for *k*'s to φ." I am hoping to do something a bit more general here in that I want to understand what judgments intuitively involving non-statistical normality have in common and so Ströβner's project is too narrow for my purposes.

[6] How do we in a principled and/or pretheoretical way work out the generics we intend to give an account for? Well, I won't decide in advance. Instead, I'll focus on some paradigm examples and decide "*a posteriori*" how far this account will go in giving a semantics for generics. There are certain cases, most famously "mosquitos carry the West Nile Virus", that make me think it won't (and perhaps shouldn't) cover every true generic.

quagmired at the level of the surface features of the linguistic phenomenon. With this in mind, I aim to focus in on a few focal cases individuated in terms of their *content* and then explore after the fact how far the analysis can extend.[7]

Let's get on the table some focal normality judgments to serve as data for our study. Each appears, if anything does, to involve a non-statistical account of normality.

(1) [Characteristic]
   a. Dogs have four legs.
   b. normal Dogs have four legs.
   c. Dogs normally have four legs.
   d. It is normal for Dogs to have four legs.

(2) [Functional]
   a. Hearts pump blood.
   b. Hearts normally pump blood.
   c. Normal hearts pump blood.
   d. It is normal for Hearts to pump blood.

(3) [Trajectory]
   a. A walnut grows into a walnut tree.
   b. A (normal) walnut grows into a walnut tree.
   c. A walnut normally grows into a walnut tree.
   d. It is normal for a walnut to grow into a walnut tree.

The analysis will be drawn from these types of judgments as its data, with each serving merely as a paradigm for its respective type of judgment.[8] Characteristic claims amount to the claim that the

[7] Some deny that normality enters at the level of truth-conditional semantics or anything resembling it. Instead, Veltman and Bastiaanse (e.g. Bastiaanse and Veltman (2016)), among others, argue that normality should be understood in terms of inference rules which don't have truth or falsity, but instead only have some form of validity or invalidity. There is, on this view, no truth-conditional semantics for phrases containing or making implicit reference to the concept of normality. I disagree: I think normality judgments are truth or false and that we can at least develop some heuristics for knowing when they are true or false. One reason to reject the inference rule interpretation is discussed in Krifka et al. (1995): there is no clear way to handle normality judgments embedded within other normality judgments (or perhaps embedded within any judgment at all).

[8] Other plausible types of normality judgment are reliability claims, powers ascriptions, competence or capacity ascriptions, intention ascriptions, causal claims, etc.

feature in question is characteristic of the kind in question. Functional claims are a species of constitutives that ascribe a function to a kind of thing. Finally, Trajectory claims are claims that absent interference things will progress in a certain way. They are claims that there is something like "metaphysical inertia" towards a particular goal or end state in that the modal profile of, say, a walnut includes the proposition that if things progress without interference, it will grow into a walnut tree.

One question for this study is whether or to what extent 4 and 5 shares the core features of normality judgments. We'll explore these categories in the conclusion:

(4) Dev is normally in his office by 2pm. [Habitual]

(5) It is normal for Computers to turn on when their power buttons are pressed. [Disposition][9]

4 and 5 will serve, therefore, as data for testing how far the account extends. Ascribing a habit or a disposition to a kind of thing seems importantly similar to the paradigmatic judgments we've introduced and so it will be interesting to see what they have in common and perhaps what they don't have in common.

Regarding (1)-(3), two things can be immediately said: a) exceptions to the corresponding universal statement[10]—call these *apparent exceptions*[11]—aren't *usually* exceptions to these generalizations, and b) the truth of each doesn't seem to rely (at least in any principled way) on any frequency or prevalence about the population over which the claim purportedly generalizes.

[9] I've opted for compressed examples here in the interest of space. One shouldn't read anything into this choice.

[10] Thanks to Sheldon Smith (*Personal Correspondence*) for the note that what others call "exceptions" aren't in fact exceptions but *would be* exceptions if the generalization in question were a universal claim; they are merely *apparent* exceptions. We can call an exception that truly undercuts a generic or normality claim a *real exception*.

[11] The question of what sorts of exceptions make trouble for normality judgments is a question worthy of its own treatment. Cf. Leslie (2007, 2008) for an extended discussion of related issues.

To demonstrate (a), we need only note that each is apparently true in spite of the fact that not all dogs have four legs, some hearts don't pump blood some of the time, and most walnuts never grow into adult trees. To demonstrate (b), we need simply note that even in a world where most dogs have been amputated, proposition (1) is still true.[12] Similarly, (2) and (3) will be true even when *all* of the hearts stop pumping blood or all of the walnuts fail to grow into adult trees—in any case they are certainly true when a great many hearts and walnuts fail to do what is normal to them. Each is therefore true somewhat independently of the actual regularity of the outcome being attributed to the kind (or individual) in the generalization.

The paradigm normality judgments—characteristics, functionals, and trajectory ascriptions—all also seem, at least at first glance, to have a *normative* inferential structure. If any individual in the generalized kind doesn't have the normal outcome[13]—if a heart doesn't pump blood, or a dog doesn't have exactly four legs, or a walnut doesn't grow into an adult walnut tree—then it will have *failed* or have been *deprived* with respect to that outcome and it furthermore *should* have that outcome insofar as it is what it is.

So (1)-(3) do not appear to be generalizations about *every* member of the class over which they generalize—i.e. they are not exception-intolerant laws or nomic generalizations. They furthermore do not seem to be judgments about a specific frequency or prevalence of an outcome in the class

[12] The truth of individual dispositionals like "My computer turns on when I press this button" seem to be grounded in the truth of the corresponding kind-level dispositional—(2) in this case. This individual claim generalizes over cases of pushing that button on my computer, whereas (2) generalizes over all computers (and perhaps all pushings of computer 'on' buttons). So the truth of (2) should be somewhat insensitive to particular button pushings and (5) should be somewhat insensitive to particular computers (and perhaps their button pushings). This is just what, I think, we find: my computer can fail to turn on in a great many cases because I forget to charge the battery and it still be the case that my computer turns on when I push the on button. For (2), even a large amount of button pushings, even over a long period of time, might not interfere with the truth of the judgment. Similarly, a great many computers can fail to turn on when their 'on' buttons are pushed—after all, it's likely that most computers in the world are broken or in a junk yard—and yet (2) will still be true.

[13] There are also "contrary normalities," which are cases where many different contrary outcomes are nevertheless normal. For now, we'll simply note that the discussion is riding rough shod over some intricacies. Later on, I will call these "incomplete normalities" to contrast them with "complete normalities"—normalities governing outcomes that have no competing alternative normal outcomes.

over which they generalize—i.e. they are not statistical or frequency generalizations. They are a distinct kind of generalization. With those observations in mind, we can develop a list of common features which are distinctly possessed by normality judgments as opposed to statistical and nomic judgments:

> (Tolerance) Normality judgments all tolerate (remain true in spite of) exceptions to the corresponding universally quantified judgment. That is, they don't have straightforward counterexamples.
>
> (Insensitivity) Normality judgments are all *relatively* insensitive to changes in the prevalence of the relevant outcome in the population over which they generalize.
>
> (Normativity)(?) Normality judgments entail that any individual over which they generalize, insofar as it is a member of the relevant kind, *should* have the outcome being attributed.[14]

The focal normality judgments have these features (though normativity is perhaps not universal), so it is tempting to think that a common account can be given of such claims. At the risk of losing readers taken with the pluralist intuition that no common account can be given of these judgments,[15] I'll press on. The proof of that possibility will be in the putting forth of a successful account.

A final clarification is required at this point. There are two easily-confused sorts of questions we might be attempting to answer: a) what does it mean to claim that flying is normal for bald eagles? and b) why do bald eagles fly?[16] I am interested in the former, conceptual question, at the expense of answering the latter, more theoretical or empirical question. I'm not convinced that there is a unified answer to (b), though I suspect that the functionality—broadly construed—of a feature is the answer more often than not to why it is normal for a kind. Instead, I'm attempting to understand

---

[14] I use the words 'outcome' and 'feature' to stand in for features, relations, activities or behaviors, and conditions. I am weary of the philosophical baggage attached to 'property', though I'm not sure anything is strictly lost by interpreting me to mean 'property' by these words.

[15] As a bandage, I can offer the intuition that it would be quite odd if the conceptual space were divided into the, I take it, clearly unified statistical/frequency judgments, the similarly unified nomic judgments, and then a grab bag of judgments with no unity or even family resemblance that nevertheless share features like insensitivity and tolerance if not some version of normativity.

[16] Alternatively, we might think this question should be stated as "why is it normal for bald eagles to fly?"

what we are saying when we say that it is normal for dogs to have four legs. What is the meaning of such a claim?

3. Addressing the Alternatives

There are a host of potential alternative accounts to my own. Here are a few examples with at least potential defenders in the literature. Normality judgments track and exclusively track:

> Features where, for instance, biologists have a suitable normality-grounding explanation of the presence of those features in the population of a species.[17]
>
> Features that arise from the nature of an open dynamical system.[18]
>
> Features that historically have been selected for in a "reproductively established family."[19]
>
> What *would* or *will* happen in ideal[20] or perhaps normal conditions.[21]
>
> *Prototypical* features of a kind of thing, where prototypicality is defined in terms of similarity to the prototypical member of the kind. [22]

---

[17] Nickel (2008, 2010a, 2010b, 2014, 2016) defends such a view at great length.

[18] Gerhard Schurz (2001a) appears to hold this view, though his target class of normality judgments is slightly different in that he is only interested in *complete* normality judgments. Cf. Nickel (2016, p. 180 and n. 3, p. 180) for critical discussion thereof.

[19] This is something close to Millikan's view in (1984). Understanding exactly what she means by 'Normal', though, is frustrated by the lack of a direct analysis or treatment of the concept itself in (1984). The central clue on which I'm basing my interpretation is the following passage, in which she offers a preamble to her definition of proper function. "Where *m* is a member of a reproductively established family *R* and *R* has the *reproductively established or Normal character C*, *m* has the function *F* as a direct proper function iff…" (1984, p. 28, emphasis mine). If we attend to the italicized phrase, it appears that she is claiming that *reproductively established* and *Normal* are more-or-less interchangeable when it comes to the character of a family. A character is a set of properties had in common by all members of a reproductively established family that were produced normally (1984, p. 25).[19] So a normal character is a reproductively established set of properties held in common in the first order case, and in the higher-order cases it is a set of properties got through normal production of members of the higher-order family. A normal feature, therefore, is a feature got as a result of one's reproductive ontogeny that has been established historically within a reproductive family.

[20] Note that ideal is still somewhat ambiguous between different sorts of idealization. This ambiguity comes into view throughout the literature on *ceteris paribus* laws. For a sampling of different potential interpretations, see Spohn (2002), 354, and also Schurz (2009a) for helpful surveys; *Ceteris Absentibus* comes from Joseph (1980, 777); '*minutis rectis'* comes from Fenton-Glynn (2016), though Cartwright (1983, p. 45) argues for a similar interpretation.

[21] Romane Clark's article (1973) on what he calls "prima-facie generalizations," gives voice to something like the view above, though not in so many words.

[22] Similarity is often defined in terms of distance in multi-dimensional space, where each dimension defining the space is a particular determinable like "size" or "limb count" or the like. The closer an individual is to the prototype, the more prototypical and hence more normal one is. Cf. for instance, (Neumann 1974; Rosch 1973, 1978; Rosch, Simpson, et al. 1976; Rosch and Mervis 1975)

In the interest of space, I'll focus all of my attention on the most recently defended of these alternatives: the first view, which is defended by Bernhard Nickel. In short, most of the others fail at the level of extension over judgements: historical and open dynamical systems accounts are too restrictive, prototypical accounts are too permissive. Ideal conditions accounts fail, by my lights, because they in principle cannot offer an illumination of what normal conditions are.

It will be helpful to recall, going into this discussion, that the account I'll offer at the end is roughly that normal outcomes are those the obtaining of which is *explained* by the kind of thing for which they obtain—as in the Aristotelian observation that identifying someone as a carpenter explains their behavior of building a house. Throughout the discussion, the contours of the view I'm developing—the Explanatory Link Account—will become clearer.

3.1 A Presence Account

The most plausible rival account is also perhaps the closest in its goals as it appears in the literature. Bernhard Nickell (2008, 2010a, 2010b, 2016) is interested in giving a semantics for bare-plural generic sentences by appeal to a non-statistical sense of 'normality'—one that plays a guiding role in our explanatory and inductive practices and is independent of prevalence.[23] In service of this end, Nickell need not give an all-purpose account of normality,[24] so he doesn't. I'd like to, conversely, test his proposal as an all-purpose account of normality.[25] My worry is that Nickel's account depends

[23] Nickel thinks that generics are existential quantifications over ways of being normal in a respect, where ways are contrary, but each normal, characteristics one can have (determinants of the determinable "respect"). Particular hair colors is an example of a way of being normal—humans are normally blonde, normally brown-haired, normally red-haired, etc. Respects are dimensions—like hair color—within which one or more value is normal for a kind and others are abnormal (respects are "determinables"). Reproductive method is one respect of being normal in that there are different ways to reproduce—internal gestation and live birth, egg-laying, etc.—and different species have different normal ways of giving birth. So, to say that "ducks lay eggs" is to say "there is a way of being a normal duck with respect to reproduction and that way is to lay eggs." I like Nickel's analysis of generics and think it's the most promising account available. I simply take issue with Nickel's account of normality.

[24] Nickel calls it "genericity", following much of the literature on generics, but as far as I can tell normality and genericity are the same phenomenon.

[25] Strevens (2012) also explores causal mechanisms as underwriters of defeasible regularities. In that respect, his discussion is similar to Nickel's: both see the defeasibility of causal mechanisms as the source of the defeasibility of certain types of normalities.

on our ability to specify which sorts of explanations can ground normalities, and there doesn't appear to be a way of doing this without introducing circularity into the account.

Nickel wants to give a semantics for bare plural generic sentences—sentences like 'dogs have four legs' and 'male platypuses have venomous barbs' and 'elephants live in Asia'—as existentials over ways of being normal in a given respect. So 'dogs have four legs' gets the semantic interpretation: "there exists a way of being a normal dog with respect to limb count and that way of being normal is having four legs." The *respect* here is something like "limb count", which is the dimension or determinable a determinant of which is the relevant way of being normal. A *way of being normal* is one of potentially many contrary properties a member of a kind can exhibit relative to a respect (see footnote 23). So far, I don't take issue with Nickel. To the contrary, I find his treatment of generics quite compelling and insightful.

It's Nickel's account of normality that's at issue. Obviously, Nickel owes us a definition of normality if he is to hang all of this semantic weight on the notion of a "way of being normal." His account has a few moving parts. He first defines characteristic properties (which are, I take it, normal properties for a kind) by appeal to explanatory strategies. An explanatory strategy is, on first pass, a set of explanations which appeal to one coherent set of explanatory factors. Evolutionary explanations, for instance, appeal exclusively to selective pressures, adaptivity, differential reproductive success, and so on. So an evolutionary explanatory strategy is one which appeals to only these factors.[26]

A characteristic property of a kind relative to a set of explanatory properties *S* is a property the presence of which can be explained by appeal to *S* (or a subset of *S*). The presence of croaking

[26] See (Nickel 2016, 178-196) for a full discussion of the definition of characteristicness in terms of explanatory strategies.

among the population of bull frogs can be explained by appeal to evolutionary explanatory strategies and so croaking is characteristic of bullfrogs (normal for bullfrogs).

What does "presence" mean? Nickel (2016, 182) precisifies the concept to "*P* is present among *K*s at [time] *t* iff at some time surrounding *t,* some *K*s are in fact *P*." The temporal extension into the future and past of "some time" is relative to the explanatory strategies being evoked. Evolutionary explanations, for example, will require long temporal extensions. Geological explanations will require much longer timespans still.

Now how do we bridge the gap between the characteristic properties of a kind and normal members of that kind? Nickel does so by appeal to causal mechanisms. Normal dogs go through a developmental process (a causal mechanism on Nickel's account) which produces exactly four legs and that developmental process is the explanation for the presence of four-leggedness in the temporally extended population of actual dogs. Dogs with four legs are normal because they participate in the causal mechanisms cited in the normality-grounding explanatory strategies employed in the explanation of the presence of four-leggedness in the population of dogs. Hence abnormal properties, on Nickel's account, are those which cannot be explained by the right kinds of explanatory strategies—he puts the analytic weight on the fulcrum of explanatory strategy selection.

How, then, does Nickel propose to restrict the sorts of explanatory-strategies that can ground the normality of a property? He supposes that in some cases a set of explanations instantiating an explanatory strategy is manifestly coherent. Perhaps evolutionary explanations are clearly coherent in that they appeal, as noted earlier, to the same sorts of explanatory factors. Other times, when it is less clear how many explanations cohere into one explanatory strategy, Nickel suggests that those implicitly appealing to explanatory strategies of which they are ignorant are also implicitly deferring to the experts—recalling Kripke (1980). More concretely, when I say that it's normal for Elephants

to live in Africa, I do not know the right explanation and so I am deferring to paleontologists, evolutionary biologists, and other experts.

Nickel's account is, near as I can tell, the most developed representative of a class of accounts of normality that I'll call Presence Accounts. They seek to define normal properties[27] in terms of the possibility of explaining (by appeal to the sanctioned explanatory strategies) the presence of some property or feature in the population of members of the relevant kind.

My account will alternatively seek to define normal properties in terms of the ability of the kind itself to do explanatory work—in terms of the explanatory link a kind has with its normal properties. This is why I call the class of accounts to which mine belongs the Link accounts. So, whereas Presence accounts want to say that it's normal for dogs to have four legs because we can explain the presence of four-leggedness in the population of dogs, Link accounts hold that it's normal for dogs to have four legs because being a dog explains why any given dog has four legs. A link account will put the analytical weight not on explanatory strategies, but on the specification of the individuals involved—specify the individual at the right level and you'll have an explanation for why that individual should have any of the corresponding normal properties.

I want to discuss two central worries for a presence account: 1) it doesn't seem to work as a simple analysis of 'normal' at first glance, whereas a link account seemingly will; and 2) it relies on a restriction of admissible explanatory strategies, which Nickel thinks can be accomplished by deference. The problem with such a move is that the deference targets (for instance, the biologists) must be devoid of tacit or explicit employment of non-statistical normality and it's at least not clear

[27] It's worth noting another difference between our accounts. Whereas Nickel is primarily interested in defining what it means to be a normal *k* with respect to a property (P-Normality), I am more interested in defining the normal properties or features. This doesn't amount to a substantive difference, but merely a difference in emphasis.

that they are.[28] If we have to choose, therefore, between link and presence accounts, I am arguing that we choose a link account.

A presence account will say that the normal is that for the presence of which we can explain by appeal to the right explanatory strategies. The worry is that it doesn't seem right to say, for example, "biology explains why cats stalk, which is to say, it's normal for cats to stalk." It seems like, on its face, this can't provide an analysis of what it means to say a property is normal for a kind. Conversely, a link account will provide an analysis that is more satisfying: "Its black because it's a raven, which is to say, it's normal for ravens to be black." Now, neither of us needs to provide an analysis in the classical sense of the concept of normality, and indeed neither of us is trying to do this. Nevertheless, it seems that an account of normality which sounds better as a *prima facie* analysis of normality seems to be to be better for it.

The deeper worry is that a presence account puts the analytic burden on the process of restricting explanatory strategies in the right way, but this process is not clearly up to the task without introducing circularity into the account. For instance, we intuitively need to disallow as a normality-grounding explanation the explanation for albinism in Ravens. Albinism causes Raven feathers to be brittle and sticky rather than durable and waxy. Albino ravens therefore cannot sustain as long of flight and are more sensitive to sunlight—both of which are antithetical to the raven form of life.[29] Nickel argues that the explanation of albinism in ravens grounds a derivative normality judgment to the effect that *albino ravens are white*. But since the explanation applies only to albino ravens—a derivative subgroup—the judgment that *ravens are white* is blocked. In talking about ravens, we tacitly defer to experts—biologists in this case—and this deference restricts the set of relevant explanatory strategies to ones concerning ravens in general.

---

[28] This does constitute Nickel's metaphysical account of *genericity*, which one wouldn't be entirely mistaken to think is—if not the same as—at least closely related to non-statistical normality.

[29] Cf. (Huizen 2015).

First, it isn't clear what makes different morphs of a polymorphism (a normal variation) distinct from black and albino ravens apart from notions tightly linked to non-statistical normality—adaptation, naturalness, defect, etc. Each is caused by a genetic variation that produces a different phenotype.[30] Either biologists, therefore, are using a conception of non-statistical normality to distinguish between normal variations and mutations proper like albinism, or they are using some other method for distinguishing them. Herein lies the problem: if biologists don't uniformly use methods other than a tacit appeal to non-statistical normality, naturalness, and the like to distinguish between polymorphisms and cases like albinism, then Nickel's presence account employs non-statistical normality—by way of appealing to deference to biologists—in giving an account of non-statistical normality. This is a serious problem, indeed.

Here is why it is plausible to suppose that biologists make appeal—explicit or tacitly—to non-statistical normality in distinguishing normal variations from mutations like albinism. The first thing to note is something Nickel himself points out: that normal variations are adaptive or at least not maladaptive whereas albinism is maladaptive and thus the symmetry between the developmental mechanisms can be broken by appeal to the ultimate, adaptive explanation. Adaptivity claims about features are claims that an organism is better off for having those features with respect to a set of environmental pressures and affordances. They're claims to the effect that the *kind's* relation to its environment (or perhaps more properly to its *niche*) sets standards which members of the kind meet to varying degrees. Since members of a kind don't all share the traits that are adaptive for that kind, but it nevertheless remains adaptive, adaptivity claims have the hallmark inferential profile of normality claims. On this basis, I'm claiming that adaptivity claims are a species of normality claim.

[30] N.B. that some polymorphisms are caused by environmental cues alone, but we can safely ignore these cases and restrict our focus on the genetically-determined polymorphisms.

Second, the task before the biologist is to distinguish cases of normal variation from abnormal variations in the same respect. The problem is that they must do so without appeal to normality. What does this entail? Well, any theorist of normality will agree that 'normal' in the non-statistical sense describes properties which characterize a kind *independently* of the statistical prevalence of that property in the population of actual members of that kind.[31] So, any principle that biologists appeal to in distinguishing between explanations that ground normal variation and those that ground abnormalities like albinism will have to make appeal to either universal laws governing the relevant biological kind or statistical generalizations about the prevalence of the properties in the population. Since biology famously doesn't conform to universal laws, we're stuck with statistical generalizations. I'm pessimistic that a distinction between black and albino ravens can be drawn on purely statistical grounds, simply since it's conceptually possible that albinism runs rampant, thus characterizing a majority of ravens. If we're pessimistic that this can be done in biology, even by appeal to partitioning and the like,[32] then we're in even hotter water when it comes to artifacts, where standards and functions are built into the kind in a more obvious and intimate way.

This is all salubrious for the explanatory link account, which thinks that the explanation of the presence of normal properties is disconnected in important ways from what it is to be a normal property. A link account holds that normality is logically prior to the actual population of members of the kind. In a slogan: no kinds without normality.[33] Nickel thinks kinds are out there and normality is grounded in explanations attaching to those kinds. This, I think, spells trouble for his account.

---

[31] Assent to this seems to be a necessary precondition to theorizing about non-statistical normality and so is perfectly neutral between accounts of normality. Either way, I think at least Nickel would agree with this minimal characterization.

[32] See Cohen (1999) for a sophisticated attempt.

[33] The converse is also true: there is no normality if there are no kinds. I think most would agree with this assessment.

On final analysis, then Nickel needs to do more to show that the experts that we're supposed to defer to in picking out normality-grounding explanations aren't themselves contaminated with normality. Regardless of the success of this enterprise, link accounts sidestep the whole problem and so are better for it.[34]

4. The Explanatory Account of Normality

The explanatory link account of normality relies on the intuition that when a normal outcome obtains, it is *explained* by the kind of thing for which it obtains. Dispositions and functions and the like are posited in part *in order to explain* the outcomes that regularly occur for their possessors, so of course whenever a normal outcome obtains (in the normal way), it will be explained by the disposition or function or the like. Furthermore, citing a disposition or function or the like is often *typing* the possessor of that disposition or function or the like by the kind for which that disposition or function or the like is characteristic if not constitutive. In the case of a normal outcome, therefore, we need only cite the kind of thing in question in order to explain the outcome. Why did that thing turn on when that button was pressed? Well, that's a computer and that's its 'on' button (part of being a computer is being disposed to spring into action at the pressing of one's 'on' button). Why does that thing have four legs? Well, it's a dog and dogs have four legs. Why is he here at 2pm? That's Dev, he's normally in his office at 2pm[35] (Dev's diachronic dispositional profile—part of his nature—explains his being in the office at 2pm). Why does this muscle group constrict like that? That's a heart, and hearts (constitutively) constrict rhythmically in order to pump blood.

[34] More broadly, I'm tempted to think that our conception of normality should illuminate what biologists are up to rather than the other way around, but one can see the merits in thinking the reverse. Note that biology will be one way we answer questions about what is normal for a given species, and why a property is normal for a given species, but I don't expect it to also play a role in our conception of what it is to be normal.

[35] It is interesting to note that we often use phrases like "he's always in his office at 2pm" to express a normality judgment. We clearly, when pressed, don't mean such statements as universal statements—even with a suitably restricted domain that rules out weekends and work holidays—since we are likely well aware that Dev is sometimes sick or on vacation or at the birth of his child on workdays.

Kind explanation—the citing of a kind or the (modally and temporally extended) nature of an individual as the *explaniens*—is closely related to Aristotle's "formal explanation," where citing the kind of thing—the same as citing the thing's essence—is used to explain a given feature of the thing.[36] To get onto the intuition behind formal explanation, it's best to think about the opposition between formal explanation and material explanations. Certain facts about me obtain because I am a human—being a human explains why I am the way I am. Other facts, though, obtain because I am a human with a particular body with a particular history and with particular abnormalities and oddities. That I have a scar in my left eyebrow isn't explained by my being a human, it's instead explained by my having run face-first into the corner of a table in kindergarten—a particular fact about my history.[37] The fact that a tree has one ring that is significantly thinner than the rest isn't explained by its being a tree, it's explained by its having insufficient material resources during one year of its development. Seeing that there is a contrast here helps us to get onto the intuition that there's a form of explanation that appeals as an *explaniens* to the kind of thing involved.

But the claim here isn't that only *essential* features are explained by the kind of thing involved (or the nature of the individual). Instead the claim is more liberal—that *normal* features are the features explained by the kind of thing involved. Some of these "*per se*" or principled features are *necessary*: some (the constitutive) are what it means to be what that sort of thing is as is for Aristotle being sensitive to an animal; others (the *propria* or proper in the sense of exclusive possession) are not constitutive of being that sort of thing, but nevertheless are always found in subjects of that kind and are only found in subjects of that kind as is for Aristotle the feature of having angles equal to

[36] Aristotle offers the somewhat unhelpful example of citing the fact that an octave in music is a numerical relation in order to explain certain facts about octaves. This is far from the most helpful example available (*Physics* II 3, 194b24).

[37] Furthermore, certain facts about me obtain because I am *who* I am and not because of some accidental fact about me—like my running into a table and scarring my brow. For instance, it seems true that I pursued a PhD in Philosophy *because that's the kind of person I am* or because I am Andrew Lavin and not someone else.

180°; still others (the generic) are necessary preconditions of being a member of that kind—and so are necessary but neither *propria* nor constitutive—as is being physically extended for being an animal.[38] There are also non-necessary but 'natural' or 'proper' (now in the evaluative sense of the word) features of members of a kind which are still had in virtue of being a member of the kind and so still *per se* features. An example of the natural is that of being a friend for humans in that the fullest and best human lives are ones in which one has friends, but one can go one's whole life without having a friend without threat to one's status as a human being (though perhaps Aristotle will question one's *humanity* or status as a full member of the species). Still others are merely *unique* to the kind for which they are *per se*. Humans are the only things (that we know of) that can find treasure[39] or have wine cellars.[40] All of these *per se* features, activities, and conditions are normal on the account under analysis here. The view therefore might be characterized as the view that the normal outcomes are the non-accidental or principled outcomes.

What to make of 'explain' here, though, when this word is now famously polysemous at best and hopelessly ambiguous at worst? For our purposes, we can get a lot of latitude out of the concept by staying at a fairly naïve level. Explanations are or are expressed by answers to 'why' questions, which themselves express or are calls for explanations. For one to give an explanation, therefore, is for one to answer an implicit or explicit 'why' question—which will determine which kind of explanation is appropriate, perhaps by explicitly or implicitly furnishing the contrast class for the *explanandum*.[41] So explanation has a pragmatic element, but perhaps is something over and above a pragmatic

---

[38] How can a *necessary* feature be nevertheless *normal*? Well, remember that the distinction between necessities and normalities *qua* types of relations between kinds and features.

[39] Is a racoon stumbling upon a squirrel's nut cache a racoon *finding treasure*? I'm not so sure. What I am sure of is that if I say "treasure was discovered in my back yard" you won't immediately think of squirrels, but will instead normally think of human contexts.

[40] Both examples and the overall thought thanks to Gavin Lawrence.

[41] Cf. for example, Van Fraassen (1977, 1980) and Lewis (1986) for a discussion of the contrastive nature of causation and therefore likely of explanation in general.

process.[42] Because here we're interested in understanding the mystery of normality by appeal to concepts that we understand better, it doesn't profit us to get bogged down in the mysteries surrounding the concept of explanation if we have an intuitively accessible naïve account of what explanation is (so long as our naïve account isn't incoherent). That is to say: I take it that we already have an intuitive understanding of what it means to give an explanation of some phenomenon (beyond the pragmatics of answering 'why' questions) and I take it furthermore that this understanding is more accessible to us than is our intuitive understanding of normality.

Having as we now do some motivation for the explanatory link account and some intuitive idea of how it might work, let us begin with an initial account and build toward something stricter:

> (Explanatory Link$_1$) Normal outcomes are those which are explained by the kind of thing for which they obtain.

Some features of a thing are had *insofar as* it is a certain kind of thing, while others are had insofar as it is a different kind of thing. As a human, it's incidental that one fail to aspirate one's H's. As a New Yorker, though, it's normal or part of the kind in question—i.e. non-accidental.

This account faces an immediate problem: some explanations are "indirect." My being a human in some sense explains a runner tripping over me. Humans are solid objects and so my being a human as opposed to being a puddle explains why that runner who ran into me as I lay reading in a field *tripped* rather than *got wet*. But we wouldn't want to claim that it's normal for humans to trip runners. It might not be abnormal, but it is intuitively too strong to say that it's *normal*.[43] With this sort of case in mind, let us posit the distinction between *direct* and *indirect* explanation.

---

[42] Whereas Van Fraassen (1980) thinks that an explanation is an answer to a why-question and that's all, I am skeptical of such a view. Explanations are answers to why-questions, but perhaps they are other things as well (relations between facts, for instance).

[43] There seem to be many examples of properties which are neither normal nor abnormal. Cambridge properties seem mostly to be neutral with respect to normality, as is the property of painting one's skin green—it's not a normal thing to do, but it's too strong to say that it's *abnormal*.

Indirect kind explanation is explanation that cites a kind (or a specification of an entity) as the *explaniens* but which has explanatory force *in virtue of* that kind bearing some connection to another kind—the kind which by itself would serve as an *explaniens*. Whatever explanatory force there is in saying that I tripped the runner because I'm a human being is had in virtue of the fact that human beings are solid objects and not liquid or gaseous objects. A direct kind explanation, alternatively, is one that cites the kind that gives the explanation its explanatory force. With this distinction in mind, let us make the following revision to our account:

> (Explanatory Link$_2$) Normal outcomes are those which are directly explained by the kind of thing for which they obtain.

In adding the condition that the explanation in question be direct, we avoid cases where the connection between the kind of thing in question and the outcome is incidental.

One might worry at this point about the following kinds of cases:

> (Adoption) Katie really likes dogs and wants to adopt a dog. She adopts Leeroy. Leeroy's being adopted by Katie is explained directly by Leeroy's being a dog. After all, Katie was set on adopting a dog rather than a cat or mongoose.

This is at least an apparent problem because the explanatory link account seems to entail that it's normal for Leeroy to get adopted by Katie.[44] "Leeroy normally gets adopted by Katie" sounds even stranger to our ears. Intuitively, there's no regularity here that could be a normality—a connection between some general entity and a feature—and so it can't be normal for Leeroy to get adopted by Katie. Part of the strangeness is that it's a relation between individuals and so appears to be at best a one-off case instead of a regularity and it seems like normalities have to be at least potentially regularities.

There's a lot to say, so it will be best to say each quickly. First, we need to get clear on the contrasts implicit in the context. "Leeroy getting adopted by Katie rather than Missy the cat getting

[44] Thanks to Katrina Elliot for these worries.

adopted by Katie" is one plausible contrast, and "Leeroy getting adopted by Katie rather than Kitty the cat lover" is another potential contrast. Leeroy's being a dog, it seems explains the *explanandum* in both cases. However, the question might have been "why did Leeroy the Labrador rather than Rufus the Rottweiler get adopted by Katie?" to which the response "because he's a dog" sounds otiose. We'll have to appeal to Leeroy's special or individual qualities to answer the question.

Second, we must pay careful attention to the work that tense does here. It might sound odd to say, "it's normal for Leeroy to get adopted by Katie" but it won't sound as odd to say, "it *was* normal for Leeroy to get adopted by Katie." This doesn't quite alleviate the worry that one-off events can't really be normalities, but it gets us some way towards understanding what is going on in this case.

Finally, it seems that perhaps the specification of the beings in question might have gone wrong. Is it normal that Leeroy got adopted by Katie or is it normal that Leeroy got adopted by a dog lover? Is it normal that Katie adopted Leeroy or is it normal that Katie adopted a dog? Again, contrasts are important to track. If the question is about why Leeroy was adopted by Katie rather than Missy the cat, then the answer will be "because Leeroy's a dog." If the question is about why Leeroy was adopted by Katie rather than Kitty the cat lover, then Katie's being a dog lover is relevant to answering the question. "It's normal for Katie to adopt dogs" seems like a perfectly good judgment and a true one at that—especially if we subjunctivize it to "it would be normal for Katie to adopt a dog." "It's normal for dogs to get adopted by dog lovers" is certainly just as good. "It'd be normal for Leeroy to be adopted by someone who is attracted to his upbeat energy and peerless loyalty" also rings true.

It's less clear, this far into the analysis, that "it's normal for Leeroy to be adopted by Katie" is a bullet that needs to be bitten by the Explanatory Link theorist because the normality in question seems not to be between individuals (or an individual and a property which makes reference to another individual) but between an individual and a general entity akin to a kind. Once we've gotten

clear on the contrasts, we'll have abstracted away from an individual-individual relationship to something that looks more like an individual-kind relationship and that's the sort of thing that can uncontroversially be called a normality. The individual-individual explanations, it seems, are indirect explanations.

Now that the basic view is on the table, it is time to test whether it accounts for the core features of normality judgments. Let's rehearse those features:

> (Tolerance) Normality judgments all tolerate (remain true in spite of) exceptions to the corresponding universally quantified judgment. That is, they don't have straightforward counterexamples.
>
> (Insensitivity) Normality judgments are all *relatively* insensitive to changes in the prevalence of the relevant outcome in the population over which they generalize.
>
> (Normativity)(?) Normality judgments entail that any individual over which they generalize, insofar as it is a member of the relevant kind, *should* have the outcome being attributed.

The explanatory account of normality judgments tells us why they should be tolerant: if the relevant feature *would* be explained by the kind *were* it to obtain, then there's simply no entailment that it should obtain in *any* particular circumstance (or potentially any *proportion* of particular circumstances). We can also account for why they should be insensitive: the claim is that there's an independent explanation for any apparent exception and there don't appear to be any principles governing how many exceptions can be so explained. So long as there's an independent explanation, the normality judgment is safe. This is what we'd expect pretheoretically given commonly held intuitions about the lax relationship between normality and prevalence.

Some normality claims also seem to be *normative* in inferential structure: at least certain normative generalizations entail *should* claims for the individuals over which they generalize. The explanatory account of normality accounts for this, but in a roundabout way. On this account, having some feature is normal for one insofar as one is an *f* if and only if having that feature will be explained by one's being an *f*. If one fails to have a feature that *would* be explained by one's being an *f*, then it's

natural to think that being an *f* won't explain one's *not* having that feature (though this isn't an entailment). Given that we need an explanation for such an outcome, we'll have to look elsewhere from the standard "x is an *f*" explanation normally available to us. We'll have to look at *what went wrong*.

Now, the shift from explanation to something going wrong isn't yet warranted. Why not simply think that we won't appeal to the kind *f*, but we can appeal to a different kind *k*? Or, alternatively, why not think that we can simply point to something that was *irregular* or *uncommon* to explain the abnormal outcome? Obviously, there is a lot of work to do here. Quickly, though, in the context of an explanation, things are specified according as they help bring about some state of affairs. And, insofar as they are so specified, it is a *bad* for them not to bring about that state of affairs. A steppingstone *qua* rock doesn't "care" how stable it is, but *qua* steppingstone it is important to its success as a steppingstone that it be stable. Something will have gone wrong at some level if the normal outcome fails to obtain because in our explanatory practices, we'll be looking for what it was that—at least analogously—went *wrong*.[45]

In addition to offering a successful account of functionals, trajectory claims, and characterizing claims, the Explanatory Link Account can also deal with less central cases. A habitual is a normality judgment on this account because Dev's work habits *explain* why he's in fact in his office at 2pm, but does not *entail* that he will be in his office at 2pm. Moreover, a disposition *explains* a particular outcome—such as a computer powering on—when that outcome occurs, but again does not entail that the disposed computer will *in fact* do such a thing.

[45] Some normalities—"complete normalities"—are such that they will obtain absent interference, while others—"incomplete normalities"—are such that if they obtain they will be normal, but there is no reason to think they will in fact obtain unless we have more information than the kind for which they are normal. The normative inferential profile can in principle only apply to complete normalities. There is no normative pressure toward the obtaining of an incomplete normality: a cat can be calico, black, white, or orange tabby without anything going wrong. In fact, those are *all* normal outcomes. So to make an inference to a 'should' claim about an individual from an incomplete normality judgment would be an error in reasoning.

We thus have an account that appears to cover the focal cases we set out to account for and has the potential to cover many more types of normality judgments. With that in mind, we've got a *prima facie* successful account of the content of normality judgments.

Chapter Four: The Explanatory Role of Normality, and Worries

This chapter will expand on the foundation laid in chapter 3 by a) exploring in more detail the role of normality in our explanatory practices in section 1 and b) responding to worries we might have about the explanatory link account of normality in section 2. We'll address them in this order because it will be necessary to gain a fuller understanding of the explanatory link account before attempting to respond to a few pressing objections to the account.

1. Further down the rabbit hole

The explanatory link account of normality takes explanation to be central to a correct understanding of normality. With that in mind, it will be fruitful to explore the role of normality in explanatory practice in the hopes of not only deepening our understanding of normality, but also of discovering more about explanatory practice. Section 1.1 will seek a better understanding of the pragmatics of explanation in normal cases, 1.2 will incorporate an insight from Pietroski and Rey about the role normality plays in explanatory practice, and 1.3 reflects more broadly on these themes.

1.1 No Explanation Called For

Some intuitive motivation for the explanatory link account can be found in a pragmatic phenomenon: we find ourselves puzzled or at a loss for words at the asking of certain explanatory questions. The explanatory link account of normality can shed light on this phenomenon, and, in turn, this phenomenon can serve as a heuristic marker of the normal.

In many cases, no explanation is called for—whether as a normative fact about how explanatory practices should go or as a pragmatic fact about conversations involving calls for and offerings of

explanations.[1] If someone with fast reflexes plays shortstop, we don't ask for a further explanation unless we don't understand the connection between the position and having fast reflexes. Once we understand this connection, though, we won't ask for further explanation.[2] If we come across a moose, we won't ask why it has four legs[3] even though we would likely ask why a three-legged moose has three legs. If we encounter an adult human being performing basic arithmetic, we won't ask how that individual is able to do such a thing.[4] In each case, it seems that—even though there are many different kinds of explanation for the outcome in question—no explanation is called for.

Questions such as the following, in the usual context (I've included examples of plausible contrasts in parentheses), are at least odd if not inappropriate and either way are often difficult to interpret:

(1) Why does that particular zebra have black and white stripes (rather than pink polka dots or being completely black or completely white)?

(2) Why is the head of that hammer made of metal (rather than gelatin)?

(3) Why did that walnut grow into a walnut tree (rather than an almond tree)?

Call the phenomenon illustrated by 1-3 the No Explanation Called For phenomenon. In each case, we are presented with a call for an explanation that, in the usual context, is either a mistake or reveals misunderstanding or ignorance. They are calls for explanations that don't seem to be

---

[1] Some—Bas Van Fraasen (1977; 1980) for instance—would argue that this is a distinction without a difference. I'm not convinced that this is accurate.

[2] Of course, there's always a scenario available in which we *would* ask for an explanation. If Tamar not only has fast reflexes but also throws a 90mph fastball, then we'll want an explanation of why she should play shortstop rather than pitcher. But absent such further information, and understanding the demands of the position of shortstop, we won't demand an explanation beyond the facts already given: Tamar has fast reflexes and so Tamar is chosen to play shortstop.

[3] Notice how there was nothing odd in this sentence's assuming that the moose had four legs. One aspect of normality seems to be: whatever "comes along for free" when we've introduced a subject of predication into a sentence. When I say "Harold is a moose" it seems that "Harold has four legs" comes for free (defeasibly of course). This inference pattern only holds for *complete* normality judgments where there is no rival contrary normality for the kind.

[4] As will become clear, we can always ask questions about why principles hold: "why do people who play shortstop have to have fast reflexes?" or "why do mooses have four legs?" or "how are humans able to engage in mathematical reasoning?"

required. Zebras as a species have black and white stripes, so asking why a particular zebra is so colored is like asking why 2+5=7. It's the nature of numbers and addition in general that explains why 7 completes the equation rather than 8, just as it is the species zebra in general that explains why Steve the zebra has stripes.[5] Asking about the individual case seems to demand an explanation at the individual level but the most natural answer is a general one: it's a zebra and zebras have stripes. The phenomenon in question is simply our puzzlement at such questions. After developing the explanatory link account, we'll gain new insights into the causes of our confusion.

Why should this phenomenon count in favor of something like the explanatory link account? The answer is found in the *cause* for the no explanation called for phenomenon itself. One plausible account has it that we are puzzled at the call for explanation because the most pragmatically appropriate explanation is *already supposed to be in the common ground*. Asking a question is a way of requesting that new information be brought into the common ground. So, if one asks a question the answer to which is already in the common ground, one is doing something pragmatically inappropriate: requesting that the common ground be updated with new information when the "new" information is already in the common ground. That Steve is a zebra explains why Steve has black and white stripes. Full stop. There are other explanations along different dimensions such as ontogeny, heredity, adaptation, and the like; but the fact that Steve is a zebra does all of the explanatory work needed to satisfy a call for explanation with a contrast such as those offered in (1).

This isn't the only possible explanation for the No Explanation Called For phenomenon, but it is the most plausible, and therefore it is sensible to hang analytic weight thereon. To finish the analysis: if the No Explanation Called For phenomenon is explained by a failure to recognize a fact in the common ground as an explanation for the proposed *explanandum*, then there must be some fact in the common ground that in fact explains the *explanandum*. The most readily available fact

[5] I am thankful to Gavin Lawrence for this point and this illustration.

available in each of these cases is the specification of the individual to which the relevant outcome attaches. It is therefore reasonable to suppose that the specification of the individual is the fact that does the explanatory work needed. There's no entailment here, but if the argument is seen along the lines of an inference to the best explanation, then the No Explanation Called For phenomenon serves as evidence for the supposition that kinds themselves do explanatory work.

Let us explore a few dimensions of the no explanation called for phenomenon and their connection to the explanatory link account before moving on.

What Question is being asked? Kinds and Levels

Consider the following plausible dialogue:

(4) Ahmed: Why is Seabiscuit running that way?
Fatima: What do you mean? That's how horses run.

The simplest thought of what is happening here is simply that the question Ahmed is asking implicitly differs from its explicit content. So perhaps the problem is *semantic*. Instead of asking why does *Seabiscuit* (the individual) run that way, Ahmed might be inquiring into the normal gait for horses and how it came to be normal for horses. So he might be asking "Why do horses run that way?" This would explain why Fatima's response is the most natural one given that Seabiscuit is running normally for a horse. Fatima's question highlights the mismatch between the explicit and implicit content of Ahmed's question.

If the question really is "why do horses run that way?" then Fatima's response merely serves to clarify the question rather than to answer it. The answer to the question might take one of a few different forms. It might appeal to the adaptive nature of the equine gait, or it might point to the lineage history that preceded the gait and the incremental—perhaps adaptive, perhaps merely accidental—changes which led to the current gait.[6] It also might do both: it came to be this way

[6] See Calcott (2009) for a discussion of lineage explanations.

through incremental changes largely owing to the structure of horse limbs and the like but also in part because it is adaptive given a certain set of constraints.[7]

Alternatively, we might think that what Ahmed is in search of is in fact an explanation of why the individual, Seabiscuit, is running the way he is. The answer here will necessarily appeal to Seabiscuit's biomechanical structure, recent history, etc., rather than the adaptivity of so running. Individual horses do not run a certain way *because it's adaptive* since facts about adaptivity are relations between species and their environments rather than individuals and their environments.[8] That is to say, the explanation on offer for an individual fact will have to be an *accidental* explanation unless it's simply "because Seabiscuit is a horse and horses run that way," which is an appeal to *principle* or the species-level. We can then further flesh out the picture by giving a principled (i.e. non-accidental) explanation of why horses run the way they do.

The upshot of this discussion is that normality judgments are paradigmatically judgments about a *kind* rather than an individual member of a kind.[9] To say "Horses walk with a four-beat gait" is to say: given that Sam is a horse, Sam's walking with a four-beat gait "comes for free" or is included—defeasibly—in the specification. In learning that Sam is a horse, we have thereby learned (defeasibly) that Sam walks with a four-beat gait. The No Explanation Needed phenomenon appears to arise in part because the call for explanation both demands we answer with facts about a kind and appears to demand that we answer with (special) facts about an individual. It appears we are being asked for a "special" explanation when no such explanation is in fact needed.

---

[7] Now horses are a domesticated species, so the story might be more complicated than that, but to the extent that their gait matches that of Zebras or the like, we'd expect adaptation to similar ecological constraints and similar biomechanical constraints to be part of the story.

[8] I take it this is relatively uncontroversial.

[9] N.B. that habituals and the like are about individuals, but we simply repeat the schema at the individual level: it's about an individual *through time* rather than *at a time*. The whole individual-time slice relation is in some ways analogous to a kind-member relation.

Common Ground and Stock Explanation

Another approach to this question is to look at what it means for an explanation to be called for or required—to look at *pragmatics* rather than semantics. What does the "called for" in "No Explanation Called For" mean?

There might be another (perhaps compatible) explanation for the Ahmed and Fatima case: Ahmed's question might sound strange because there's a stock explanation available in the common ground and thus his question sounds like it's implying that Seabiscuit runs abnormally because to ask such a question would otherwise flout a Gricean maxim when the stock explanation is ready to hand. It might be that no explanation is needed and so to call for one inadvertently (or for that matter advertently) implicates abnormality. To unpack this idea a bit, it might be the case that Fatima is assuming that Ahmed already knows that horses normally run that way and so Ahmed's question seems odd. After all, Ahmed's question, being as it is about an individual rather than about the kind, only seems to make sense against a background of understanding that Seabiscuit is a member of a kind the members of which normally run a particular way. If it truly is common ground, then Ahmed's question flouts what might be called the maxim of salience—the maxim that one should not make salient what is in the common ground unless necessary—in that it asks a question the answer to which is available in the common ground. Alternatively, we might simply think this it is the Maxim of Quantity being flouted in that the question is more specific than necessary or of Relation in that the relevant facts aren't being explicitly included in the question. Either way, such flouting in this case produces the implicature that Seabiscuit is running abnormally for a horse, which is false. Hence Fatima's confusion and perhaps exasperation.

In other words, Fatima and Ahmed's exchange may be analogous to the following:

> Sam: I know squirrels are normally grey, but why is that squirrel grey?
> Elika: Ummm… because her parents were grey? I'm not sure what you're asking.

In this case, Sam makes explicit that the stock explanation is common ground and then proceeds to ask the question anyways. The implicature seems to be that the squirrel serving as the subject of Sam's inquiry is abnormal, but Elika doesn't have a clear route to understanding in what sense it could be abnormal. Is it that the squirrel's parents are brown or red or white, and so it is abnormal for it to turn out to have grey fur? Or perhaps there's reason to think that this squirrel is gray for a nonstandard reason? Sam's question is thoroughly confusing in perhaps the same way that Ahmed's question is confusing: each seems to ask for something available in the common ground to be made explicit and in so doing seems to implicate that a normal case is abnormal.

"Special Explanations"

We're now in a position to bring the explanatory account of normality into the discussion. The explanatory account holds that the normal for kind *k* is what, when it obtains for an individual *k*, is explained by that thing's being a *k*. If there's no explanation called for in certain cases—especially if these are cases involving normalities—then it will be tempting to think that there's no explanation called for *because we already have an explanation*. Intuitively, all things that can be explained call for explanation in some wider sense. If there is no explanation in our science or theory, then in some sense there should be (provided it falls within the domain of the science or theory). If this is true, then if something doesn't call for explanation in the pragmatic sense we've been discussing, then it would have to be because we already have an explanation for it. This is just what the explanatory account claims: when there's no explanation called for in the case of a normal property instantiation or the like, it's because the kind of thing in question—which feels in our examples like it's already in the common ground—does all the explanatory work we need.

If I'm going to hang analytic weight on the notion of a 'special explanation,' then I owe an analysis of what this means. In its simplest formulation, a special explanation is one which isn't already in the common ground. There is, however, a bit more to say.

Let's start with a particular insight about the kinds of explanation required by different cases and build our way towards an understanding of what a "special" explanation is. First, consider the following cases:

(5) Lottery1: Mohinder wins the raffle prize in a random ticket drawing raffle.

(6) Lottery2: Marisol wins the raffle prize in a random ticket drawing raffle.

(7) Lottery3: Nobody wins the raffle prize in a random ticket drawing raffle.

Note that the explanations for $Lottery_1$ and $Lottery_2$ can be structurally identical:

(8) $LotteryExplanation_1$: In the random ticket drawing, Mohinder's ticket was drawn.

(9) $LotteryExplanation_2$: In the random ticket drawing, Marisol's ticket was drawn.

And the explanation for $Lottery_3$ must be structurally distinct:

(10) LotteryExplanation3: A ticket was accidentally put into the tumbler the pair of which was not purchased.[10]

The same can be found in cases where the outcome isn't probabilistic:

(11) $Lock_1$: When I hit the lock button on my wireless dongle, the car locks.

(12) Lock2: When I hit the panic button on my wireless dongle, the car alarm is tripped.

(13) Lock3: When I hit the unlock button on my wireless dongle, the car locks.

Now look at the most natural explanations on offer:

(14) LockExplanation1: The dongle and car security system were designed and built so that the lock button would lock the car.

(15) LockExplanation2: The dongle and car security system were designed and built so that the panic button would trip the alarm system.

[10] Why not explain it in the following way? In the random ticket drawing, a ticket belonging to no one was drawn. This isn't clearly structurally distinct from the others. It also doesn't cite the important fact in explaining the event: what happened that lead to nobody winning? $LotteryExplanation_3$ does this in that it cites the event which made our candidate explanation here true. Intuitively, we wouldn't expect this explanation to satisfy the inquiry whereas we likely would expect $LotteryExplanation_3$ to satisfy it. $LotteryExplanation_3$ itself calls for explanation, whereas the other two do not.

(16) LockExplanation3: The dongle or car security system has a crossed wire which is corrupting the designed structure of the circuit.

Again, the intuitively satisfying explanation for the abnormal outcome is structurally distinct from the explanations for the normal outcomes, which are all structurally similar. No doubt there is more to say on this and more work to be done in terms of precisifying the claim being made here, but this will suffice for our purposes here. The upshot of such observations is that abnormal instances seem to require *structurally distinct* explanations from normal instances, whereas all normal instances will perhaps be structurally identical. Now we of course have yet to *demonstrate* this principle, but we've at least given it some initial motivation.

Anscombe's concept of a "special context,"[11] as I read it, is alternatively meant to be relatively intuitive. The idea seems to be that there is a "usual context" of utterance for most if not all sentences such that if I present you with the sentence out of context, you will implicitly call to mind that context in interpreting the sentence. So, if I present you with the following sentence:

(17) The door hinge squeaks.

You'll (likely at least if not normally) supply a context according to which it is a *bad* or *inconvenient* thing that the door hinge should squeak. You might even feel compelled to offer the advice that the speaker use some graphite or liquid spray to lubricate the hinge. Yet if I offer a *special* context, the meaning changes: the speaker has just gotten through talking about how they like having noisy things around the house since their guard dog passed away and they now rely on the noise from squeaky hinges and creaky stairs for security. Now the squeak of the door hinge is *convenient* in that it accomplishes something for its owner. I take it that this illustrates the idea of a special context as Anscombe intended it: a context of utterance other than that which automatically comes to mind in

[11] We also see this idea employed by Geach (1956, p. 41), Foot, and others around the same time.

interpreting an otherwise out-of-context utterance. In this case the normal context is one that takes seriously the intended functioning of a hinge rather than the intentions or needs of its user.[12]

In light of this, how do we make sense of the notion of a "special" explanation? Well, the first pass answer seems to be relatively straightforward: a special explanation is an explanation other than the one which automatically comes to mind in seeking to explain an event for which no explanation is explicitly given. The car door unlocks when I hit the lock button on my wireless dongle and *vice versa*. The normal explanation doesn't apply since the dongle was designed and built to produce the opposite effect. Thus, we must appeal to a special explanation: wires became crossed in the dongle so that the structure of the circuitry has been corrupted.

In light of the insight about structural similarity from the beginning of this section, we can precisify the connection between special explanations and normality. A special explanation will be one that is structurally distinct from the set of normal explanations governing outcomes attaching to a particular antecedent.[13] The structural distinctness can serve as a mark of the special explanation. So operationally we'll start with the explanation of a normal outcome, and then we'll look at other possible outcomes and produce the most natural explanations for them. If any given explanation can be expressed in a structurally equivalent way, then it's (at least heuristically) a normal outcome. Whereas if it cannot be, then it (again perhaps heuristically) is an abnormal outcome.

1.2 Independent Explanations

An alternative route onto the intuition behind the explanatory link account starts with an insight from Pietroski and Rey's (1995) discussion of *ceteris paribus* claims that I believe extends to certain

[12] Whether the normal context is as a rule the context in which things or at least salient or central things fulfill their proper functions is an interesting question worthy of further exploration.

[13] There's a choice to be made here about what the set of explanations is that we're interested in. What's the comparison class? Here, since I've built the triggering conditions into the antecedent and so have relatively fine-grained antecedents (but not so fine grained that abnormal outcomes won't be included), I think I can just define the group as the set of outcomes attaching to a particular antecedent.

normality judgments—namely, certain *complete* normality judgments. A normality generalization is in part a promissory note that an independent explanation will be available for any apparent exception. An *independent* explanation is an explanation which makes appeal to phenomena or facts already verified by or employed in the explanation of *some other phenomena* than the one being explained. An example of an independent explanation is the following:

> A spot develops on my arm after years of exposure to sunlight. I ask my dermatologist why this happens, and she explains that the ultraviolet waves from the sun often damage the DNA inside of our skin cells, causing mutations which result in moles and sometimes in malignant skin cancer.

The appeal to UV rays is an independent explanation because we already have countless reasons to posit UV rays in explaining a variety of other phenomena. Bees apparently see UV light, UV rays can be used to sterilize dental equipment, and so on. Since we have all of this independent verification for the existence of UV rays, we can use them in the explanation of abnormal events like skin growths.

Pietroski and Rey offer a more outlandish example to illustrate the notion of a non-independent explanation. Consider the following case:

> My friend Adisa claims to be clairvoyant. Adisa, on more than one occasion, fails to see what clairvoyance should enable him to so. He appeals to "disturbances in the ectoplasm" to explain his failings.

This explains the exception to his clairvoyance, but it is neither posited in the explanation of any other phenomena nor is it verified independently of Adisa's failure to see in this instance what he supposedly can see in normal conditions. This is not an independent explanation because the only reason for positing ectoplasm in the first place is to explain Adisa's failings. In making certain complete normality claims (or a *ceteris paribus* claim in Pietroski and Rey), we are making a guarantee that any apparent exception will be explicable by appeal to factors not posited solely to explain those exceptions.

This insight is summarized nicely by Martin Smith (M. Smith n.d.):

> "In this sense of 'normal' it could be true that Tim is normally home by six, even if this occurrence is not particularly frequent. What is required is that exceptions to this generalisation are always explicable as exceptions by the citation of independent, interfering factors – his car broke down, he had a late meeting, etc. If this condition is met, then the best way to explain Tim's arrival time each day is to assign his arrival by six a privileged or default status and to contrastively explain other arrival times in relation to this default."

It could be true that my computer normally turns on when I press the power button, even in the strange circumstance where it didn't happen the last hundred times I've pressed it. This would require, though, that *for every one of those hundred instances (or for every* type *of instance) I can provide or at least posit an independent explanation for the failure*. In making this claim, I am therefore making a sort of guarantee that we will have or have the possibility of finding a special explanation.

It's worth noting a bit of reservation about this claim as I have presented it here. I don't see a need to reject a normality claim if there's a case the deviance of which we can't explain nor have the conceptual resources to imagine what the explanation might be. Some of this hinges on the relation between epistemology and ontology: there will be many cases the explanation to which we do not know and even for which we lack the conceptual resources needed to posit a potential explanation, but this shouldn't have an effect on what we think is in fact normal or what *ceteris paribus* claims we in fact endorse. This worry is easy enough to accommodate, but I have the lingering feeling that even though we want there to be an explanation of a great many abnormal cases, there doesn't seem to be anything in the concept of normality that guarantees that there shouldn't be unexplained abnormalities. Furthermore, I think cosmic accidents in a certain sense can't be explained (there are no principles governing them) and they happen quite often.[14] If that's true, then there are unexplained abnormalities all throughout time and space.

[14] This intuition relies on the distinction between explanation proper and merely describing what happened. Very roughly, explanation is supposed to tell us *why something should have happened,* while a mere account or description tells us *what in fact happened* without giving us a clear sense of why it should in principle have happened the way it did. Some things "just happen" and so their occurrence is a brute—in the sense of inexplicable—fact of the universe.

1.3 Normality in Explanatory Practice

The explanatory account is part of an overall picture of the triangular relationships between three core concepts: normality, explanation, and expectation. In giving this general picture, we'll get a clearer sense of the attractiveness of the explanatory account. None of this is strictly entailed by the original formulation of the explanatory account. Instead, we'll be building out a picture, with choice points all along the way.

In short, the normal is that which is explained by the kind of thing in question, which means that a) an explanation isn't called for pragmatically in conversations about individuals of that kind, b) the normal outcomes are the explanatorily privileged outcomes—those which will be assumed initially when constructing an explanatory narrative for a particular outcome. As a result of this connection between our explanatory practices and normality, there is a connection between normality and *expectation*, mediated by the connection between *explanation* and expectation. We are *pro tanto* or defeasibly warranted in expecting whatever would be explained by the kind of thing in question insofar as that thing is a thing of that kind. So, completing the triangle, the normal for a kind will be whatever one has *pro tanto* warrant to expect on the basis of an individual's being that kind of thing, and the normal for a kind is also whatever is explained by an individual's being that kind of thing. We've already got on the table a picture of how normality and explanation are connected, but this needs to be fleshed out by looking into the role normality plays in explanatory practices as an addition to our discussion from before of the role explanation can play in giving account of normality. This section will further flesh out the role of normality in explanatory practice.

Pietroski and Rey's discussion offers an additional insight: the normal outcomes are the explanatory starting point and so enjoy explanatory privilege over more abnormal outcomes. The first set of factors or conditions that our explanatory practices will make appeal to is the most

normal set. If that should fail, we'll introduce or seek more abnormalities until we've got a set of circumstances that explains the outcome we are in the business of trying to explain.

For example, if we are trying to explain why a man beat a cheetah in a footrace,[15] the order of possible explanations will likely proceed from the normal case—where both are running full speed and are unhindered—to incrementally more abnormal factors to investigate: say, the cheetah is distracted or uninterested in the race, or it has a tight tendon in one leg, or it's been hobbled in some way, or it's missing a limb, and so on. The last scenario we'd be likely to appeal to would therefore be the most abnormal—perhaps the cheetah was dead the whole time or the cheetah was in fact a stuffed children's toy. The intuition here is quite compelling: at the failure of the normal circumstances to explain a particular outcome, we'll introduce abnormalities until we have a set of circumstances that explain the outcome. If the normal circumstances in fact explain the outcome, then the outcome is itself normal.[16] If the cheetah wins the footrace, then we won't go searching for an explanation because it's exactly what we would've expected in normal circumstances to have happened. This is, it seems, an important part of what it is to be normal conditions.

So far, the discussion has centered on normal *conditions*, but what of normal *features* or *activities*? They should, if this insight is to help us understand normality writ large, enjoy similar explanatory privilege. Well, one thing to note is that there isn't a deep metaphysical cut between the different categories of outcomes: normal conditions for human beings involves the sun *doing* a great deal and the surfaces of objects *having* certain reflectance features, among countless other examples. Nevertheless, features and activities do enjoy explanatory privilege apart from their participation in explanations of other outcomes—it just won't show up as being the default assumption when constructing an explanatory narrative. Instead, their explanatory privilege consists in their status as

[15] Example comes from Romane Clark (1973).

[16] Note that this isn't meant to be an analysis of normal outcomes. I would expect this to work as a rule, though, and so for there to be extensional convergence between normal conditions and normal outcomes.

the contrast case that sets the context of explanation for other relevant outcomes. The contrast case for why questions or calls for explanations in a given context is—by default at least—the normal case. So, when asking "why did the man win the footrace?" the question will have an implicit contrastive tag to the effect of "why did the man win the footrace *rather than the cheetah*?" In normal circumstances, we'd expect the Cheetah to win, so a) we'll implicitly contrast any other outcome with this most normal outcome, and b) we'll go searching for the abnormality that explains the abnormal outcome if the man in fact wins the race.

So normalities are privileged in our explanatory practices in two ways: first, normal conditions are prior in our practice of building explanatory narratives and second, normal outcomes of all kinds are the default and often implicit contrast class for calls for explanations for relevant alternative outcomes.

Before moving on to address some worries for the account, it will do us well to take stock with a summary of insights from the preceding sections:

A. The explanatory link account illuminates the No Explanation Called For phenomenon. In turn cases where the relevant intuition is activated by a call for explanation are heuristically cases where the feature serving as the *explanandum* is normal for the kind of thing involved. For example, "why is that animal striped?" sounds odd if it's part of the common ground that that animal is a zebra. Zebras are normally striped, so it feels as if asking why an individual zebra is striped must be the result of some confusion.

B. The kind concepts in the common ground often come with their own explanatory resources and indeed often have explanatory power in themselves. An implicature is therefore triggered when one asks for an explanation that is already in the common ground.

C. Normal outcomes are those which require no special explanation, where "special explanation" is understood in terms of the automatic context one applies to a statement as a native speaker of a language and an inculcated participant in a culture. Special explanations will likely resist formulation in a structural identical form from normal explanations; whereas all possible normal explanations of a given outcome in a given respect for a given kind will likely allow structurally identical formulations.

D. Most complete normality statements seem to make a sort of guarantee that there will be an independent explanation of any deviant cases.

E. Normal outcomes will be explanatorily privileged in two ways:

a. Normal conditions are assumed when building an explanatory narrative, and we proceed to less normal conditions until we find a narrative which explains the outcome in question. That is, given two sets of conditions which would explain an outcome, we'll privilege the more normal conditions.

b. Normal conditions, features, activities, and relations will all serve as the default implicit contrast in calls for explanation without explicit contrasts.

2. Worries for the Explanatory Link Account

Now that we have a more complete picture of how normality and explanation are supposed to be related according to the explanatory link account, we have the resources on hand to address some lingering worries for the account.

2.1 Genera and higher-level kinds

One of the deepest worries for the explanatory link account of normality is that it *prima facie* rules out the possibility of true normality judgments at the level of a basic kind or species entailing or supporting weaker inferences to higher taxa which contain that basic kind. We'll use 'genus' loosely throughout the discussion to refer to higher taxa than a species, and we'll use 'species' a bit loosely as well to refer to the basic kind.[17] As an example, it seems true that:

(18) Sparrows fly.

(19) And that

(20) Birds fly.

Now, as we will see in the final chapter, (20) does not have the inferential profile that many philosophers and artificial intelligence researchers have thought. Nevertheless, it seems to be true that birds fly.

This presents a problem, however, because it's unclear at best that being a bird *explains* one's flying abilities and behavior. Being a *sparrow* seems to do explanatory work that being a *bird* never

[17] A basic kind might also, again echoing an Aristotelian perspective, the kind that is responsible for the creation or being of a thing—akin to a substantial form. When a metalworker creates a fork, they are first and foremost thinking of the thing *as a fork* rather than *as an instrument* or *as a utensil*. Cf. (Rosch 1978) and (Rosch, Mervis, et al. 1976).

could do. After all, birds also swim, waddle, and run. So even though flying is a somewhat unique and striking feature of birds as a taxonomical class, it doesn't clearly seem to be an activity or ability of an individual bird that can be explained by appeal to *that* kind. The explanatory link account claims that what it is to be normal for a kind is to be explicable with reference to that kind, but it seems uncontroversial that it's normal for birds to fly while also being at best controversial that any individual bird's flying can be explained by appeal to its being a bird.

Other examples give somewhat different results but appear to deepen the worry. Consider:

(21) Platypuses lay eggs

(22) Mammals lay eggs

(21) is fairly uncontroversially true, but in this case (22) is difficult to assess for truth. It is clear that some species of mammals lay eggs as a matter of course. It is less clear that those mammals lay eggs *because they are mammals*. It's intuitively unclear, furthermore, that it is correct to claim that mammals lay eggs.

How does this deepen the worry? The explanatory link account claims, roughly, that one's normal features are those explained by the type of thing one is. So, if it's not one's being a mammal that explains one's laying eggs—which seems correct to say of female platypuses—then it can't be normal for mammals to lay eggs. The result would be that it is not normal for mammals to lay eggs, but it would be odd to claim that a normal platypus was an abnormal mammal. Normalities attaching to higher taxa shouldn't overrule normalities attaching to the basic kind since the basic kind—e.g. the species—is tied in intimate ways to the niche of the animal and many more traits than are higher taxa.

Now that the problem is clear, let us explore a few strategies for dispatching it. As a preliminary, we'll need a provisional definition of species and genus that applies not only in the biological domain but also in the artifactual domain among others. In Aquinas, we find a definition rooted in

the Aristotelian categories. A species is defined by a substantial form shared by all of the members of the species. A genus, alternatively, is defined merely by shared material facts between species. All animals, for instance, have similar structures that are analogous across the genus of animals. They form at worst a family that bears resemblance relations to one another, but at best they share a great number of analogous structures the possession of which is definitive of being an animal.

Furthermore, the commonalities between member species of a genus are limited and abstracted. Mammals tend to have *auricularis* muscles—the muscles that move one's ears. In humans they are vestigial, but especially in prey species like mule deer, they are an important means of collecting acute auditory information from all angles. From an Aristotelian perspective, ancestry plays very little role, and so the genera are merely analogous groupings according to similarity of physical structure and other material features. From a modern scientific perspective, though, the explanation for these analogous features is owed to a common ancestry that branched out into different species to occupy and construct different niches in the ecological landscape. Today, therefore, we can give some explanatory weight to a taxon like *mammalia* beyond the mere analogy of material features. The features that can be explained by being a mammal, though, must be limited as there are countless and quite diverse species of mammals.

There are two visions for what, say, a genus is. According to one view—call it the Aristotelian perspective—the genus has no "real definition": it is a mere (ahistorical) list of criteria for grouping or material similarities. On that view, the only features that could possibly be explained by being a mammal would be the list of criteria for being a mammal and any features that would follow—logically or in terms of natural causation or the like—from the possession of the criterial features. According to a modern scientific perspective, a genus is a historical lineage with a connection to certain selection pressures and a wide range of member species with their particular niches and selection pressures. A genus has a *real definition* on this picture in that one can discover surprising

facts about a genus. We discovered not too long ago that a monotreme was a sort of mammal. For scientists, it was the lineage that decided in favor of grouping the platypus with its closest cousins—other mammals. This would be impossible on the Aristotelian view since it involves denying that certain features that are *definitive* of mammals are in fact criteria for belonging to the group—Aristotelians throughout much of history denied relations of lineage between species and so nothing but the definitive features could group species into higher taxa.

The Aristotelian vision of the higher taxa has other drawbacks, but rather than adjudicate these two pictures, it will be more fruitful for our current purposes to continue on as if the Aristotelian view is incorrect. Genera and other higher taxa are more than mere sets of features: they are explanatorily potent, even if quite limited, categories that play an important role in our explanatory practices—in part because they are historical families in addition to being similarity classes. There are three questions which now present themselves: Why think that the species is distinct from higher taxa? Second, why think, then, that higher taxa have explanatory relevance? And Third, why doesn't a higher taxon explain an individual member's possession of a feature which is proper to a member species?

What makes a species or basic kind distinct from the higher taxa of which it is a member is its constitutive relation to the niche occupied thereby. Mammals don't occupy a niche in any thick sense of the word since mammals from the blue whale to the tiny Etruscan shrew share little if anything in common in terms of environment, food source, or territory. On the conception of kinds at play in this study, a kind and a niche are deeply interconnected in that the kind is essentially a particular range of responses to a niche that has been carved out over time by members of that kind.[18] The niche places ecological demands on the species and the species members respond in a variety of ways to those demands. The selection pressures of the niche shape the morphology and

[18] Cf. (Odling-Smee et al. 1996, 2003)

behavior of the individual members of the species, while the morphology and behavior of the individual members of the species in turn shapes the niche into something more compatible with the thriving of the species. Mammals as a kind didn't carve out a niche for themselves, the individual member species carved out their own niches. If that is true, then many of the selection pressures and demands that make species such explanatorily robust kinds aren't present for a genus and therefore the genus concept does less explanatory work.

Yet this presents a problem: do higher taxa carry *no* explanatory weight? It seems that being a mammal *can* explain certain features like a suckling reflex or the possession of fur. It'd be odd, therefore, if what has been said so far rules out the possibility that genera have explanatory power. Being *silverware* explains why a fork should be on the dining table just before dinner, or so it seems.

The account sketched thus far does not, in fact, rule out such a possibility. Even though genera do indeed have *less* explanatory power than their member species, they don't *lack* explanatory power. The few features that characterize the genus—apart from features "borrowed" from its member species—are indeed explicable with reference to the genus. In lieu of a more complete treatment of the relationship between higher taxa and normality judgments, there are a few things to say to dispense with the worry that certain seemingly true normality judgments are not classed as such by the explanatory link account.

Being a mammal explains the presence of mammary glands in a female individual, one's giving birth to live young, and perhaps a few other features. It also explains anything that follows of necessity or by natural causation from the possession of those features. Mammals arrived on the scene many epochs ago, producing milk for their liveborn offspring. This isn't universal to kinds of mammals, but it is an important and salient phenotypical and behavioral trait that helps us to understand how mammals survived and how, for instance, many of them survived a catastrophic asteroid impact.

Being a mammal doesn't appear to explain much that being a member of an individual species can explain. For instance, mammals don't as a (statistical) general rule fly. Being a fruit bat, though, would explain why an individual is able to fly. Mammals do, as a statistical matter, tend to be land dwellers, but it's unclear how being a mammal explains one's being a land dweller. Being a whale explains one's dwelling in the sea just as being a camel explains one's dwelling on the land (perhaps even a quite specific type of land), but neither of these appears to be explained by either's being a mammal.

Thus far, it seems clear that some features can be explained by being a member of a higher taxon, while others cannot. Higher taxa have relatively weak explanatory power. It does seem true, though, that:

(23) Birds fly

(24) Mammals give birth to live young

(25) Trees have green leaves.

None of these behaviors, upon reflection, seems clearly explicable by reference to the taxon cited since each is merely one of a number of contrary possibilities present in the genus. There are a few strategies for dealing with this tension (other than conceding that the explanatory link account is inaccurate) worth considering: a) perhaps these statements simply aren't true, b) maybe the generic forms of (23), (24), and (25) are stand-ins for *prototypical judgments* rather than for normality judgments, or c) perhaps these are true as a result of a process whereby higher taxa "borrow" normalities from their member species.

The first strategy (a) is, as far as I can tell, a non-starter. It is true that birds fly and that mammals give birth to live young. These are not marginal cases that theory can decide downstream—they are basic cases that must be dealt with in an account such as the one on offer here. Furthermore, these cases are exemplars of but one grammatical construction used to express normality judgments. As

we'll see in the discussion of the second possibility, other grammatical constructions fair better than the generic form. Finally, if we deny the truth of these judgments or at least analogous judgments, we must also deny the truth, so it seems, of all incomplete normality judgments. I, for one, am not willing to abandon the truth of (26):

(26) It's normal for cats to be black

Even if claiming that this is false does away with more problematic normality judgments about genera, it is so clearly true that this would be a price not worth paying.

According to the second strategy (b), we can make a distinction between normality judgments—which have the inferential profile and content we've so far discussed—and prototypical judgments—which ascribe not normal or normative features to a kind but instead characterize a kind in terms of *salient* or *common* features. (23) is true not being flying is normal for birds, but because flying is a striking or salient feature of birds. Leslie (2007, 2008, 2012a) has theorized that this is one of a few uses of generic judgments in human discourse and psychology. This doesn't, however, clearly have much to do with the sense of normality at play in normality judgments, where salience and particularly prevalence play little role if any. On this strategy, then, generics as a grammatical category sometimes express normality judgments, but sometimes express prototypicality or perhaps even statistical judgments. This claim is quite plausible, in fact.

This may be all well and good, but what of grammatical forms like the following?

(27) It is normal for birds to fly.

(28) It is normal for mammals to give birth to live young.

(29) It is normal for trees to have green leaves.

One can imagine natural contexts surrounding each that make each clearly a true statement. A context where someone is surprised to see a flying bird or a mammal giving birth to live young or a tree with green leaves will be a context in which saying something along the lines of (27)-(29) would

be an appropriate and true response. Each therefore seems uncontroversially true, but each also seems to have to do with normality in the sense under analysis. There's no getting around the facts that a) these are true normality judgments, and b) these attach to higher taxa while also perhaps more fundamentally attaching to member species of those taxa. The problem is that it isn't clearly correct to say that sparrows fly *because they're birds* or that sycamores have green leaves *because they're trees*. These sorts of cases seem to require something along the lines of the third strategy introduced two paragraphs above.

On such a strategy (c), there is a process by which a genus can "borrow" a normality from its member species. Sure, it is true that it is normal for mammals to have venomous barbs: that's one normal way to be a mammal, after all! But it nevertheless seems odd to say that being a mammal explains one's having a venomous barb. This seems to be because there are lots of normal ways for a mammal to be: both possessing *and* lacking a venomous barb seem to be perfectly normal ways to be a mammal. Being a mammal, that is, seems to be indifferent to possessing a venomous barb, just as it is indifferent to possessing a mane or dwelling on land. If the kind is indifferent to a feature, then anything goes and it is therefore odd to say anything at all about the connection between mammals and venomous barb possession. It may not be right to say that male platypuses have venomous barbs *because they're mammals*, but it is nevertheless normal for mammals to have venomous barbs in virtue of the facts that male platypuses are normal mammals and that they normally have venomous barbs. But this is okay, since the fact about mammals is simply derived from the fact about species rather than being a basic fact. Genera both have their own basic normalities and derive normalities from their member species.

The explanatory link account is primarily focused on understanding normality judgments at the species level. If, therefore, the account cannot entirely account for normality judgments at the level of higher taxa, then the account may either have to expand to encompass the complexities involved

with genera and the like or develop a principled reason for rejecting the truth of these normality judgments. Here is one potential solution to some of the complexities introduced in this section. No doubt there are other complexities worthy of their own study.

One potential upshot of this discussion is that the explanatory link account needs to be revised. It appears to be too restrictive in classing as normal only those features which are explained by their possessors being the kind of thing for which that feature is normal. In addition, one must allow for a different sort of normal feature:

> For a higher taxon *g*, a feature or outcome *φ*, one of *g*'s member taxons *s*, and an individual *i,*
>
> *φ* is normal for *g*'s iff *i*'s possession of *φ* is explained by *i*'s being a *g* OR by *i*'s being an *s*.

This account holds that a feature is normal for a kind if and only if that kind explains its possession or a sub-kind of that kind explains its possession. Faced with the dilemma of denying that it's normal for mammals to fly and allowing that "mammals fly" is a true normality judgment at the cost of the original formulation of the explanatory link account, I've chosen to revise the original formulation in order to maintain the apparent truth of many normality judgments about higher taxa.

This addition has consequences for the inferential structure of normality judgments, which will be explore further in chapter 7. Of note here, though, is that there are a variety of normality judgments which have different status in regards to both their content and inferential structure.

- Complete non-derivative normality judgments [there is only one way to be normal, and the kind explains the outcome]: dogs have four legs, knives are sharp, cats stalk.
- Complete derivative normality judgments [there is only one way to be normal, but the normality of the outcome is derived from a or some member species]: not clearly possible except by cosmic accident.
- Incomplete non-derivative normality judgments [there are multiple contrary ways of being normal, and the kind explains the outcome]: cats are black, cats are calico, cats are grey.
- Incomplete derivative normality judgment [there are multiple contrary ways of being normal, but the normality of the outcome is derived from a or some member species]: birds fly, birds swim, birds run, birds waddle.

This depth of distinction is a boon to the account. Intuitively there are different sorts of genus-level normality judgments and so a picture of how normality judgments function in terms of content and structure is better for including such a distinction. Rather than work around it, the explanatory link account actually offers the tools necessary to make the necessary cuts in conceptual space.

Again, this is a preliminary revision based on a few core cases. If it turned out to have unsavory consequences, then very little stands or falls with it and so it can be rescinded without dire consequences. It does, at least as a patch, appear to answer the initial worry that being a member of a genus does not explain one's having certain features that appear to be normal for members of that genus.

2.2 Accidental Success

The explanatory link account holds that normal features are the feature that are explained by the kind of thing involved. There is a *prima facie* problem with this account, though, in that it's not merely in principal possible, but indeed actual that there would be accidentally successful cases—cases where an individual has a feature that is normal for the kind of thing involved, but which does not have the feature on account of the kind of thing it is.

We can imagine coherent scenarios in which a piano's strings are hammered in a sequence perfectly matching a player's key strokes, but not *because* of those key strokes (perhaps it's a player piano whose keys are causally disconnected from the strings).[19] In this case, it seems like the account will mark those as normal outcomes because they're the kind of outcome that is explained by the kind of thing when they obtain (perhaps with the rider "by default" attached). Why are the piano strings hammered at the same time that the corresponding keys are pressed? It's a piano!

But what about cases where the accidental connection isn't between two parts (conceptual or mechanical or otherwise) of the thing in question, but between the kind and the individual outcome?

[19] Nickel (2016, p. 54) discusses the example of an albino raven being painted or dyed.

It's normal for ravens to be black, and any particular black raven's being black is explained by its being a raven. But what of an albino raven that has fallen in a vat of black paint and emerged otherwise unscathed? It's being black isn't explained by its being a raven; but it is black, and ravens are black. So, the problem here is supposed to be that we have an outcome which is the normal outcome—a raven is black—but which can't be explained by appeal to the kind of thing involved—i.e. "it's a raven." We can't say that a normal outcome is that which is explained by the kind of thing involved and then specify this particular outcome as "being black on account of being a Raven." That would just be to specify this outcome as "being black in the normal way" and then our analysis is unilluminating since it doesn't distinguish *in a principled way* between truly normal outcomes and accidentally "normal" outcomes.

There's a relatively simple way out of this problem: the normal outcomes are all and only the ones that are explained by the kind of thing involved, so the strings of a piano being hammered when the corresponding key is struck is only a normal outcome when it is explained by the kind of thing involved—which it surely isn't if there's some extra mechanism involved which automatically hammers the strings in a sequence that happens to correspond to the striking of the keys.

The worry was that an accidental success case violates an entailment of the Explanatory Link Account to the effect that

> (30) If it's normal for an individual to have a particular feature, then that individual's having that feature will be explained by its being a member of that kind.

So if it's normal for ravens to be black, then any raven's being black will be explained by its being a raven. This is a problem, though, because it is normal for ravens to be black but there are at least imagined cases aplenty of ravens that are black, but whose blackness is not explained by their being a raven. Straightforward counterexamples. Fortunately, there simply is no such entailment. One entailment of the account is:

(31) if an individual's having a feature is explained by its being a member of a kind, then that feature is normal for that kind.

If a raven's being black is explained by its being a raven, then it is normal for ravens to be black.

Another reverse entailment is specified slightly differently:

(32) If a feature token is normal for a kind of thing, then its possession will be explained by its possessor's being that kind of thing.

This entailment seems to hold for the cases we've thus far considered. The status of being a raven of an individual raven that is black in the normal way—that has a normal token feature of being black in the normal way—explains why that individual raven is black. The claim is that it will hold for all such features.

2.3 Unique Features

Let's start with a class of cases that on the one hand appear to be perfectly normal, but on the other hand aren't obviously accounted for by the explanatory link account. Here are a few such cases:

(33) Why does that animal play the guitar (rather than not play any instrument at all)?

(34) Why does that animal write poetry (rather than not writing anything at all)?

(35) Why does that animal have a vocation (as opposed to merely surviving)?

In each case, the why question calls for an explanation of a particular behavior or feature which is unique to humans. The answer to such questions, though, doesn't seem to be "because they're a human." It might be obtuse to provide such a response, but nevertheless the feature or behavior in each case is uncontroversially normal for human beings. If the kind of thing in question doesn't explain the feature in question, then the explanatory link account will not judge these features and behaviors to be normal. If Jimi doesn't play the guitar *because he's human,* then it isn't normal for humans to play guitar. This would be a serious problem indeed.[20]

[20] One might distinguish between two sorts of normalities: features that are compatible with or not in tension with the nature of the kind of thing involved and features that are determined or caused by the nature of the kind of thing involved. Being a human in some sense *causes* one to have four limbs, but does not likewise cause one to play the guitar.

Part of the problem here is again looseness with respect to the questions that can be answered by appeal to bare specification—by simply specifying the individual in question.[21] So three things could be happening with "why does that animal play the guitar?" First, the question might, as discussed earlier, be a question about why *humans* play the guitar. If that is the case, then there is no problem because playing guitar is normal for humans and the answer to "why do humans play guitar" has many different sorts of answers. There's no worry here for the explanatory link account. Second, the question might be asking why *Jimi* plays the guitar, in which case an individual historical answer seems appropriate and merely saying that Jimi is a human seems strange at best and entirely unhelpful at worst. Third, the question might be the sort of question we're interested in. It might be that the questioner is unclear about either or both of the facts that a) the individual in question is a human, and b) that it's normal for humans to play the guitar. In this case, the answer of "that's a human being" along with the associated implicature "it's normal for humans to play the guitar" explains the *explanandum* perfectly well.

It's the second possibility which creates the apparent problem. If the explanation-seeker asks why Jimi plays the guitar, the natural answer is an individual historical[22] or even perhaps biomechanical answer (like he has a weak diaphragm and so can't play the woodwind). At best, this would support the judgment that it was normal for Jimi to play the guitar. But the idea was supposed

---

Nevertheless, it is certainly normal for humans to play the guitar. The explanatory link account will hold that even in the cases where the feature is merely not in tension with the nature involved, that nature nevertheless still explains why the individual should be doing that thing. Playing the guitar requires have a particular physiology and bodily comportment with a guitar, as well as having a sense of rhythm and musicality. All of these features are determined by—not merely compatible with—human nature.

[21] It won't help solve the particular problem in question, but it is important to note that a different question can, I think, unambiguously be answered by bare specification. "Why is Jimi able to play the guitar?" has the strange implicature that marks the No Explanation Called For phenomenon. It seems like there must be something special about Jimi which the asker is trying to ascertain, but in fact Jimi is able to play the guitar because Jimi is a human, not because Jimi is special in any way (of course Jimi Hendrix is able to play so well because he is special, but he's not special in being able to play full stop).

[22] A memorable cola commercial offers the imaginative explanation that Jimi Hendrix preferred one brand of cola and so instead of buying the off brand near the accordion shop, he crossed the street to buy the name brand and as a result laid eyes on the Stratocaster he was famous for playing. The advertisement was good, but I don't remember which brand of cola it was, so it was only so effective.

to be that Jimi plays the guitar *because* he's a human. That is, it's normal for humans to play the guitar in that individual humans play the guitar *because they're human*. Furthermore, answering the question "why does Jimi play the guitar?" with "because he's a human" seems strange or obtuse given what might be considered the standard interpretation of the question. Jimi, according to one explanation, plays the guitar because he's human. But it doesn't seem that "because he's human" answers the question "why does Jimi play the guitar" and so it doesn't seem like the explanatory link account can handle the datum.

Not to fear, though, since careful attention to the contrasts involved help to clear things up. Recall the three possible interpretations of "why does that animal play the guitar?" from before. If we add contrasts to each, it makes it a bit clearer what is going on between the cases.

(36) Why does that animal play the guitar (rather than not being musical at all)?

(37) Why does that animal play the guitar (instead of some other instrument)?

(38) Why does that animal play the guitar (rather than use the guitar to build a nest)?

Wondering why an individual plays the guitar rather than another instrument demands an explanation at the individual level (there aren't species-specific instruments, after all). So (37) is simply asking the wrong sort of thing for an appropriate response to be: "because he's a human." For (36) as well as for (38), though, the response "because he's a human" seems to answer the question just fine. Jimi plays the guitar. Playing the guitar is normal for humans. Jimi's playing the guitar (rather than not being musical at all) is explained by his being a human. Jimi's playing the guitar rather than the oboe, on the other hand, is explained by individual historical and biomechanical facts about Jimi the individual. So, the problem dissolves. What is normal for Jimi *qua* human is what is explained by Jimi's being a human.

2.4 Ecology

At the level of the ecosystem, kinds begin to take on an irreducibly relational and holistic character. This raises a potential problem for the explanatory link account: some deer don't die *because they're deer* but instead die because they are predated by a mountain lion or wolf. But the relationship between wolves and deer is in some sense *internal* to each kind and so it feels as if it might be normal for deer to be predated by wolves. The two species have a sort of symbiotic relationship where the predation by the wolves helps to cull the deer population to such a degree that the resources of the deer habitat are sustainable. Conversely, the wolves are fed to an appropriate degree, but unfortunately deer are scarce enough and difficult enough to catch that the wolf population remains low through starvation. This balance has been achieved at great cost throughout the evolutionary history of the two species within this ecosystem. So, where's the problem for the explanatory link account?

If it turns out that a particular deer doesn't die *because it's a deer*, but because of the complex ecological relationships that deer have to their environments, then the explanatory link account may not call it normal for a deer to be predated by a wolf. But this seems like it might be normal if for no other reason than that it is necessary for the continuance of the current state of the species. So, an apparent problem arises for the account: it might not be normal for an individual deer to be predated by a wolf, but deer are predated by wolves. What is normal isn't what is explained by the individual's being the kind of thing it is, or so it seems.

Again, there appears to be a bit of a type-token ambiguity at play here. An individual deer, going through its normal deer activities, won't normally get eaten by a wolf. The trajectory of its life and all of its functions are geared toward survival and reproduction rather than towards predation. That being said, deer get eaten by wolves all of the time. And this seems to go beyond a statistical fact. If we see it happen, we might think it's tragic, but we won't think "how awful! It doesn't have to be this way. It's just not normal." We'll instead think "that's nature at work." There stands, therefore,

the possibility that one must take into consideration two different sorts of things in deciding whether or not it's normal for deer to be eaten by a wolf. An individual deer isn't *supposed* to be predated by a wolf, but as a pair of kinds deer and wolves share a sort of symbiosis such that deer *are* supposed to be predated by wolves.

Another piece of the puzzle is the possibility that there are two notions of a kind at work here. One would be a narrow conception according to which only what happens inside of the boundaries of the body of a deer count as being part of the kind. So the deer's normal movements, development, etc. are all part of the kind irrespective of facts about the environment or of ecological relations or the like. Another conception of a kind would be a broad view according to which a kind is defined in part by its ecological relations, environmental conditions, etc. There may also be a different conception of the kind that is the kind viewed merely as an ecological entity and no longer as a collection of individuals with their own particular goods and harms. Because of this potential ambiguity, we might find ourselves with divided intuitions about certain cases which involve ecological relations. Broadly conceived, deer normally get eaten by wolves. Indeed, they *need* to be predated on in order to maintain balance in their environment. In some way, it is what is *supposed to happen*. Narrowly conceived, something has gone deeply wrong for the deer in its fatal encounter with the wolf. It's not *supposed* to happen.

This is a delicate topic and one requiring more care than we are able to give it here. These are some preliminary thoughts which alleviate the initial tension, but no doubt introduce tensions of their own. As per usual, this issue is worth a study of its own and this ought not to be read as a settled matter.

2.5 Covers Nomic Cases Too?

Our original division was between statistical, nomic, and normality generalizations. If it turns out that the explanatory link account of normality will classify certain nomic generalizations[23] as normality generalizations, then there's a *prima facie* problem with the explanatory link account: it counts as normality judgments certain judgments which are not normality judgments because they are universal laws.

The first question to address is whether or not it is the job of the account of the *content* of normality judgments to distinguish between nomic and normality generalizations when we a) already have a means of distinguishing these judgments by means of their structure, and b) the point of an account of the content of such judgments is—at least primarily—to understand the class of judgment as a family rather than to distinguish it from other classes. The full account might just as well be that a normality judgment is one that has the content dictated by the explanatory link account *and* has the inferential structure outlined in both the first and last chapters of this dissertation.

Nevertheless, it feels like we'd be better off if the account in fact doesn't cover certain normic judgments. That is, if *content* of the nomic judgments in question is captured in an account of the content of normality judgment, then it seems the account of normality judgments is too broad. Here are a few cases to consider:

(39) Prime numbers are divisible only by one and themselves.

(40) Squares are composed of two right triangles

(41) Fir trees are plants.

[23] There appear to be accidental and non-accidental nomic generalizations. It's not entirely clear how to deal with the accidental nomic judgments—judgments like "all lions live on Earth" or "all US supreme court justices have odd social security numbers." They may be universally true (suppose the second one is in fact true of all US supreme court justices) but they are about as modally "fragile" as anything: send a lion to the International Space Station or appoint a justice with an even social security number and the judgment falls apart. These judgments, therefore, I am tempted to treat as a special class of *statistical* judgment where the statistical value is 100%.

(42) Bishops move diagonally.

In each case, though it feels odd to say that it is normal for these kinds of things to have the features they do—indeed, it is necessary that they have these features—the explanatory link account will apparently call this a normality because the kind explains the feature in question. Why is 13 divisible by only one and itself? Because it's a prime number. Why is it that one can divide this shape into two right triangles? Because it's a square. Why is that thing a plant? Because it's a fir tree and fir trees are plants. At least that's the answer we're supposed to accept if this worry is to have the force it is supposed to have. Why did that player move that piece along that diagonal? Because it's a bishop and bishops move along diagonals.

To be sure, these question and answer pairs seem more natural when put a different way: Why is 13 divisible by only one and itself? It's a prime number. *Every* prime number is only divisible by one and itself. And so on for the other examples. The quantified forms involving 'all,' 'every,' 'each,' 'only,' or the like are no doubt more natural than the generic forms. Nevertheless, there is a problem here: many nomic connections or necessities hold as a result of or in such a way to ground the power of the kind in question to explain the outcome; but nomic cases shouldn't count as normality cases because nomic judgments are fundamentally different. Here are a few strategies for addressing this problem.

First, we might think that there are actually three distinct uses of the word "normal" or other words in the intensional neighborhood: a statististical use, a nomic use, and a "normic" use. It's already clear that there is a statistical and a normic reading of "normally". Take, for instance, the following judgments:

(43) Mammals normally give birth to live young.

(44) A piano player will normally know where middle C is.

(43) doesn't attach to any biological norm: it seems to be a claim about *most* mammals or mammal species. Few would claim that there is anything abnormal in the normative sense about platypuses. They're odd or peculiar mammals, not defective or abnormal. (44), conversely, seems grounded in a basic fact about playing the piano: identifying middle C is a fundamental or basic skill in understanding the layout of a keyboard and its relation to musical notation.[24]

The response would then be to posit a third use of the words 'normal' and 'normally' that has a nomic inferential structure and perhaps content. A statement like (45) would then have the same content as the nomic judgment (46).

(45) Bishops normally move diagonally

(46) Bishops only move diagonally

It would simply be expressed using a word that *appears* to have something to do with non-statistical, non-nomic normality, just like (43).

I'm tentatively unconvinced that this is a good strategy. It just doesn't seem possible to use these words without inviting considerations of normality and without activating the inferential profile of normality judgments—with the exception of statistical uses of these words and phrases.

Second, we might note that many normality judgments are *compatible* with the truth of a corresponding nomic judgment even if they have different content. For instance, here's a normality judgment:

(47) It's normal for bishops to move diagonally.

Of course, it *is* normal for bishops to move diagonally: that's the *only* normal way for them to behave. This at least appears to be true. It's awkward, though, because there is simply something much stronger to say:

[24] Note that generic forms like those in (39), (40), (41), (42) seem to admit of a nomic interpretation. Generics, it appears, are a slippery group being as they are individuated along grammatical rather than semantic grounds. Grammar won't always be a good indicator of the sort of content involved in a group of statements, claims, or judgments.

(48) (46) Bishops only move diagonally.

There seems to be another pragmatic phenomenon at play here. (47) appears to be implicating that there *isn't* something much stronger to say—that (46) isn't also true. Perhaps the maxim of quantity is being flouted here. This strategy seems to be sound: normality judgments are often compatible with the truth of the corresponding nomic judgments, but because of the pragmatics involved, they come across as strange statements at best.

The third strategy would be to precisify the explanatory links involved. On this strategy, (49)will be cashed out along the lines of (50):

(49) All cats are mammals

(50) Any individual cat being a mammal is explained by the conceptual relation between CAT and MAMMAL.

Alternatively, (51) will, just as always, be cashed out in terms of the explanatory link between being a cat and stalking behavior.

(51) Cats stalk

It's not a matter of conceptual truths or universal laws, but as a matter of what it is to be a certain kind of thing. It's no conceptual truth that cats should stalk, but it is a truth grounded in the nature of cats—the contingent and multifarious nature of an actual lineage and population of living creatures in multiple ever-changing environments.

Finally, we might question whether the normality judgment form of these statements is actually true. After all, it is quite odd to say Bishops *normally* move along diagonals, because the rules of the game dictate that they will *only ever* do so.

In sum, though it is in fact a problem that the content ascribed to normality judgments by the explanatory link account appears to capture certain nomic judgments, there are a few strategies for

dealing with these cases that undercut the force of the objection. At this preliminary stage, it is unclear which strategy is preferable, but there does in fact seem to be *a* response to the objection.

Chapter Five: Aquinas, Normality, and Regularity

Sometimes when we use the word 'normal' or 'normally,' we mean to express the judgment that some feature is frequent, ordinary, average, or prevalent amongst a kind of thing. Alternatively, we often use these words—as well as sentences with other surface grammars—to express a different sort of judgment. These latter judgments have more to do with fittingness, naturalness, or functionality than with frequency. The relation in extension between these two concepts (and, no doubt, in intension as well) is complex and multifaceted, and I intend to make some progress toward understanding it.

If our goal is to gain some understanding of the link between normality and frequency, we would do well to consult the Aristotelian tradition, where resources other than those found in relevant contemporary discussions are available and relatively untapped. Since Aquinas was among the first major Aristotelian thinkers in the Latin West, his corpus is a natural starting point.[1] The first goal is therefore to settle what sorts of uses Aquinas makes of words associated with normality and frequency. In service of this goal, it will serve us to look at previous scholarship on Aristotle's use of the from *hos epi to polu*—often translated as "for the most part."

Aquinas seems to think—seemingly following Aristotle—that normality and frequency are tightly linked at least in the actual world. This presents a problem, since most people would agree today that a feature can be normal without any particular frequency threshold being satisfied; and that a feature can be overwhelmingly frequent or even *de facto* or actually universal to a kind of thing without being normal for that kind of thing. Furthermore, most will agree that the normal is often not actual. It appears, therefore, that Aquinas is out of sync with highly plausible philosophical

[1] There are many earlier discussions of similar themes, to be sure. See Martin (2003) for a discussion of abnormality in Abelard.

beliefs about the relationship between normality and frequency. Furthermore, it seems that many things that Aquinas himself says run counter to the claim that they are tightly linked in extension. The second goal is therefore critical and philosophical in nature: to attempt to understand how he could hold the view he does while apparently being aware of potential counterexamples to that view.

Section 1 will look briefly at Aristotle's use of the Greek phrase standardly translated as "for the most part". Section 2 will explore Aquinas' use of similar Latin phrases, which presents two problems. Section 3 will explore the first problematic and section 4 the second problematic. Section 5 will present some upshots of the discussion.

1. Aristotle

Aristotle famously used the phrase *hos epi to polu* or "for the most part" throughout his corpus, particularly in discussions of biology, ethics, and epistemology. Some things in nature happen all the time or of necessity. For example, grave things *always* try to move downwards, and light things *always* try to move upwards. The heavenly bodies *always* move in perfect circular motion. Other things, though, happen only for the most part: grave things *in fact* move down only for the most part because they are sometimes impeded or subject to violent motion. Yet both sorts of generalizations are for Aristotle the object of scientific knowledge. We can demonstrate with them and they can be the objects of demonstration.[2] Thus, for Aristotle generalizations about what happens "for the most part" are serious generalizations rather than being merely heuristic or pragmatic versions of the true, universal generalizations.

Did Aristotle, though, use the phrase '*hos epi to polu'* (hereafter HETP) univocally? Scholars have already demonstrated that he did not. While it seems relatively clear that there is one central reading

[2] There are different routes we can go in understanding Aristotle here. We might think, with Anagnostopoulos (1994), that Aristotle had two different conceptions of demonstration only one of which can involve *hos epi to polu* propositions and the other of which is exclusively scientific in the fullest sense. Alternatively, with Reeve (1992), we might think that Aristotle had two different conceptions of science. There are other interpretations one might hold as well.

of the phrase, Aristotle appears to have employed it to different effects in his corpus. Henry (2015) offers the following taxonomy of uses:

A. Causal Regularities
   A1. *Ceteris Paribus* Laws
   A2. More and Less
B. Correlation without Causation

Most centrally, Aristotle uses HETP to describe regularities secured through a consistent cause. These are non-accidental and so *per se* causal regularities, but they nevertheless are not necessary connections. The focal case of HETP causal regularity for Aristotle is what Henry (perhaps contentiously) calls a "*Ceteris Paribus* law" (hereafter 'CP law'). For Henry, a CP law is a "a general law whose scope is restricted to what occurs under normal circumstances and therefore excludes cases where certain factors interfere with the normal causal relationships."[3] It is an understatement to say that Henry's use of CP laws as an interpretive tool for Aristotle's HETP is worth a second look, given that there are all sorts of generalizations other than CP laws which are, in their own right, qualified candidates for the interpretation of Aristotle's HETP. That being said, the focus here is on Aquinas' view, and so we need to bog ourselves down in the interpretive details of Aristotle's use of the phrase.

Henry's next category is less fraught. These are generalizations about what happens only for the most part because whatever the cause is—perhaps an underlying mechanism—is variable in output. Humans have a certain range of hair follicles on their head, but many have less due to a certain process being interfered with, or a certain resource being in short supply during critical periods, or simply because the mechanism(s) that determines the count of hair follicles is variable in its output based on a whole host of factors or perhaps simply because it's indeterministic. Many, to be sure, also have more than the statistical norm of hair follicles on their head. There isn't too much to say

[3] *Op. Cit.,* n11, p. 171.

about this sort of generalization other than to say that for Aristotle every process, capacity, power, disposition, or mechanism has a proper level of specificity in its effect and so it's far from surprising that Aristotle should think demonstration is possible even regarding causes with varying effects.

Finally, Henry—contra Winter (1997) and Ferejohn (1991)—notes that Aristotle in a few places uses HETP to describe accidental connections—in this case connections not secured by a cause. Aristotle uses HETP to describe the connection between being a fissiped and giving birth to only one baby at a time, which he acknowledges is accidental. I think we can safely set aside such uses as slips in an otherwise relatively uniform use of HETP to describe regularities that attach to causes of different sorts. This may be controversial, but it seems that the general consensus among scholars is that Aristotle (for the most part) intended HETP to describe non-accidental connections—connections held on account of some principle or cause.

2. Aquinas

We are trying to find out what Aquinas' considered view is on the link between normality and frequency. To that end, we'll do well to first get clear on what sorts of Latin phrases Aquinas uses to express these two concepts. First, the English phrase "for the most part" is typically used to translate Aristotle's Greek phrase '*hos epi to polu.*' The situation in Aquinas' Latin, however, is not so clear. Translators have used that same English phrase—'for the most part'—along with others to

translate a variety of Latin phrases: *'ut in pluribus,*[4] *'frequenter,*[5] *'in maiori parte,*[6] *magis,*[7] *maxime,*[8] and the antonymous *in paucioribus,*[9] and *in minori parte.*[10]

Fortunately, in the vast majority of these cases, the somewhat ambiguous "for the most part" seems to map on to Aquinas' somewhat ambiguous meaning despite his use of different Latin phrases. In short, there don't appear to be *systematic* differences in the use of each phrase, so we can't draw any clear conclusions about distinctions between the use of each phrase *as such*. In absence of such conclusions, it seems safe to conclude that Aquinas uses these phrases more or less interchangeably or at minimum that context is a better cue to his meaning than is his selection of phrase. Additionally, context makes clear that these phrases are each used to mean more or less the same thing at different places in Aquinas' discussion. In particular, we find each phrase linked with *semper*—as in "always or for the most part" and those passages form the core of the data for the conclusion I draw here. There are many others worth looking at, but the central phenomenon at issue is the use of phrases that translate with some fidelity as "always or for the most part."

What, then, of words or phrases that express the concept I've called *normality*? Here, Aquinas uses again a variety of phrases,[11] but the meaning of each makes relatively clear that he is expressing at minimum a concept in the same intensional family as normality or its antonym abnormality. Most

---

[4] *In Metaphysica* 6, lectio 2, 1183, 1190; Foreword by St. Thomas in *In Post. Analytica*; *In Phys* 2, Lectio 12, 253; *Summa Contra Gentiles* 2.24.4; *Summa Contra Gentiles* 3a.5.2; *Summa Contra Gentiles* 3b.86.4; *Summa Contra Gentiles* 3b.125.10; *Summa Contra Gentiles* 1.44.7; *Summa Contra Gentiles* 1.3.5; *QDdV* 3.1.corpus

[5] Foreword by St. Thomas in *In Post. Analytica*; *In Post. Analytica* Lectio 32, caput 20; Com met. 1, lectio 6, 107; *Summa Contra Gentiles* 3a.6.4-7; *Summa Contra Gentiles* 3b.99.9; 1 *Perih.* 14 f

[6] Cf. for a sampling *In Metaphysica* 1, Lesson 5, 99; *Summa Contra Gentiles* 1.42.7; *Summa Contra Gentiles* 3.42.11; ST III.66.4.ad5; *QDdV* 5.2. corpus; *QDdV* 5.3. corpus; *QDdV* 12.3. corpus; *QDdV* 22.1. corpus; *QDdP* 3.6.corpus

[7] *In Metaphysica* 6, lectio 2, 1190

[8] Ethica 1. Lectio 16. 192

[9] *In Phys* 2, Lectio 8, 208; *In Phys* 2, Lectio 12, 254; *Summa Contra Gentiles* 3a.5.2; *Summa Contra Gentiles* 3a.6.4; *Summa Contra Gentiles* 3b.99.9; *QDdV* 3.1.corpus

[10] *In Phys* 2, Lectio 8, 209

[11] There is a rather obscure word '*normalis'* which is related principally to a "right angle" or a square. It doesn't appear in Aquinas' corpus and only rarely occurs in the Latin world in general, and so unfortunately isn't useful here.

of the phrases, for instance, have a connection to '*natura*': *secundem naturam* and *contra naturam*; *ex principiis sua naturae*; *per modam naturae*; *contrariantur iuri naturalem; praeter naturalem consuetudinem*; *contra naturalem inclinationem*; and simply *naturale*.[12] Other words and phrases make no mention of *natura*, but instead involve words closely associated with the foregoing list: *conveniens* and *inconveniens*; *per se sequiter*; *quod pertinet ad habitum virtutis; ordinate* and *inordinate*; and *habet inquantum est.*[13] *Nota Bene* that the claim here isn't that Aquinas means the same thing in his use of these different words and phrases. Instead, the idea is that lessons can be drawn on Aquinas' view of the relation between normality and regularity by treating these phrases as indicating *various species* of normality.

There is a view that shows up in the generations after Aquinas' death according to which God can impose order on the universe apart from the regularities that spring from the natures God imbues in natural entities. This is fairly clearly not Aquinas' view: he thinks that whatever law-like regularities there are in the universe spring from the actions and passions that characterize things of a given nature. Normalities cannot be grounded in divine fiat. So whatever is normal is what attaches to a nature—there is no other grounding for a normality in Aquinas' universe.

With the vocabulary on the table, we can begin drawing some conclusions from the corpus as a whole. Aquinas appears repeatedly to rely on the following two claims in his arguments:

(1) Frequency is (often conclusive) evidence that there is a corresponding normality (a per se cause).[14]

---

[12] Cf. for examples, respectively: *In Physics* 2 Lectio 14, 267 and also *Summa Theologiae* I-II, q. 84 a. 4 co.; *Summa Theologiae* I-II, q. 71 a. 2 arg. 3; *Summa Theologiae* III, q. 57 a. 4 ad. 4; *Summa Theologiae* I-II, q. 51 a. 3 co.; *Summa Theologiae* II-II, q. 60 a. 5 ad 2; ST 1.105.a7.ad2 and also *Summa Contra Gentiles* 3b.99.9; *Summa Theologiae* I, q. 63 a. 9 co; and finally *Summa Contra Gentiles* 1.44.7 and also *In Metaphysics,* Prologue.

[13] Cf. for examples, respectively: *In Posterior Analytics,* Foreword; *Summa Contra Gentiles* 1.4.5 (for a use of *inconveniens* that is only arguably connected with abnormality); *Summa Theologiae* I-II, q. 20 a. 5 co; *Summa Theologiae* I-II, q. 51 a. 3 co.; *Sententia Metaphysicae*, lib. 1 l. 6 n. 1 for both *ordinatum* and *inordinatum*; and *Super Sent.*, lib. 2 d. 1 q. 1 a. 1 ad 3.

[14] This claim is made many times, but for a sampling, see *Summa Contra Gentiles* 2.24.4; *In Phys* 2,12,253-254; *Summa Contra Gentiles* 1.42.7,1.44.7; *In Physica* 2.8.214; *In Metaphysica* 1. Lect 6.107; *In Phys* 2,14,267.

(2) Infrequency is (often conclusive) evidence that there is no corresponding normality (that there is at most a per accidens cause).[15]

Given these two principles, there is still the question of what the link is between frequency and normality apart from the former's evidential status with respect to the latter. Aquinas could be committed to any one of the following (or different claims at different points in his corpus).

B1. Normality *is frequency*: what it means to be a normality is simply to happen most or a great deal of the time
B2. Normality and Frequency are linked in extension: they're different things, but the normal and the frequent come together to some degree.
    a. They *very often* come together.
    b. They *generically* or "as a rule" come together.
    c. The extensional link goes one way but not the other.
    d. They *always* come together.
    e. They *must* come together (necessary connection)
B3. Normality and Frequency have no principled extensional link in the actual world (and therefore no modal link).

Aquinas clearly doesn't accept B3, as will become clear. But what about B1?

Given a number of passages, we are forced to rule out B1 as an interpretation of Aquinas. Here is one such passage that can serve as a representative:

> I answer that more angels endured than sinned. Since sin is against their natural inclination, yet that which is against nature happens in few cases; for the effect of a nature follows from it either always or for the most part. [16]

Angels have sinned, according to Christian tradition; but at minimum it must have been a minority, if not a very small minority. Sinning isn't *normal* for angels—it is against their nature—and so it can't happen in a majority of cases. Instead, they will "stand firm" or "endure" against sin always or for the most part. Though there are perhaps clearer translations of "*ut in pluribus*," I've chosen the

[15] Again, see for a sampling: *Summa Contra Gentiles* 3a.5.2; *Summa Contra Gentiles* 3a.6.6-7; *In Phys* 2,8, 208 & 215; *In Metaphysica* 1, lectio 6, 107.

[16] *Summa Theologiae* I, q. 63 a. 9 corpus: Respondeo dicendum quod plures Angeli permanserunt quam peccaverunt. Quia peccatum est contra naturalem inclinationem, ea vero quae contra naturam fiunt, ut in paucioribus accidunt; natura enim consequitur suum effectum vel semper, vel ut in pluribus. (All translations are mine unless otherwise noted.

purposefully ambiguous "for the most part" to highlight the parallels between this passage and others like it.

Regardless of how we translate these terms though, the lesson is clear: Aquinas can't think that to be normal *is* to be frequent since he relates them in extension in this passage. This argument faces two potential worries. First, there's a hidden assumption in the inference: that reading phrases involving *natura* as 'normal' as we have is licensed. I think, though, that this is fairly safe. We've said at the outset that there are two uses of the word normal: a frequentist use and a more normative use. Anyone who disagrees with this claim has gotten off the bus long ago, and it'd be a very different conversation if I were trying to convince them. Instead, I will assume that there are two concepts here and that we are trying to understand the connection between normality as a somewhat normative concept and frequency. With that in mind, when Aquinas uses a phrase like "*contra naturalem inclinationem,*" it is difficult to see how we could read it as expressing a concept other than abnormality.

Second, perhaps the inference is still too quick: Aquinas could simply be claiming that there's an extensional relationship between the two without implying or even implicating that they are not simply the same thing. I might, for instance, say "Everyone in jail is incarcerated" without implying that being in jail and being incarcerated are different things at the level of intension. The apparent moral: claiming an extensional link between two terms doesn't necessarily commit one to the intensional inequivalence of the two terms.

Fair enough, but it would be uninformative to make such a claim. Since Aquinas makes many such claims, it becomes increasingly odd with each successive passage to think that Aquinas means to make a merely extensional point without also implicating that there is an intensional inequivalence between the two concepts. Thus, I think it's safe to conclude that for Aquinas, frequency and

normality are different concepts and therefore that he does not accept B1. Now the challenge is to discern what sort of extensional relation he thinks holds between them.

Let us recall the remaining possible interpretations:

B2a. They *very often* come together.
B2b. They *generically* or "as a rule" come together.
B2c. The extensional link goes one way but not the other.
B2d. They *always* come together.
B2e. They *must* come together.

Aquinas certainly would affirm B2a given passages we've already seen and will see, but that's far too weak a claim with which to characterize Aquinas' considered view. Let's get some passages on the table before proceeding:

> (In Phys 2.9) For every *per se* cause produces its effect either always or for the most part.[17]
>
> (In Phys 2.13) Therefore, those things which happen always or for the most part happen for the sake of an end. Now whatever happens according to nature happens either in every instance or in most instances, as even they admitted. Therefore, whatever happens by nature happens for the sake of something.[18]
>
> (In Phys 2.8) Since those things which come to be on purpose or by nature come to be on account of an end. And those that come to be on account of an end come to be either always or frequently.[19]
>
> (ST 1-II 71) Nothing that is counter to a nature is found in things of that nature for the most part.[20]

[17] *In Phys* 2.L9.220: "omnis enim causa per se producit effectum suum vel semper, vel ut frequenter."

[18] *In Phys*. 2.L13.256 "ergo ea quae fiunt semper vel frequenter, fiunt propter aliquid. Sed omnia quae fiunt secundum naturam, fiunt vel semper vel frequenter, sicut etiam ipsi confitebantur: ergo omnia quae fiunt a natura, fiunt propter aliquid."

[19] *In Phys* 2.8.213 quia ea quae fiunt a proposito vel a natura, fiunt propter finem; et ea quae fiunt propter finem, fiunt semper aut frequenter.

[20] *Summa Theologiae* I-II, q. 71 a. 2 arg. 3. Praeterea, nihil quod est contra naturam, invenitur in habentibus illam naturam ut in pluribus.
It's unclear how much authority to put in such a passage: it's an objection the premise to which Aquinas doesn't deny in his response (in fact, he takes it on board in that he offers an explanation the necessity of which relies on the veracity of the principle). The use of the quantifier '*nihil*' here, though, is instructive.

> (In Meta 1.5) Yet, this is shown to be false, in that good dispositions of this kind come about either always or for the most part. Yet those which come about by chance or fortune are not so always, nor frequently, but rarely.[21]
>
> (ST I 23) In response to the third objection it should be said that the good that is proportioned to the common state of nature happens for the most part and what is missing that good happens for the less part. But the good that is outside of the common state of nature happens for the less part and the defect of that good happens for the most part.[22]

This survey gives us a sense of the range of claims Aquinas makes throughout the many works of his corpus. They aren't clearly unified or consistent, but perhaps we can still draw some lessons therefrom.

Starting at the extremes, we can narrow down to what Aquinas appears to think. First, it appears that we can rule out B2e. B2e is the claim that the regular and the normal are not only connected in extension in the actual world, but that they are *inevitably* connected even if the world were quite different (perhaps within certain limits like an accessibility constraint or the like). God not only chose to create a world with such harmony that the normal is frequent and the frequent is normal but also would have created such a world even if God had chosen to make a very different world than the one God in fact chose to make.

B2e is simply too strong. *Summa Contra Gentiles* 3.99 in particular (see below) speaks against this interpretation by introducing the possibility that a natural cause might upset the natural order. If there is a clear relationship between normality and frequency, it isn't one of necessity.

---

[21] Met.1.5.99: Quod tamen patet etiam esse falsum, per hoc quod huiusmodi bonae dispositiones inveniuntur vel semper, vel in maiori parte. Ea autem quae sunt a casu vel a fortuna, non sunt sicut semper, sed nec sicut frequenter, sed ut raro.

[22] *Summa Theologiae* I, q. 23 a. 7 ad 3. Ad tertium dicendum quod bonum proportionatum communi statui naturae, accidit ut in pluribus; et defectus ab hoc bono, ut in paucioribus. Sed bonum quod excedit communem statum naturae, invenitur ut in paucioribus; et defectus ab hoc bono, ut in pluribus.

B2c is also ruled out, given that Aquinas makes claims in *either* direction at different times to suit different purposes: Aquinas appears to think that the extensional link between normality and regularity is bidirectional.

Grammatically speaking, we seem to find passages that speak in favor of both the generic claim B2b and the universal claim B2d (see above for complete citations):

Generic Passages: In Phys 2.8, In Meta 1.5, ST I 23.

Universally Quantified Passages: In Phys 2.9, In Phys 2.13, ST 1-II 71.

In light of this evidence, and the fact that the grammar of the universal claims doesn't appear to be accidental in context, we have a choice to make: either leave our interpretation ambiguous between the two or decide in favor of one or the other. As I noted, it's not clear that we can get away with claiming that Aquinas didn't intend to make universal claims in the passages above, since they are explicitly universally quantified and since the universality of the claims are important for the relevant inferences (for instance, In Phys 2.9 is making an analytic-conceptual argument and so appears to *need* a universal claim). If we can't overlook the grammar of the universal claims, then we are either stuck with a sort of inconsistency between the generic and universal claims, or we are forced to overlook the weaker claims in favor of the stronger claims.

As it turns out, one passage in particular speaks *against* the universal claim, and so the situation is further complicated:

> (SCG 3.99) Therefore, if through some created power it can happen that the natural order is changed from what happens for the most part to what is rare without any change in divine providence; much more can the divine power sometimes make something, without prejudice to its providence, outside of the natural order imbued in things by God. Indeed, God does this very thing sometimes in order to manifest God's Power.[23]

[23] *Summa Contra Gentiles* 3b.99.9: Si ergo per aliquam virtutem creatam fieri potest ut ordo naturalis mutetur ab eo quod est frequenter ad id quod est raro, absque mutatione providentiae divinae; multo magis divina virtus quandoque aliquid facere potest, sine suae providentiae praeiudicio, praeter ordinem naturalibus inditum rebus a Deo. Hoc enim ipsum ad suae virtutis manifestationem facit interdum.

This passage is central in understanding Aquinas on this topic. Aquinas makes the claim here that a natural cause can prevent a normality from being regular. There are two ways of understanding such a claim. One centers on examples such as the case where a species becomes endangered and so the birth of an animal which was once quite plentiful becomes rare. The other centers on examples such as the case where for a time something like building houses becomes rare because of a war. A natural cause, in each case changes what is regular and does so either, in the first case, permanently or, in the second case, temporarily. Since Aquinas has virtually no discussion of the engenderment of species or of extinction, it seems that we can interpret him as intending temporary disruptions in the otherwise stable normal order like wars, famines, droughts, etc. The upshot, though, is that Aquinas appears to be explicitly claiming that the normal *need not* be regular all of the time and therefore that the normal is not *always* regular.

The *Summa Contra Gentiles* comes just before the *Sententia Super Physicam* in Aquinas' life,[24] and so one might think that there was an evolution toward universal thinking and away from mere generic thinking between the two works.[25] Nevertheless, it's difficult to make sense of how he might admit the possibility of something in the *Summa Contra Gentiles* and then back off of that claim just a few years later without comment or reason. With that in mind, I am proposing that we take this passage seriously.

There is also some scope ambiguity at play here. If Aquinas holds a universal view, is his considered view about the relationship between regularity and normality Wide or Narrow?

(Wide) All normalities are regular and all regularities are normal.

(Narrow) The normal is always regular and the regular is always normal.

[24] See Torrell (2005) for a detailed discussion of the chronology of Aquinas' works.

[25] Speaking in favor of this view is the fact that the *Summa Theologiae* makes a universal claim and is more contemporary with the *Super Physicam*.

The passage above speaks against Narrow in that he seems to claim that natural causes can make the normal irregular without violence to the normal order of things. In fact, it also seems to speak against a hybrid reading:

(Hybrid) All normalities are always regular and all regularities are always normal.

Because any narrow scope reading appears to be ruled out by SCG 3.99, we are left with Wide as our only plausible universal reading of Aquinas on the relation between regularity and normality.

We still have two live readings of Aquinas on the extensional link between normality and regularity, neither of which we seem to be able to rule out conclusively:

B2b. Normality and regularity *generically* or "as a rule" come together.

B2d. Normality and regularity *always* come together.

We won't attempt to adjudicate further. But it will be important moving forward to note that the generic reading is meant to have extensional import: what generically comes together, comes together in a vast majority of cases barring unanticipated interventions. Aquinas appears to hold one of these two principles or is himself undecided between the two.

3. Problematic I

Aquinas appears to think either that the normal is always frequent and the frequent is always normal; or perhaps that the normal is generically regular and the regular is generically normal. This sort of view, on whichever interpretation, faces two clear problems: one external to his view and another internal to it. We'll address the internal problem in the next section. The external problem is simply that there appear to be clear counterexamples, many of which would have been part of Aquinas' everyday life. Here are a few such counterexamples:

C1. Olives, Almonds, Walnuts, etc. all don't grow into their respective trees a vast majority of the time, but that's what normally follows from being an olive, almond, or walnut.

C2. Humans are overwhelmingly likely to be vicious, ignorant, and to follow their baser desires., but they are normally supposed to be virtuous, wise, and rational—that is, these things follow from human nature.[26]
C3. Infant sea turtles are predated in overwhelming numbers before reaching the ocean, but they normally live to be 80 years old or more.[27]
C4. If dogs should suddenly be subject to a worldwide epidemic disease that caused gangrene in their limbs, it wouldn't any longer be common to see a fully intact dog, but arguably would still be true that dogs normally have four legs.

Aquinas wouldn't have been aware of C3, one imagines, but he was evidently aware of C2 and would have almost certainly been aware of C1. C4 certainly doesn't seem to have occurred to him, even though it's obviously a possibility even on a pseudo-temporal conception of modality like Aquinas arguably had. Nevertheless, the problem in each case is that the world doesn't appear to obey Aquinas' principles: the normal is sometimes not frequent and the frequent is sometimes not normal.

Each case presents a different sort of difficulty and so the solution to each case is somewhat distinct. The solution to C1 won't convince a modern skeptic that Aquinas' principles are consistent with the world, but nevertheless would likely be Aquinas' response to the case. Instead of accepting the premise of the apparent counterexample, Aquinas would simply deny that an almond is the sort of thing that a normality can attach to. Normalities for Aquinas clearly attach to natures and to natures alone, and an almond isn't a nature, but instead is a means by which a nature reproduces. The almond tree uses its fruit to reproduce—an almond doesn't use an almond tree to reproduce itself. We might call the general strategy employed here *Respecification*. Aquinas has principled reasons for thinking that we ought to respecify cases like this to ensure that the regularity we're interested in

[26] See the following passages for each respective claim; Aquinas appears to be fully aware of this fact. Humans are vicious: *Super Sent.*, lib. 1 d. 39 q. 2 a. 2 ad 4; humans are ignorant and misguided: *Summa Contra Gentiles* 1.3.5 and 1.4.5; and humans follow their baser desires: *Summa Contra Gentiles* 3a.6.8 and cf. Super Sent., lib. 2 d. 1 q. 1 a. 1 ad 3 for a more complete discussion.

[27] This example is discussed often in the literatures on generics and *ceteris paribus* laws. Strößner, for instance, claims that the sentence "turtles live to be 100 years or more old" indicates a potential age turtles can reach rather than an actual fact about the history of the species (2015, p. 803). Schurz (2009b) suggests that it's a claim about having the "capacity" to live that long.

is one attaching to a nature and not to some accidental entity or mere instrument of a non-accidental nature. In this case 'reproductive tool of an almond tree' isn't a nature and so can't have normalities attach to it.[28]

John Duns Scotus, in the generation after Aquinas, wouldn't likely have accepted this strategy, and much less need someone today. It's hard to think about an almond as having anything other than a nature. They're produced by a reliable process to reliably have very similar chemical and macro structures and thereby to have a characteristic set of actions and passions—what we'd now call dispositions. It is reasons like this that would have moved Scotus to posit a nature or 'substantial form' in this case—indeed he's not allergic to positing natures *all throughout* the natural world—and it is similar reasons which make Aquinas' solution unattractive to a modern mind.

Nevertheless, there is some intuitive pull to at least the *priority* claim that Aquinas is making: Almonds don't *use* almond trees. Just the opposite, in fact, Almond trees have the disposition to create loads of almonds in the hopes that many will sprout, fewer will grow into saplings, and even fewer into adult trees. With that in mind, it feels less strange to say that it is normal for countless almonds to fail to grow into almond trees.

One of Aquinas' direct defenses against the challenge posed by C2 is found in the *Summa Theologiae, prima secundae*:

> In response to the third objection it should be said that there are two natures in a human being: rational and sensitive. And since it is through a sense operation that a human comes to an act of reason, therefore there are more who follow the inclinations of sense than the order of reason. For there are more who arrive at the beginning of a thing that those who reach the end. Yet vice and sin come to be in

[28] A related but quite distinct strategy is employed by Thompson (2008) in dealing with the problem that the statement "Mayflys breed shortly before dying" is apparently true even though a vast majority of mayflys never get to mate. Thompson's response is to offer a shift in understanding the content of the statement: instead of 'mayflys' being the subject, one should treat "mayfly breeding" as the subject and interpret the statement to mean "mayfly breeding happens just before the death of the breeding partners."

> humans from this: that they follow the inclination of their sensitive natures against the order of reason.[29]

Humans sin as a matter of course, so the worry goes. "Vice is found in human beings for the most part,"[30] claims the objection, and Aquinas never denies this premise in his response. Instead he appeals to the multifaceted nature of humankind and make a sort of apology for the viciousness of humanity: we are sensuous creatures and so often are wont to ignore our rational intentions. But just above, in the corpus, Aquinas claimed quite clearly that vice and sin are against the nature of humanity. On the interpretation schema I've proposed, this means that they're abnormal. So, Aquinas accepts that in some sense what is abnormal is overwhelmingly common, but he doesn't stop there in defeat. Instead, he points out that human beings are rational *and* sensuous, and being sensuous means that we're going to irrationally follow our baser desires much of the time. It will be quite few who are able to act rationally and so to always or for the most part act with virtue.

This is an intriguing response strategy: accept that something abnormal happens regularly but then explain the regularity by appeal to a different nature. It does, however, have some awkward implications for Aquinas' broader thought about the relation between frequency or regularity and normality. For one, we might have thought that whatever doesn't happen always or for the most part *can't* be normal. But here, it seems that Aquinas has bitten the bullet and accepted that something irregular can be normal so long as its irregularity is explained by its rivalry with the regularity attaching to a different nature. More concretely, Aquinas is claiming that even though it

---

[29] *Summa Theologiae* I-II, q. 71 a. 2 ad. 3: Ad tertium dicendum quod in homine est duplex natura, scilicet rationalis et sensitiva. Et quia per operationem sensus homo pervenit ad actus rationis, ideo plures sequuntur inclinationes naturae sensitivae quam ordinem rationis, plures enim sunt qui assequuntur principium rei, quam qui ad consummationem perveniunt. Ex hoc autem vitia et peccata in hominibus proveniunt, quod sequuntur inclinationem naturae sensitivae contra ordinem rationis

[30] Article 2, Corpus: "Sed vitia inveniuntur in hominibus ut in pluribus." See also, *Super Sent* 2 d. 1 q. 1 a. 1 ad 3, where Aquinas makes a similar claim about the connection between the specification of an individual and the regularity of a particular effect and then appears to posit a sort of justification of this move. Here's the justification: Si est unius naturae, sicut in Angelis, sic constat quod in pluribus est consecuta operatio recta secundum convenientiam naturae ipsorum, et peccatum ipsorum fuit ut in paucioribus. Si autem est plurium naturarum, sicut est homo, qui est compositus ex natura intellectuali et sensitiva, potest considerari dupliciter.

doesn't happen often that humans should be virtuous or righteous, it is still normal for humans *qua* rational beings because the rarity of it can be explained by appeal to a different competing regularity: humans are regularly subject to sin and vice because they are *sensuous* beings. This is an intuitive solution, but is incompatible with the claim that infrequency is conclusive evidence for abnormality. Aquinas does, however, use the infrequency of phenomena as conclusive evidence for its non-normality quite often. It suffices to say, for now, that this appears to be an inconsistency in Aquinas' thought. Similar themes will again arise in our discussion of the second problematic.

In answering the challenge presented by C3, it's tempting to take a similar tact to that provided in our discussion of C1: why not simply think that a mother turtle is trying to produce one healthy offspring and in doing so is forced by ecological considerations to produce 8 at a time? But this simply won't work, given that a baby turtle is a substance if anything is and therefore has a substantial form if anything does. How, then, do we answer the challenge on behalf of Aquinas? We'll offer a potential solution to C3 in response to C4.

There is an interesting question as to how Aquinas can conceive of C4. How might someone with as underdeveloped a notion of modality as Aquinas conceive of such a counterfactual? Regardless of how we interpret Aquinas on modality, luckily C4 can be read as a statement of future contingent possibility. Read that way, we avoid getting lost among the modal weeds and can stay in comfortable territory for someone with Aquinas' conceptual resources. Aquinas himself, in fact, admits that such a possibility might be actual in a passage quoted earlier:

> Therefore, if through some created power it can happen that the natural order is changed from what happens for the most part to what is rare without any change in divine providence; much more can the divine power sometimes make something, without prejudice to its providence, outside of the natural order imbued in things by

> God. Indeed, God does this very thing sometimes in order to manifest God's Power.[31]

If it can happen that a disease can cause widespread limb loss in the worldwide dog population so as to make rare that a dog should have four limbs when it used to be for the most part without us thinking, "oh, I guess God didn't make a universe where dogs normally have four limbs," then it's no worry that God should more directly perform miracles. Aquinas appears to be happy to admit that there is such a possibility and to do so in a passage where the very problem of exceptions to what normally happens is at issue. There can be no pleading ignorance, it seems.

The solution to both C3 and C4 will be to appeal to an underappreciated but central aspect of Aquinas' natural philosophy—and his philosophy more broadly. When an answer is not available at one frame of reference, we must shift our frame of reference until the answer becomes available. Why is that musician building a house? Well, because being a musician is just her hobby, she's a carpenter during the day. Why is that *motive being* sitting down right now rather than running or jumping or laying? Well, he's also an appetitive being and he's got the appetite for sitting at the moment. Why is it that I'm always picked last for softball? Well, Stacy was captain last week and she has a grudge against you, and Gurpreet was captain the week before and she is best friends with the other three players that are equal in skill to you, and the week before you came out to play too late to get picked sooner. When we are stuck with a *per accidens* description or explanation of a phenomenon or when an explanation is otherwise mysterious or unavailable, it means we are looking at the phenomenon from the wrong frame of reference: shift the frame of reference and all will become clear.

[31] Si ergo per aliquam virtutem creatam fieri potest ut ordo naturalis mutetur ab eo quod est frequenter ad id quod est raro, absque mutatione providentiae divinae; multo magis divina virtus quandoque aliquid facere potest, sine suae providentiae praeiudicio, praeter ordinem naturalibus inditum rebus a Deo. Hoc enim ipsum ad suae virtutis manifestationem facit interdum. *Summa Contra Gentiles* 3b.99.9s

The problem in both C3 and C4 is that what appears to be normal is nevertheless quite rare and Aquinas seems to think that this isn't supposed to happen. The block quote above indicates that Aquinas is comfortable admitting that there are exceptions to these rules and not only in the case of miracles, but it still seems odd to say that something as baked into the form of life of a sea turtle as their hatching and travelling to the shore is somehow a mere exception to God's otherwise orderly providence. Perhaps there's a single solution to both worries, and perhaps that solution appeals to frame dependence.

First, what do I mean by 'Frame Dependence"? Many elements of Aquinas' natural philosophy are relative to a particular frame of reference (something I've elsewhere called an 'organization'). See, for instance, the following telling passage:

> Therefore, Avicenna said that in those things which are indeterminate, things come to be by fortune, like things that happen in a minority of cases. And it is no objection to this that we don't say that it's fortune that Socrates sits, given that it is indeterminate. Since—allowing that it is indeterminate with respect to Socrates' motive power—it is not indeterminate with respect to his appetitive power, which determines tends toward the one thing. And if something should happen outside of that thing, then it will be called 'fortune.'[32]

Here, we see Aquinas defending a principle of Avicenna by appeal to frame dependence. Socrates, relative to one frame, is indeterminately sitting because we've got little in the way of explanatory information available about why he should be sitting rather than standing. Relative to another frame, though, it is determinate that he should be sitting—those who want to sit do so and they do so non-accidentally. Accidentality and "per seity" is relative to the frame of reference or the context with respect to which we specify and therefore understand a thing or process. Because so much in

[32] Avicenna ergo dixit quod in his quae sunt ad utrumlibet, contingit aliquid esse a fortuna, sicut ea quae sunt in minori parte. Nec obstat quod non dicitur esse a fortuna quod Socrates sedeat, cum hoc sit ad utrumlibet: quia licet hoc sit ad utrumlibet respectu potentiae motivae, non tamen est ad utrumlibet respectu potentiae appetitivae, quae determinate tendit in unum; praeter quam si aliquid accideret, diceretur esse fortuitum. *In Phys* 2, Lectio 8, 209.

Aquinas' worldview hangs on the distinction between the per se and the per accidens—most notably evaluations of normality and abnormality—a great deal hangs on the choice of frame of reference.

Aquinas is also comfortable relegating the ultimate explanation of some regularity to one frame of reference while admitting that it in some sense shouldn't be regular with respect to another frame. Though the biology and associated anthropology here is—to put it lightly—morally problematic, the passage is nevertheless quite instructive as to how Aquinas is thinking about causes and frames of reference. Here, one cause considered in itself would normally produce one type of effect, but since it regularly doesn't produce that effect and instead produces a different sort of effect, Aquinas searches elsewhere for a cause of the apparent deficiency.

> Differentiation of the sexes must be attributed to celestial causes. Our reason for saying this is as follows: Every agent tends to form to its own likeness, as far as possible, that which is passive in its respect. Accordingly, the active principle in the male seed always tends toward the generation of a male offspring, which is more perfect than the female. From this it follows that conception of female offspring is something of an accident in the order of nature-in so far, at least, as it is not the result of the natural causality of the particular agent. Therefore, if there were no other natural influence at work tending toward the conception of female offspring, such conception would be wholly outside the design of nature, as is the case with what we call "monstrous" births. And so it is said that, although the conception of female offspring is not the natural result of the efficient causality of the particular nature at work—for which reason the female is sometimes spoken of as an "accidental male"—nevertheless, the conception of female offspring is the natural result of universal nature; that is, it is due to the influence of a heavenly body, as Avicenna suggests.[33]

If something is accidental and therefore would be irregular or infrequent with respect to one efficient cause, but nevertheless is regular or frequent, then we must look for another efficient cause which is securing the regularity in question. In this case, the single seed theory of reproduction leads

[33] *QDdV* 5.9.ad9 Ad nonum dicendum, quod oportet sexuum diversitatem in aliquas causas caelestes reducere. Omne enim agens intendit assimilare sibi patiens, secundum quod potest; unde vis activa quae est in semine maris, intendit conceptum semper ducere ad sexum masculinum, qui perfectior est; unde sexus femineus accidit praeter intentionem naturae particularis agentis. Nisi ergo esset aliqua virtus quae intenderet femineum sexum, generatio feminae esset omnino a casu, sicut et aliorum monstrorum; et ideo dicitur, quod quamvis sit praeter intentionem naturae particularis, ratione cuius dicitur femina mas occasionatus, tamen de intentione est naturae universalis, quae est vis corporis caelestis, ut Avicenna dicit.

Aquinas to think that a male efficient cause of reproduction normally produces male offspring; but the universal order as secured by the efficient influence of a heavenly body dictates that the male efficient cause should be partially ineffective about half of the time.

From the perspective of the sea turtle itself more or less isolated from its environment, it can't clearly be called normal that it should be predated. The snuffing out of a natural thing earlier than its longer-lived conspecifics by an order of magnitude cannot, intuitively, be normal for that kind of thing. If we, however, "zoom out" to the ecological perspective and see the form of life of the sea turtle as having been shaped by eons of evolution in response to the looming threat of predation, then suddenly it doesn't seem as clearly normal for them to survive *en masse*. It's normal, in fact, for a great many of them to die in infancy just as it is normal for some proportion of all prey populations to be predated. Moreover, it feels odd to make normality claims *without* taking into account this ecological perspective. It would be analogous to claim that it's normal for humans to live to 170, but we are just faced with all of this interference and radiation and damage and infestation, etc. Without all of that, after all, we'd live to be much older! This, though, simply doesn't sound true. Of course, things would go differently in a truly ideal environment for human life, but that's not the environment we evolved with nor is it a "normal" environment. Nevertheless, we wouldn't want to say, I take it, that sea turtles, given the pressures of their normal environments, are normally predated or, even worse, that "it's normal for sea turtles to be predated". Instead, it feels like we must use a frequency word like 'regularly' or 'quite often' to make such a sentence ring true to our ears.

The example of the sea turtles is difficult to make sense of even with the tool of frame dependence at our disposal. Ultimately, it appears that what is normal in the non-statistical sense—living a long life—simply isn't what actually happens in most cases. There doesn't appear to be a turtle-centric normality that can map onto the regularity in question. We *can* however, shift frames

so as to decenter the turtles and in doing so perhaps find a better fit with the regularity. It's normal for sea gulls to eat small animals on beaches. They predate crabs, baby sea turtles, and so on quite regularly. It might not be normal for the animals to be eaten, but it is—it seems—normal for the sea gulls to eat small animals. The generic "sea gulls eat baby sea turtles," after all, seems clearly to be true. So, we were able to find a normality which attaches to the regularity in question; we simply had to place the sea gulls at the center of the frame rather than at the periphery and suddenly the regularity comes into focus.

A promising principle for selecting the salient property—predating vs. being predated—is to select the property that is *active* or is an *activity*. Regularities which arise from background conditions—like the sea turtle's beach being replete with predators—are to be explained by appeal to the activity of some actor in those background conditions. Alternatively, regularities which result from the activities of an organism or other active process or object are to be explained by appeal to those very activities. Of course, there is much more to say, and these distinctions are far from clear and unproblematic, but perhaps a principled account can be developed with effort.

Similarly, in C4, when we shift frames to an epidemiological frame—where the disease and not the carriers takes center stage—it feels more natural to say something like

(3) The disease normally causes limb loss in roughly 90% of the population.

It might even be okay to say (4):

(4) It's normal these days for dogs to get disease X and lose a limb.

Furthermore, the frame relativity allows us to say different things with respect to different frames: (4) makes sense relative to an epidemiological or pathogenic frame of reference, whereas (5) makes sense relative to a frame that centers an individual dog or the natural kind dog.

(5) It still isn't normal for a dog to get such a horrible disease.

This appeal to different frames of reference may alleviate some of the strain from Problematic I, but, as we'll see, it potentially creates problems of its own for Aquinas.

4. Problematic II

The second problem for Aquinas' view on the connection between normality and frequency is that it appears to be incompatible with other aspects of his worldview. In particular, the "frame dependence" of his natural philosophy appears to make it all but impossible that the frequent should always be normal or that the normal should all be frequent; nor that the infrequent should be abnormal or the abnormal infrequent.

Having introduced frame dependence in the previous section, we can begin to see why it might pose a problem for Aquinas' claims about the links between normality and frequency. It is normal for water to move downwards according to the elemental physics that Aquinas inherits. Whether it is abnormal or merely neutral that it should be subject to violent motion which causes it to move laterally is an interesting question, but either way it seems safe to say that it is abnormal for it *qua water* to move upwards. But if we shift a tree to the center of our ontology, that water moves upwards is necessary for it to survive and so decidedly normal for it qua nutrient for the tree. Capillary action forces water upwards through the vessels in a tree so that water and the minerals it carries can reach the outermost limbs. The water is doing something at once normal and abnormal.

With this in mind, it becomes easy to see how Aquinas' claim that whatever is frequent is normal and whatever is normal is frequent—even if it should be weakened or qualified—is under threat. How can it be that something can be overwhelmingly likely to do two different things? Furthermore, if we just consider water *qua water*, then it seems to be overwhelmingly likely not to do what is normal for it: to move downwards. Aquinas himself seems to be aware of this sort of phenomenon,

claiming in places that the deficiency of an efficient cause and the inconvenience of a material cause are commonplace.[34]

The first thing to note in defense of Thomas is that when he talks about frequencies he seems almost exclusively to be talking about regularities: one thing following another thing with some amount of frequency or often one thing happening at fixed or predictable intervals. We can even restrict the candidate regularities further by requiring that a property be regular *for members of a kind* before we'll accept it as normal. This only gets us so far, though, because a sea gull attacking a baby sea turtle sadly results in its death in an overwhelming majority of cases (one presumes)—the death of the sea turtle follows the cause, but the death of the sea turtle isn't clearly normal. It's tempting to think that one could develop principles which would get around such problems. In restricting what regularities can count as normalities, we might make some progress, but will only ever solve the "whatever is frequent is normal" problem since we'll be restricting the domain of the universal or perhaps generic claim. But we can't make progress on the other problem in this way—the problem that Aquinas seems to claim something like "whatever is normal is frequent" in the face of evidence to the contrary.

In attempting to solve the latter problem, the parallel strategy would again be domain restriction: restrict what sorts of normalities imply regularity and therefore save the principle that whatever is normal is frequent. It's unclear, though, how this might work. How might we, in a principled manner, restrict the normalities which are supposed to imply regularities? It's not at all clear that such a restriction can be effected and less clear still that it can be motivated.

Never mind all that, though, since the way that Aquinas would defend against the problem of frame dependence as a whole will be by appeal to his unitarianism about substantial form. Aquinas holds the somewhat counterintuitive view that there is only ever one substantial form governing the

[34] Cf. *Summa Contra Gentiles* 3b.99.9; *QDdV* 5.9.ad1; and S*ententia Metaphysicae*, lib. 6 l. 3 n. 22.

changes and activities of a substance *and all of its integral parts*. Water, becoming an integral part of a tree when the tree absorbs it, loses the substantial form of water and gains a form which is included in the substantial form of the tree as a whole. There is no problem with the water in a tree moving up through capillary action if it ceases to be substantially water when it becomes part of the dynamics of the tree. If, instead of being water, it is now a *nutrient* or the like, then its activity of moving up the tree through capillaries is perfectly normal; in fact, it is what is required of it insofar as it is what it substantially is.

Attached to this unitarianism is an overall view of the harmony of the natural world, which is secured by a provident and omnipotent creator. If God creates the world and in so doing has cognizance of even the minutest of accidental happenings therein, then God will have the choice between a world wherein things are rivalrous and essentially opposed or a world wherein things are created to harmonize together in a global order. Again, there will be the choice between harmony between *kinds*, where no kinds are essentially opposed to one another[35] and harmony between *individuals*, where no individuals are essentially opposed to one another.

Aquinas appears to hold that the world has been harmonized in that for the most part we find things have a sort of harmony:

> Moreover, according to the Philosopher, "it is the office of a wise man to set things in order." For things can be ordered only by knowing their relation and proportion to one another, and to something higher, which is their end; for the order of certain things to one another is for the sake of their order to an end. But only a being endowed with intellect is capable of knowing the mutual relations and proportions of things; and to judge of certain things by the highest cause is the prerogative of wisdom. All ordering, therefore, is necessarily effected by means of the wisdom of a being endowed with intelligence. Even so, in the world of the mechanical arts, the planners of buildings are called the wise men of their craft. Now, the things produced by God have a mutual order among themselves which is not fortuitous, since this order is observed always or for the most part. That God brought things

[35] i.e. where the existence of one kind interferes with the existence of the other. So predator-prey relations between kinds do not count, since they are in some sense symbiotic: each needs the other at the level of ecology.

> into being by ordering them is thus evident. Therefore, God brought things into being by His wisdom.[36]

There are many things in the world, but we find among them a sort of balance and harmony. The wolf doesn't exist for the destruction of the mule deer, but instead relies on the continuance of the deer population to supply it with prey. "If the wolf had it its way", in other words, the world would be much as it is lest its offspring should starve. Aquinas uses this observation toward a theological end, but his argument is not our focus here. Instead, the *datum* is our concern: the world isn't marked by essential opposition, but by mutual interdependence.

Is this a harmony between individuals or between kinds? As I see it, the only way to make sense of this and other passages in Aquinas is that individuals find themselves at odds with one another as a matter of course, but *kinds* are locked into harmonious mutually dependent relations with one another. The *individual* deer is not a fan of the wolf, for the wolf's being in the deer's surrounding territory is a threat to the deer's very existence. The deer population as a whole, though, relies on predation by the wolf population to keep it from overusing resources and upsetting the balance of its ecosystem.

This sort of interpretation helps us to make sense of other passages where it appears that there is supposed to be rivalry and competition in nature. The following passage from *de Veritate* is one example:

> Damascene intends to exclude from heavenly bodies only first causality or any causality which would introduce necessity into bodies here below. For, even if heavenly bodies always act in the same way, their effects are received in lower bodies

[36] *Summa Contra Gentiles* 2.24.4: Amplius. Secundum philosophum, in I Metaph., ordinare sapientis est: ordinatio enim aliquorum fieri non potest nisi per cognitionem habitudinis et proportionis ordinatorum ad invicem, et ad aliquid altius eius, quod est finis eorum; ordo enim aliquorum ad invicem est propter ordinem eorum ad finem. Cognoscere autem habitudines et proportiones aliquorum ad invicem est solius habentis intellectum; iudicare autem de aliquibus per causam altissimam sapientiae est. Et sic oportet quod omnis ordinatio per sapientiam alicuius intelligentis fiat. Unde et in mechanicis ordinatores aedificiorum sapientes illius artificii dicuntur. Res autem quae sunt a Deo productae, ordinem ad invicem habent non casualem, cum sit semper vel in pluribus. Et sic patet quod Deus res in esse produxit eas ordinando. Deus igitur per suam sapientiam res in esse produxit (Translation by James F. Anderson). Cf. also *Summa Contra Gentiles* 1.13.35.

> according to the manner of these lower bodies, which are frequently seen to be in a contrary state. Consequently, the forces exercised by the heavenly bodies are not always able to bring about their effects in the bodies here below, because a contrary disposition prevents them from doing so. This is why the Philosopher says that signs of storms and winds frequently appear, but the storms and winds do not take place because the contrary dispositions are stronger.[37]

If there were perfect harmony—if every cause found itself in ideal material conditions for the realization of its effect—then every cause would be a necessary cause. But the natural world is full of contingency, and this is because the world is harmonious not at the level of each individual, but instead at the level of kinds of things. The heavenly bodies are able to exert some contingent influence such that when something like a storm arises, it will be on account of that influence; yet they are not able to impose their effects on the sublunar in every case because—even though the oceans and atmosphere of Earth is the right *kind of matter* in which the heavenly bodies can produce tides and storms and winds—the natural material isn't in each and every instance of a harmonious disposition.

Why think, though, that the harmony between natural kinds, or between causes and material conditions, needs to lead to a principled connection between normality and regularity? The deer might have only needed a small amount of predation on occasion in order to maintain balance. The heavenly bodies might have only been able to produce winds as rarely as they in fact produce eclipses (and perhaps sporadically as well as rarely). In short, the normal might have only happened *irregularly* without doing damage to the claim that the world is marked by harmony between kinds.

[37] *QDdV* 5.9.ad1: Ad primum igitur dicendum, quod Damascenus intendit a corporibus caelestibus excludere respectu horum inferiorum causalitatem primam, vel etiam necessitatem inducentem. Corpora enim caelestia etsi semper eodem modo agant, eorum tamen effectus recipitur in inferioribus secundum modum inferiorum corporum, quae in contrariis dispositionibus frequenter inveniuntur; unde virtutes caelestes non semper inducunt effectus suos in his inferioribus propter impedimentum contrariae dispositionis. Et hoc est quod philosophus dicit in Lib. de somno et vigilia quod frequenter fiunt signa imbrium et ventorum, quae tamen non eveniunt propter contrarias dispositiones fortiores (Translated by Robert W. Mulligan, S.J.).

Again, we must understand the distinction between frequency and regularity. Eclipses are quite rare—they are infrequent—but they are very nearly perfectly *regular* in that they follow a predictable pattern. Let us explore two passages from the *Summa Theologiae* to further bring out this point:

> And such future effects can be foreknown and predicted with certitude by the consideration of their causes, as astronomers predict future eclipses. But some causes produce their effects not by necessity and always, but for the most part, failing but rarely. And through such causes it is possible to foreknow future effects—not with certitude, but by a kind of inference, as the astronomers can foreknow and predict wet and dry weather through the consideration of certain stars, and as the medics with health and death.[38]

We may not be able to know with certainty that something will happen in the future given its cause—indeed that is the very nature of a contingent cause—but this does not mean that we cannot know *at all* that the effect will happen. Furthermore, when the effect does happen, it will be because of its cause, but this still does not mean that it *will* happen. The relevant lesson here is that what is normal—what follows from a *per se* cause—will generically happen *regularly*: it will follow from its cause.

Aquinas faced two problematics. External to his worldview, it seems like regularity and normality come apart in the actual world and seem even more likely to come apart in terms of what is possible. But Aquinas seemed to think that the normal and the regular share an extensional covariance—at least generically if not universally. Aquinas has the resources to defend himself against this charge of inadequacy to the evidence, and we explored some of the responses he in fact gives and others he would likely give if faced with the problematic. Ultimately, he must appeal to the frame dependence inherent in his worldview to make sense of how something that is quite common

[38] *Summa Theologiae* II-II, q. 95 a. 1 co.: Et huiusmodi effectus futuri per certitudinem praenosci possunt et praenuntiari ex consideratione suarum causarum, sicut astrologi praenuntiant eclipses futuras. Quaedam vero causae producunt suos effectus non ex necessitate et semper, sed ut in pluribus, raro tamen deficiunt. Et per huiusmodi causas possunt praenosci futuri effectus, non quidem per certitudinem, sed per quandam coniecturam, sicut astrologi per considerationem stellarum quaedam praenoscere et praenuntiare possunt de pluviis et siccitatibus, et medici de sanitate vel morte (Translated by Fathers of the Dominican Province). Cf. also *Summa Theologiae* I, q. 86 a. 4 co.

might still be abnormal: it's only abnormal with respect to one frame or cause, but it is normal with respect to another. Nevertheless, there appear to be some inconsistencies between Aquinas' explicit claims and his appeals to frame dependence.

On that note, the problematic that occurs *within* Aquinas' belief system has to do with the relation between frame dependence and Aquinas' beliefs about the relation between normality and regularity. It appears that frame dependence makes it all but impossible that the normal should be frequent, since making the normal frequent with respect to one frame interferes with the normal being frequent with respect to another frame. Aquinas' unitarianism about substantial form and belief in the harmony of nature helps to alleviate the strain here, but ultimately there is a tension that deserved more attention from the *Doctor Angelicus*.

5. Normality and Regularity

The focus has been on making sense of Aquinas' considered view on the relationship between normality and regularity. What philosophical lessons, though, have we learned from this discussion?

First, there's a sense in which the predictability of the world stands in tension with the idea that the normal need not be regular. It seems that the regular and the normal *must* coincide to some considerable degree simply because of the relationship between the normal and the predictable or expectable.[39] There's an ambiguity between predictable in the sense that we could have predicted that someone in particular would lose if they bid on black 20 in roulette and in the sense that we can predict an eclipse. The probabilities are so vastly different, for one thing, but there's also a sense in which when one loses at roulette, it's a bit of an accident; whereas when an eclipse happens, it's no accident. The eclipse happens as a result of regular happenings in the solar system and has a set of causes that secure the particular outcome in question. In contrast, the roulette spin by its nature

[39] A number of philosophers have argued for this claim. Among them are (Pearl 1988; Schurz 1997, 2009b; Strößner 2015)

secures no particular outcome and so whether the marble lands in the black 20 slot or not is—idiomatically—a crap shoot.

If we focus on the eclipse sense of prediction, it seems that what we're doing is extrapolating trajectories and understanding natural causes in order to judge that an eclipse will happen in a certain way and at a certain time. We're doing something in some sense *normative* in that we're ignoring the possibility of massive collisions between the moon and an asteroid, which would throw off our prediction but would also be extraneous to our understanding of the moon's *normal* motion.[40] In predicting an eclipse, we're tacitly ruling out interfering factors and interference is a normative concept—it is parasitic on the notion of an activity or behavior with at minimum some end state towards which it is tending. We form expectations and make predictions on the basis of the normalities we understand to govern the world around us and so if we are to be any good at predicting and forming expectations, the abnormal must be somewhat rare.

Second, it seems surprisingly plausible to claim that if something is normal but nevertheless rare, then its rarity can be explained by appeal to another nature or normality. It's rare for sea turtles to grow past infancy, but it's not normal for them to die. The rarity of their longevity is owed, it turns out, to the nature of predatory birds, who feed on the helpless hatchlings. Without the predation of the birds, it would be quite odd to say that the turtles normally live to a long age if it turns out they just die in large numbers because of an internal cause like widespread heart defects or the like.

Third, the question at hand itself transforms a bit as we advance in our understanding thereof. Rather than being interested in the relationship between normality and bald frequency, which isn't terribly robust, we are instead interested in the relationship between normality and *regularity*. Once this shift is made, we can see how there might be a tighter relationship than once thought. If the

[40] As Wallace (1974) summarize one of Aquinas central claims about natural philosophy: in doing natural philosophy, we're considering what would have to be the case for the end state to take place rather than predicting going forward what will in fact happen.

question is what sort of extensional relationship obtains between that which follows from a natural cause with some regularity and that which *should* follow from a natural cause, then it becomes more natural to think that there is such a relationship.

Chapter Six: The Structure of Normality Judgments

Having already fixed a view about the content of normality judgments, one task remains before we've given a more or less complete account of what a normality judgment is. We need to go beyond *what it means* to make a normality judgment, toward an understanding of *what sort of structure* characterizes a normality judgment.

What do we mean by the "structure" of a judgment type? We could mean any number of things by such a phrase. We could be referring to the *propositional* structure of a judgment: what are the syntactic components and how are they grammatically related? Then again, we might mean the *conceptual* structure of such a judgment: what concepts are (perhaps necessarily) involved in such a judgment and what does it mean to competently make such a judgment? Yet again, we might also mean the *extensional* structure of a judgment: what range of cases is it supposed to cover and how does it divide those cases? Alternatively, we might mean the *modal* structure of such judgments: what sort of intensional profile does the judgment have, perhaps conceived of in terms of a range of truth values across a range of possible worlds or situations? Finally, we could mean the *inferential* structure of such a judgment: what inferences warrant these judgments and what inferences do they in turn warrant?

We'll focus our attention on the inferential structure of normality judgments.[1] First, I will highlight some basic inferences they warrant and otherwise participate in in section 1, and then discuss some more controversial inferences in section 2, before challenging what I take to be errors in the various literatures that bear on the topic at hand in section 3.

[1] Toulmin (1958, p. 166) argued that a particular sort of normality judgment, at least, is fundamentally "substantial" or material and so could not be formalized. I don't intend here to offer a formalization of normality inferences writ large. Far from it. Instead, the goal is to elucidate whatever general inference patterns appear to hold for normality judgments—however limited that class ends up being.

1. Basic Inferential Structure

Normality judgments are defined, at an intuitive level, by two inferential features I call insensitivity and tolerance. They remain true in spite of fluctuations of the actual prevalence of the predicated feature in the population of the kind being predicated of: they are therefore *insensitive* to changes in prevalence. Additionally, they remain true in spite of potentially many apparent counterinstances: they *tolerate* apparent counterexamples. We might call these "essential" or "constitutive" inferential features of normality judgments, but it's not yet clear that we should think about them this way. How we choose to define normality judgments as a family will determine whether these are constitutive of being a normality judgment and I'm not yet convinced that they ought to be defined by appeal to this shared inferential profile. It may, instead, be a mere intuitive heuristic for figuring out what sort of judgment we are dealing with in any given instance.

For instance, we might think that what actually *defines* this set of generalizations is that they involve to some degree a non-statistical concept of normality. If this is true, then it's their content and not their inferential structure that defines the class and so the inferential profile they have in common is an interesting commonality rather than a constitutive fact about them. Alternatively, we might think that this is truly a *family* of generalizations and so both the structure and the content *characterize* the family, but neither *defines* it. I'm tempted toward the latter thought, but I'll leave it as an open question for our purposes here.

1.1 Complete vs Incomplete Normality Judgments

As we discussed earlier in the dissertation, there are two sorts of normality judgment: complete and incomplete. The basic case seems to be one in which the outcome or feature ascribed to the kind is what *will* happen absent interference or abnormality. This is the class of normality judgment to which most *ceteris paribus* claims belong. There is, however, a derivative case evoked particularly clearly by a particular grammatical construction exemplified by "it is normal for piano players to be

able to play *Für Elise*." It is one among many normal outcomes of piano training that one should end up being able to play such a song, but it isn't the sort of thing that *ought* to happen or that *will* happen absent interference. Similarly, it's normal for cats to be calico, but this isn't the only normal fur color for cats. To make such a judgment complete would involve a rather long disjunctive predicate of the form seen in "it's normal for cats to be calico, or orange tabby, or white, or black, …" Complete normality judgments are marked by a distinctive inferential profile: only they, if any normality judgment does, have a normative inferential profile; and only they have the implication that there is no other normal way to be a member of the relevant kind in the relevant respect and thus that any individual that lacks the feature cited in the complete normality judgment is thereby abnormal in the relevant respect. So there is a distinction between normality judgments that are complete and those that lack such an inferential profile because they are merely stating one of potentially many normal outcomes in the relevant respect for a member of the relevant kind.

Schurz claims that an analysis of the concept of normality must require "overwhelming selection" of the supposed normal trait because "without this condition, one may end up with the result that both the possession and the non-possession of a trait could count as prototypical for a given species—which is conceptually incoherent, because it violates minimal logical principles of the concept of normality" (2001a, p. 489). Ströβner (2015, p. 795) makes a similar argument. Schurz and Ströβner, therefore, appear to be rejecting the possibility of incomplete normality judgments. I disagree. I think it can be normal for lions to both possess and lack manes. No conceptual incoherence here: different features can be normal for a kind of thing provided one does not take "it's normal for lions to have manes" and "it's normal for lions to lack manes" to entail "normal lions must have and lack manes." There is no such entailment for incomplete normality judgments.

Ströβner (2015) maps this distinction onto a grammatical distinction—that between statements of the form "It is normal for *k*'s to φ" and those of the form "*k*'s normally φ." As compelling as an

initial investigation of these two grammatical forms turns out to be, I am hesitant to hang too much weight on grammar. Instead, the real distinction seems to be based on the content and inferential structure of the judgments: certain normality judgments about a kind *k* consist of the claim that there is one (perhaps disjunctive) way of being a normal *k*, while other normality judgments don't have the implication that there are no other ways of being a normal *k*.

One reason not to hang too much on grammar is simply that there are more than two grammatical forms in which normality judgments can be expressed. One can use a bare plural generic sentence, for instance, for either sort of normality judgment and so grammar won't decide in favor of either. For this reason and others, I don't go with Ströβner in tying the distinction to a grammatical distinction.

1.2 Basic vs Derivative Normality Judgments

As discussed in previous chapters, one must make a distinction between basic and derived normality judgments. The truth of basic normality judgments is grounded in the explanatory link between the kind that serves as the subject of the judgment and the outcome ascribed to that kind in the judgment. A judgment like "dogs have four legs" is a basic normality judgment because the truth of this claim is grounded in the fact that being a dog explains why an individual would have four legs. Alternatively, the truth of 'derivative' or 'derived' normality judgment is grounded in the explanatory link between a member species and the outcome ascribed to the genus.

There are different inferential structures to basic and derivative normality judgments. For one, virtually all derivative normality judgments are incomplete normality judgments. They therefore will have the same inferential structure that any incomplete normality judgment will have insofar as they are incomplete normality judgments. Moreover, basic normality judgments will entail statements directly implied by the semantics of normality judgments, whereas derivative normality judgments won't. For instance, the claim "dogs have four legs" entails that "an individual's having four legs in

the normal way can be explained by appeal to its being a dog." The claim "mammals have four legs" apparently does not have such an entailment. There is more to say here, but this suffices to get the distinction on the table.

For the sake of efficiency, we'll move quickly through some basic inferences before settling into a more extensive critical discussion.

1.3 Respects and Individuals

Some basic inferences in the profile of normality judgments involve normality in a respect and the normality of individuals.

Asher and Morreau (1991) seem to suppose that an individual is a more normal member of a kind the more normal properties of that kind they possess. This is a tempting thought. It's not, however, so tempting to affirm a less relativized version of the same thought, which would be an alternate reading of their text: an individual is more normal *tout court* the more normal properties (of kinds they belong to) they possess. One might be skeptical of an idea such as "normal individual *tout court*" if one thinks that normality is *always* relative to a kind. There is no such thing as being a normal individual (insofar as "individual" isn't itself a name of a kind) if one can only ever be a normal wallaby, or a normal dentist, or a normal employee whose employer forgot to sign their paycheck. On such a view, there is no sense to the question "how normal is that individual?" since no kind is specified, and individuals are only ever normal members of kinds. I think this view is correct, and so I don't think there's any sense to the notion of normality *tout court*.

One reason for thinking that the normality of individuals is always being a normal member of this or that kind comes from the exploration of what happens when two individuals have abnormalities as members of different kinds. There are two river otters, let's say, Audrey and Gregory. Audrey is an albino otter and so is abnormal in many respects as an otter. Gregory has a few abnormalities as an otter—he's missing a forelimb and an ear from a terrible fight in his

youth—and he lives with a family as their pet and so has a few abnormalities as a pet. Pets don't normally live in dens like a river otter does. Pets don't normally (we might think) eat crayfish. Let us say that when we count up the abnormalities, we get an equal number. Which is less normal? I'm not sure there's an answer to the question until we restrict it to a kind.

Similarly, one might think that normality is relative to a *respect*.[2] Audrey is abnormal with respect to fur color, but is perfectly normal with respect to diet, say. Gregory is abnormal with respect to limb count but is perfectly normal with respect to fur color. Again, it is tempting to think that comparison of normality across respects is as fraught as it seems to be across kinds. If this is true, then the normality of individuals is always normality as a member of a kind in a respect.

1.4 Normativity

At first, it appears that all normality judgments have what might be called a normative inferential profile. On reflection, however, we find that it's at least possible that certain of them may not. First, what is a normative inferential profile, and second, what are some examples of normality judgments that may not have normative inferential profiles?

A normative inferential profile is one that, centrally, includes a licensed inference to a normative claim about individual members of the kind being generalized over. Paradigmatically, this means that a generalization like "dogs have four legs" licenses the inference to a claim like "any given dog *should* have four legs." What exactly the sense of this 'should' is a discussion for a different day, but it does have some normative force no matter what. As a first pass, the idea is that something will have gone wrong for the individual *qua* member of the relevant kind if it fails to have the normal property it should have.

[2] Though the idea doesn't seem to have originated with him, Nickel (2008, 2010a, 2016) is my centrally influence in thinking that normality is always relative to a respect.

Apart from the central 'should' claims, a normative inferential profile includes inferences about defects, privations, abnormalities, interferences, and other normative concepts. A three-legged dog has suffered some *defect*, or been *deprived* of a leg, or alternatively something has *interfered* with a process that was *supposed* to produce or was *aimed at* producing a four-legged creature.

It appears to be far from settled whether or not *all* normality judgments have a normative inferential profile. For instance, it appears to be normal for a hurricane to travel along a certain path given some set of parameters. Has something, though, gone *wrong* for the hurricane if it encounters a perpendicular weather system that blows it off course? I'm not so sure. One can certainly say that it has been *interfered with*.

If *interference*, though, is a normative concept, then the suggestion that all normality judgments have a normative inferential structure is much more plausible. Hurricanes travel along predictable paths unless something interferes, meaning that we can understand the motion of the hurricane with respect to one end state and thereby understand any unforeseen factor as an interference in the normal path.

These inferences to defects, privations, failures, and the like are all implicitly relativized to a respect and a kind. If I say, for instance, "this toaster is defective: it doesn't heat up," what I in fact mean to suggest is that the toaster is defective *qua* toaster with respect to its heating ability. A thing is always abnormal in a respect and relative to a kind if it is abnormal at all.

1.5 Invalidities

As one might expect, given the unique nature of normality judgments, there are countless monotonic and classical validities that do not hold for normality judgments. Modus ponens, for instance doesn't hold since the hallmark of a normality judgment is the possibility of its "antecedent" being true while its "consequent" is false. Contraposition similarly doesn't hold since the "consequent" of a normality judgment is not a necessary condition. If it's true, for instance, that

"When I eat out, I normally eat at Aca Taco", then it does not follow that "If I don't eat at Aca Taco, then normally I don't eat out at all."[3] There are many others, but the idea is sufficiently clear without discussing more of them: many valid inferences in classical logic do not appear to hold for normality judgments.

1.6 Defeasible Modus Ponens

Many researchers in this area hold that the following inference is defeasibly valid (the symbol '⟿' will represent the "defeasible entailment" relation):

(1) Lions are dangerous, Lizzy is a lion ⟿ Lizzy is dangerous.

As with all non-monotonic inference, the addition of new information will defeat the inference. For instance, if we learn that Lizzy is 8 days old, or is catatonic, or has no teeth or claws, then the inference may not go through.

This inference pattern is restricted to general premises expressing *complete* normality judgments. We certainly cannot infer that a cat is calico upon learning that it is a cat. We might, however, accept the following inference and accept it as licensed *by the normality* rather than by some statistical fact or the like:

(2) It's normal for piano players to be able to play Beethoven's Moonlight Sonata,
Chela is a piano player ⟿ Chela can play Beethoven's Moonlight Sonata.

I'll admit to being unsure about what to say about this case. If this supposition is warranted, it might be because it is merely *common* for piano players to learn this song, but it also might be because it is *normal* in Western society to learn this song when learning to play the piano. The strongest version of the normality would obtain if it were somehow a rite of passage to learn this specific song; but this is not the case, so some weaker normality holds between being a piano player and playing this song. It is thus not clear what exactly licenses the inference in (2), if it is licensed at all.

[3] Boutilier p. 92 discusses contraposition. He also discusses modus ponens and other classical inference patterns.

1.7 Irrelevant Information

Asher and Morreau also suppose that one can add irrelevant information into the premises of a defeasible inference involving generic claims without defeating the inference. Their example is as follows:

(3) Lions are dangerous, Lizzy is a lion, Lizzy is five years old ⟿ Lizzy is dangerous.

Lizzy's being five years old counts out one potential defeater: Lizzy is an infant, but it is otherwise irrelevant to the inference in question. After a short preliminary survey of possible counterexamples, this inference pattern seems to work so long as the information added is truly irrelevant. Furthermore, it being a content-based inference pattern, grammar seems to have no effect.

1.8 Defeasible Chaining

The following inference pattern presupposes that the transitive property holds for normality statements:

(4) Birds fly, Sparrows are birds, Tweety is a sparrow ⟿ Tweety flies.[4]

As I'll discuss below, I don't think this particular example works because I don't think that birds "by default" fly and so the defeasible inference isn't warranted. Here's an alternative example that does appear to work by my lights:

(5) Female mammals have mammaries, beavers are mammals, Beata is a female beaver ⟿ Beata has mammaries.

One must be quite careful with the normality statements one assents to when the subject kind is a higher taxon than a species or its non-biological equivalent. If one is suitably careful, though, it appears that transitivity should hold.

1.9 Weakening of the consequent

A basic inference from classical logic is available for certain normality statements. Any inference of the form:

[4] Again from Asher and Morreau (1991)

(6) As are B ⊢ As are B or C.

Should be not only be defeasibly acceptable, but indeed deductively valid. For instance, Asher and Morreau (1991) offer the example of:

(7) Lions have manes ⊢ Lions have manes or wings.[5]

This is all true provided that the premise of the entailment is a *complete* normality judgment. Consider the following claim, which at least invites false implicatures if it isn't itself simply false:

(8) Cats are calico ⊢ Cats are calico or teal.

The implicature of the conclusion seems to be that there are two normal ways of being cats and those ways are being calico or teal.[6] This is clearly false, so the conclusion is false: the entailment does not hold. It follows that weakening of the consequent does not hold for incomplete normality judgments.

There are many deductively valid inferences available through strict entailment to certain classes of normality statements. It won't do to discuss all of them, but it is worth noting that though the logical form of normality judgments is at best complicated, certain inferences appear to be valid that rely on a logical separation between the kind serving as subject and the predicate.

Interestingly, though, weakening of the consequent isn't available for all grammatical forms used to express normality judgments. Consider the following:

(9) (a) (?) It's normal for lions to have manes or wings
(b) It's normal for a lion to have a mane or wings
(c) Lions normally have manes or wings
(d) Normal lions have manes or wings

[5] Asher and Morreau suppose, in calling this 'weakening of the consequent' that generic sentences have a conditional structure. This is a tempting thought, but it's worth noting that the entailment to "Lions have manes or Lions have wings" is available, perhaps even as an interpretation of the above entailment. This would allow neutrality with respect to the logical form of normality statements. A conditional form shifts down to the quantificational, individual level; but we might suppose that normality statements are about kinds, not individuals, and so we would want to resist imposing a conditional form on normality statements.

[6] A deliberate echo of Nickel's (2016) semantics for bare plural generic sentences.

(9)a isn't clearly a felicitous grammatical construction, given the pressure "it's normal for" puts towards a singular subject as in (9)b. Nevermind that, though, since here we're interested in the fact that 3a-d do not appear to have the intended reading. The implication in each seems to be that lions can have manes or wings, no matter which one. It does, however, matter a great deal which one a lion has in deciding whether it is a normal lion!

There is one way of escaping this problem, and that would be to interpret each, respectively, along the following lines:

(10) (a) It's normal for lions to have manes or it's normal for lions to have wings.
(b) It's normal for a lion to have a mane or it's normal for a lion to have wings.
(c) Lions normally have manes or Lions normally have wings
(d) Normal lions have manes or Normal lions have wings

(10)a-d not only seem to be true, but follow simply by propositional logic from the premise in (7). If this is how we're meant to read (9)a-d, or even the conclusion of (7), then the point is somewhat uninteresting: propositions expressing normality judgments obey propositional logic in that they can be combined using standard propositional connectives. If the inference pattern is supposed to be more interesting, then it only seems to apply to the generic form.

1.10 Dudley Doorite

Here is an exemplar of a similar inference pattern, again discussed by Asher and Morreau:

(11) Quakers are politically motivated, Republicans are politically motivated,
⊢ Quakers and Republicans are politically motivated.

We can conjoin subjects, it seems, and so the subjects must play an independent role in the logical form of normality statements.[7] In this case, we're not faced with the dilemma from the previous inference, since it seems to work for normality statements of each grammatical form. Conjoining subjects of judgments with identical predicates seems to be valid on all of the following:

[7] Again, we're faced with the problem of identifying whether the conclusion should be "Quakers and Republicans are politically motivated" or "Quakers are politically motivated, and Republicans are politically motivated." It is in principle possible that the conclusion of this inference appears to follow because it is shorthand for the longer conjunction.

(12) (a) It's normal for Quakers and Republicans to be politically motivated
(b) Quakers and Republicans are normally politically motivated
(c) Normal Quakers and Republicans are politically motivated

It is somewhat unclear, though, what to make of:

(13) It's normal for a Quaker and a Republican to be politically motivated.

The sense of such a statement might be that there are two specific people, or one specific person who is "a Quaker and a Republican." Either way, the case is less clear than those of 6a-c.

1.11 Competing Normalities

A classic example from the nonmonotonic logic literature—in particular those discussing default logics—deals with competing normality judgments.

(14) Republicans are nonpacifists, Dick is a republican, Quakers are pacifists, Dick is a Quaker,
➥ Dick is a nonpacifist,
➥ Dick is a pacifist

Since the two judgments are complete normality judgments (there is no other normal way to be a Republican with respect to pacifism vs. militarism; similarly for Quakers), the inferences seem to work. The logic of normality judgments, it seems, won't decide such a case for us, so we must withhold assent from the incompatible conclusions until the facts of the world tell us which is true.

There's a connection between these sorts of inference patterns and normativity in that insofar as Dick is a Republican, Dick *should* be a hawk; but insofar as Dick is a Quaker, Dick *should* be a pacifist. Something has interfered with either normality if the normal outcome doesn't obtain.

2. Perhaps controversial

Having settled on the most basic and uncontroversial aspects of the inferential profile of normality judgments, the next task is to explore some inference patterns that are still quite basic to the kind of judgment but may not be quite so *prima facie* plausible as those we've just discussed. In fact, some may not hold for normality judgments at all.

2.1 Exceptions

What sorts of inferences regarding apparent exceptions are available to normality judgments? On the way to answering such a question, we'll do well to understand what sorts of exceptions are problematic for normality judgments and what sorts are consistent therewith.

As a rule, apparent exceptions to normality judgments do not refute or falsify them. There are, however, a few sorts of exception that do appear to interfere with the normality judgments to which they are exceptions. For one, as Leslie (2008, 2012a, pp. 26–27) points out, the difference between positive counterinstances and instances of mere lack (negative counterinstances) appears to be relevant in deciding whether an exception is merely apparent or is a true exception. Consider:

(15) Lions have manes.

This is a true normality judgment apparently because female lions simply lack manes. They don't have mohawks or something similarly positive. Consider, alternatively:

(16) Dogs weigh over 40 lbs.

This is a false normality judgment perhaps because weighing under 40 lbs. is not only common among dogs but is itself a positive way to be: it's not a mere lack.

One might worry, however, that it's the *naturalness* of being a small dog that refutes (16) rather than the positivity of weighing under 40 lbs. After all, it's a perfectly normal way to be a dog to be a Jack Russell Terrier. This, however, won't explain the phenomenon illustrated by the truth of (15): it's a perfectly normal or natural way to be a lion to be a female lion who lacks a mane, but that does not refute (15). It seems, therefore, that one must posit an alternative distinction between (15) and (16) that explains why the former is true and the latter false. The distinction on offer from Leslie is that one has a positive counterinstance while the other has counterinstances that are mere lacks. This is the best account available in the literature, even if one would want to know whether there are cases where a refuting counterinstance is a mere lack and whether there are cases where a non-refuting counterinstance is a positive case. More work must be done here.

Other sorts of apparent counterinstances, because the conditions present are not normal, ideal, or specifically correct, fail to refute the relevant generics. Consider, for instance:

(17) Water is liquid

This claim appears to be true on its face, but it relies on "normal" conditions where room temperature and pressure or some "normal" range of thermodynamic conditions is assumed to hold. Ice is no counterexample because ice only occurs outside of the temperature and pressure conditions assumed by (17). Similarly, the famous

(18) $f = ma$

Or "force equals mass times acceleration" ends up accurately describing virtually no physical systems because, as Smith points out (2002), it does not function as an empirical generalization in the practice of physics. Only in ideal conditions does it hold as a description and therefore, though it is true, it appears to only hold given ideal conditions. Of course, (18) may or may not be a normality (as opposed to nomic) judgment, but regardless it serves to illustrate the idea that certain normality judgments would only be refuted by counterinstances that held in suitably similar and therefore suitably ideal conditions.

There is, no doubt, more to be said about varieties of and conditions on counterinstances to normality judgments,[8] but this will suffice for this brief survey.

2.2 Statistical Consequences

Some[9] have defended what Schurz calls the *Statistical Consequence Thesis*: the claim that complete normality judgments have reliable statistical implications. This is supposed to be important because it gives complete normality judgments empirical content and therefore makes them answerable to empirical science. Pearl (1988, pp. 477–80) also argues that having statistical consequences maintains

[8] Cf. Greenberg, Y. (2007), and Strahovnik (2016) for more discussion.

[9] Rescher (1994, p. 10ff) and Schurz (2009b).

the usefulness and reliability of complete normality judgments.[10] It is important to note, however, that the implication is asymmetric: normality judgments imply statistical claims but not vice versa.[11] One consequence of this is that a transposed form of the statistical consequence should be available. "Dogs have four legs" implies something like "most dogs have four legs"; therefore "a minority of dogs have four legs" implies that it is false that "dogs have four legs" (supposing we understand the latter as a complete normality judgment).

It turns out to be a bit difficult to specify a precise statistical value. One might tend toward a simple majority, but counterexamples abound. Schurz presents some preliminaries (2009b) the implication of which is that what he calls 'normic laws' have "high conditional statistical probabilities, but their numerical values are usually unknown." He also concludes that the minimum acceptable statistical probabilities vary from domain to domain.

As we discussed in the previous chapter, careful study reveals surprising links between normality judgments and statistical judgments about what is actually the case. Nevertheless, the thesis that complete normality judgments *entail* particular statistical judgments is too strong and raises more problems than it solves. First, on the explanatory link account normality judgments aren't primarily useful because they make empirical predictions. Instead, they are useful because of their role in explanatory practice and because of the way they track explanatory links between kinds and outcomes. So the fact that they do not entail even vague statistical judgments doesn't count against their usefulness.[12] Second, statistical judgments are by their nature tied to the way the world *actually is or has been*. The problem arises again and again, though, that there's no guarantee that the actual

[10] Cf. also Schurz (1997) for similar arguments and Schurz (2009b) for an updated discussion.

[11] At least not synchronically or directly. Schurz (2009b) has a discussion of the possibility that diachronic changes in the statistical facts about a population may serve as evidence for or even entail changes in the relevant normic laws.

[12] There's an ambiguity between a normative and an epistemic use of "expect" and "expectation." Even if one overcomes this ambiguity, however, there are at least three grounds for epistemic expectations: regularities, necessities, and normalities. Expectations are grounded in more than just statistical probabilistic judgments.

should be representative of the normal or "in principle."[13] Hence, there may very well be an intractable problem: perhaps one can never offer a generally acceptable statistical entailment from normality judgments.

Nevertheless, the thought that over time even abnormal conditions may come to be seen as normal conditions and so the individuals that are adapted to those new conditions may become the normal ones while those that are only adapted to the old conditions may become abnormal is a quite compelling thought. Furthermore, since our understanding of kinds intrinsically involves the conditions in which the individuals of those kinds live, it is quite plausible to suppose that as a rule statistical claims can be derived from normality judgments. Perhaps this entailment isn't terribly modally robust, and so perhaps it's no strict entailment at all.[14] Nevertheless, at least as a matter of practical expediency and at best as a matter of good natural philosophy, the normal will by default be quite common.

2.3 Kind Percolation

Bernard Nickel has defended the claim that there is a logic to bare plural generic sentences such as (15), (16), and (17) above. Nickel claims (2014, 2016) that an inference holds for generics such that, loosely, the attribution of a normal property to a kind "percolates" up to super kinds like genera and other higher taxa. So, from:

(19) Penguins swim.

The inference to (20) is warranted.

(20) Birds swim.

---

[13] McCarthy (1986, p. 91); Reiter (1988, pp. 149f, 180f); and Millikan (1984) reject the statistical consequence thesis on more or less exactly these grounds.

[14] Veltman (1996) and Boutilier (1994) take it that the world is normal—that things are normally normal or that the actual world is normal most of the time. This is a tempting thought, but the world is also replete with abnormalities tragic and trivial. One begins to wonder whether the world is in fact normal. A deeper concern is rooted in the idea that what is normal is always relative to a kind specified in a particular context and so the idea of an "objectively normal" world may not be sensical at all. The rabbit must die "before its time" if the fox is to feed its pups.

This should work, he claims, for all *kinds*, but not for any set at all. "It is widely and correctly observed," Nickel writes, "that generics do not in general allow an inference from subset to superset" (2014). When dealing with proper kinds such as natural kinds, however, the inference is warranted. So while one cannot infer from "albino ravens are white" to "ravens are white," (albino ravens do not constitute a kind, but instead merely a derivative subclass), one should be able to infer from "dogs have four legs" to "mammals have four legs" (Nickel 2014).

How does Nickel dispatch with apparent counterexamples? After all, they are legion—even given the restriction to kinds rather than classes or sets more broadly. One seemingly cannot make the following inferences:

(21) Platypuses have venomous barbs ⊬? Mammals have venomous barbs.

(22) Portabella mushrooms are edible ⊬? Mushrooms are edible.

(23) Emus are flightless ⊬? Birds are flightless.

Nickel responds to these apparent counterexamples by positing an illicit implicature present in the conclusions of these inferences. The implicature is supposed to be something to the effect of "and that's the whole truth" (Nickel 2014, 2016). This is a problematic implicature indeed, since not all species of mammal have venomous barbs, not all species of mushroom are edible, and not all species of bird are flightless. In short, these are not the only ways of being a normal mammal, mushroom, or bird. Nickel's claim that this is an illicit implicature, though, avoids the implication for his account of the logic of generic. For instance:

(24) Mushrooms are edible ⊬ non-edible mushrooms are abnormal.

This inference is not warranted, though it is a natural *implicature* from the premise. Any of what I have been calling the *incomplete* normality judgments can in certain grammatical forms and in certain contexts appear to implicate that the whole story has been told. It sounds odd to say, out of context:

(25) Cats are calico.

It does not sound so odd to utter (25) as a part of a longer claim:

(26) Cats are black, and cats are calico, and cats are white and…

Now it is always possible that this context has shifted our natural interpretation of these disjuncts to be claims involving the quantifier "some". If that is the case, then the example does nothing to show that bare plural generic forms of incomplete normality judgments can stand true on their own. Perhaps the following context is different:

(27) Birds get around in many ways. Birds fly. Birds swim. Birds run. Birds waddle.

Now it seems perhaps a bit clearer that each generic sentence stands on its own as a true statement. No sentence, I suppose, is true apart from some linguistic context or another; so it shouldn't be a bother if these sentences appear to be true only given certain linguistic contexts.

That's the case Nickel offers for his claim that there is a perfectly general inference pattern holding for all bare plural generics about kinds. Let us explore a few objections before making a final assessment of the relevance of such an inference pattern for normality judgments writ large.

David Liebesman's (2017) review of Nickel's (2016) monograph points out what appears to be an awkward result of Nickel's analysis. If this illicit implicature is present for cases of uncommon but nevertheless normal properties, it should hold just as well for cases of common and paradigmatic normal properties. So, on Liebesman's analysis,

(28) Birds fly.

should ring just as odd to our ears as:

(29) Birds run.

but it does not sound just as odd. What's more, the implicature should occur in both cases and so (28) should seem false; knowing as we all do that birds also swim and run and waddle. (28), though, is a fairly uncontroversially true generic.

To rehearse, the claim is that there should be the same implicature for each case. (28) and (29), therefore, should implicate, respectively:

(30) Flying is the only means of locomotion for birds.

(31) Waddling the only means of locomotion for birds.

Both of these implicated claims are false, but the original generics seem both to be true. This is what Nickel's account predicts. (29) is supposed to sound odd. If it does, that is supposed to be explained by appeal to the false quantity implicature captured by (31). The same goes *mutatis mutandis* for (28), though, so Liebesman argues that the account Nickel has given predicts that "birds fly" should sound odd to our ears, but it does not.

I'm not entirely convinced that Liebesman has identified a real problem for Nickel's account here. (28) doesn't sound odd to us for psychological as well as for semantic reasons. Our minds appear to be built to process paradigmatic properties more quickly and easily,[15] but this shouldn't clearly have implications for the semantics or logic of generic sentences. Moreover, and perhaps because my pretheoretical intuitions have vanished after years steeped in various theoretical approaches to generic sentences, but "birds waddle" doesn't sound at all odd to my ears any more than "birds fly" does. Perhaps the example itself is unclear.

Michael Strevens (Strevens 2014) discusses an alternative example: the leucistic raven. A leucistic raven is neither black nor albino. It is white, but not due to a lack of melanin throughout its body including its internal organs. Instead, a leucistic raven has muted pigmentation in its plumage alone and so suffers few of the ill effects that accompany albinism in ravens. It is, therefore, not clearly abnormal for a raven to be leucistic. Let us assume, for the sake of the study, that this is merely a normal variation in ravens—albeit a statistically rare variation. If this is allowed, then (32) might sound odd to our ears, but should be counted true on Nickel's analysis.

[15] Though there is ample evidence for this conclusion, the *locus classicus* for the claim is Rosch (1978).

(32) Ravens are white

There is a normal way of being a raven with respect to coloration and that way is being leucistic or white, and so even though it initially may appear to be false, it nevertheless must be counted among the true generic sentences. Nickel offers an explanation for why it appears to be false: it implicates illicitly the following proposition:

(33) White is the only normal coloration for a raven.

Liebesman enters at this point in the analysis and points out that:

(34) Ravens are black.

Should also implicate something analogous. Perhaps:

(35) Black is the only normal coloration for a raven.

And so, on final analysis, (34) should sound false even though it is among the most widely accepted generics in the literature and in empirical studies. Again, though, Liebesman's argument does not appear to hold water, in that the psychology behind why we might initially reject (32) while initially accepting (34) has more to do with general ignorance of leucistic ravens (or the non-salience of swimming birds, spotted zebras, and other statistically rare but still normal specimens) than it has to do with deep facts about the truth of generic sentences. With that in mind, it appears that Nickel's implicature-based explanation is at least a provisionally acceptable explanation of why "incomplete generics" appear to be false while being in fact true.

One more case of apparently illicit kind percolation is worth investigating. Krifka, et al. (1995) observe that the following inference is not monotonically valid:

(36) Berber lions are well-adapted to cold weather ⊬ lions are well-adapted to cold weather.

On Nickel's account, though, this should be a case where kind percolation holds. Lion is a species, Berber or Barbary lion is a subspecies. They both appear to be proper kinds—the sorts of things

that Nickel would allow as kinds at any rate. Yet the move from the conclusion of (36) to (37) or something to the same effect seems in this case to be more than a mere implicature.

(37) Lions not well-adapted to cold weather are abnormal

Whether it is the evaluative nature of the antecedent in (36), or the fact that being "well-adapted" is wrapped up with being a kind of thing more intimately than is "being white" or "swimming," it seems that (37) is part of the content of the claim. If this is true, then the consequent of (36) is false (for (37) is surely false) and the inference in (36) is indeed a faulty one. Nickel's kind percolation, then, wouldn't hold for all bare plural generics—nor for all normality judgments.

Nickel could offer one of three responses: a) deny that subspecies are proper kinds in his sense, b) argue that (37) is not an implication of the consequent of (36), or c) simply accept (37). No one will go for the third option, since it is patently false that lions adapted for the Tanzanian savannah are abnormal lions. Nickel would certainly go in for the second option (though there is also reason to think that he might go in for the first option as well).[16] According to his semantics for generic sentences, there is no entailment from "lions are well-adapted to the cold" to "lions that aren't well-adapted to the cold are abnormal." If this implication held for all bare plural generic sentences, then they could never be used to express incomplete normality judgments and it seems that they can.

At long last: what implications does this analysis of the logic of *bare plural generics* have for the logic of *normality judgments* as a whole? The first thing to note is that there is one grammatical construction for a sentence expressing a normality judgment that appears not to have the illicit implicature identified above. Namely, that found in an antecedent of (2), which is the form:

(38) It is normal for k's to φ

This form tends (or so it seems after some informal study) not to activate the intuition that:

(39) *φ*ing is the only normal way to be a *k*.

[16] See the discussion of albino ravens in Nickel (2016, pp. 202–208).

Other grammatical forms that express normality judgments, however, are more prone to implicate something of the form of (39), and so are likely to give rise to corrupted linguistic intuitions.

There is a sense in which kind percolation—according to the explanatory link account—simply *cannot* hold for normality judgments since the specification of the individuals involved is supposed to explain the individuals' possession of the normal feature in question. That a bird is a penguin explains its swimming, but its being a bird, it seems, cannot explain its swimming since the class *aves* has no direct connection to the *niche* occupied by penguins. The class simply doesn't offer the explanatory resources necessary to explain the means of locomotion. There's a deep sense in which the higher the taxon one uses to specify an individual, the fewer explanatory resources one has in order to explain particular features, activities, and conditions. Many features, it seems, can *only* be explained by appeal to the species (in some sense of species) and not by appeal to any higher taxa. If this is all correct, then kind percolation does not hold for normality judgments for semantic reasons.

If this isn't correct, or if there are true normality judgments that attach to higher taxa *as well as* lower taxa, then the explanatory link account is faced with a serious challenge. Let us explore whether there are cases of true normality judgments at both higher and lower taxa and how the explanatory link theorist might deal with this problem.

(28) Birds fly.

(40) Sparrows fly.

(41) Female mammals have mammary glands.

(42) Female beavers have mammary glands.

(43) Bacteria are smaller than animal cells.

(44) Staphylococci are smaller than animal cells.

It seems, in fact, that there may be *countless* true normality judgments at higher taxa that are also true of lower taxa. This is a problem because the explanatory link account holds that there must be an explanatory link between the kind and the feature for the feature to be normal for the kind.

In Chapter 4, we discussed a few possible strategies for dealing with this problem. In the end, room has to be made for an operation whereby higher taxa or genera can borrow normalities in a derivative way from their member species. So it is true that mammals lay eggs just as it is true that mammals give birth to live young: both are derived from normalities attaching to the species that make up the genus of mammals.

Nickel thinks kind percolation holds for any bare-plural generic sentence so long as one restricts oneself to dealing with kinds rather than mere sets. The strategy just discussed—allowing for derivative normality judgments attaching to higher taxa—allows for Nickel's kind percolation to hold for normality judgments, so long as one recognizes that there are different statuses of normality judgment. Some are complete and some are incomplete: some ascribe the *only* normal outcome in a respect to the kind whereas others ascribe merely one of multiple contrary normal outcomes. Some are basic and some are derived: the truth of some normality judgments is grounded in an explanatory link between *that* kind and the outcome ascribed whereas the truth of others are grounded in an explanatory link between their member sub-kinds and the outcome ascribed. Some genus-level normality judgments are basic, like mammals have mammary glands, whereas others are derivative, like mammals lay eggs. Both, though, appear to be true.

3. Mistakes

Having introduced what I take to be the essential and characteristic features of the inferential profile of normality judgments, it will be fruitful to spend some time exploring what I take to be mistakes in exploring their inferential profile further.

3.1 Background Knowledge

The first mistake in the field is a tendency to allow our own background knowledge of a particular case infect our intuitions about the sort of judgment in play. For instance, Nickel claims that the generic "ravens are white" is false and furthermore that the generic "ravens are black" seems to mean that *almost all* ravens are black.[17] I'm not so sure. First, leucistic ravens are not clearly abnormal and so it simply may not be true that "ravens are white" is false. Second, if it isn't true *in general* that a generic of the same form as "ravens are black" requires that there be near universality, then we'd need a much stronger reason than linguistic intuition to think that this generic is markedly different in its content. Linguistic intuitions are sensitive to background knowledge and thus the fact that nearly all ravens are indeed black seems to push our intuitions towards thinking that the statement itself *means* that almost all ravens are black. We shouldn't be too quick to trust our linguistic intuitions about seemingly outlying cases. The statement could just as easily *mean* something much weaker even if our acceptance of related claims—like, for instance, ravens are white—is influenced by apparent knowledge we have of ravens.

3.2 Content Sensitivity

On a related note, certain inferences involving normality judgments appear to depend on the content of the statements involved and so seem not to be good formal inferences at all. There may be principles governing which are good material inferences and which are not, but the inferences do not hold as a matter of form.

[17] "Among the exceptions to a generic, there are exceptions that, were they actual, would serve to falsify the generic. Call these genuine exceptions. For example, the generic ravens are white is false because there are genuine exceptions to the generic. In the framework of normality that I am proposing, that generic is false in virtue of the fact that there are normal ravens that aren't white. [n. 4 more precisely: because there are no ravens that are normal in any of the ways a raven can be normally colored and that are white]. It is the fact that the non-white ravens are normal that makes them genuine exceptions. Then there are exceptions that, even if they are actual, leave the truth of the generic untouched. Call these merely apparent exceptions. Thus, ravens are black has exceptions, and given that the generic is true, we can tell that they are all merely apparent. In the framework of normality, that is to say that none of the non-black ravens are normal" Nickel (2016, p. 54).

For instance, we can strengthen the "antecedent"[18] of a normality judgment in some ways, but not in others. To strengthen the antecedent of a proposition is to add a qualifier to the subject term of that proposition such that the domain becomes smaller. Here are some examples of apparently good and bad antecedent strengthening:

(45) Birds fly.
⊢ white birds fly.

(46) Birds fly.
⊬ dead birds fly.

The content of the strengthened antecedent in (46) seems to block the acceptability of the inference, even if the inference pattern works in (45) and many other cases. We might think this is a sign that this isn't so much a matter of form as a matter of the material content of the propositions involved.

This example, though is problematic in that dead birds are arguably not birds at all. Some adjectives change the very nature of their subjects, while others don't. It is quite plausible to presume that some inference patterns will distinguish between these two types of adjectives. This distinction might save antecedent strengthening as a valid inference pattern for normality judgments.

Let's look at two more:

(47) Platypuses have venomous barbs.
⊢ male platypuses have venomous barbs.
⊬ female platypuses have venomous barbs.

This inference might not work at all since the alternative conclusion about female platypuses isn't true. The fact that the analogous conclusion isn't true suggests that the strengthening in (45) and (47) only *appears* to work because the conclusion is true and not because the inference is generally

[18] Again discussed in Asher and Morreau (1991). (45) and (46) are their examples.

acceptable.[19] We can therefore conclude that antecedent strengthening does not appear to work for normality judgements—whether complete or not.

3.4 Ignoring incomplete normalities

Often, in the literature, the idea that normality judgments often express *incomplete* normality judgments is overlooked, circumscribed, or explicitly disallowed. I regard this as a mistake. Consider the following proposition:

(48) Cats are black

The claim cats are black, seems to be false taken on its own. This may be, Nickel (2016) argues, due to an implicature to the effect that the speaker has stated all there is to know about the coloring of cats. When we conjoin the statement with others like it, though, it appears that each might be true on its own:

(49) Cats are black, and cats are brown, and cats are orange tabby, and cats are calico, and cats are gray, and cats are white, and cats are…

This seems fairly convincing. If it is acceptable with a shift in context, then it seems that the *content* of (48) doesn't include any sort of exclusivity claim to the effect that cats aren't also brown or white or grey. (49), that is, appears to have essentially the same content as:

(50) Cats are black, or white, or brown, or orange tabby, or calico, or gray, or white, or…

We could understand this phenomenon according to the idea that grammar is often deceptive. If it looks generic, in other words, it might still be statistical in semantic content. Instead of conjoining a series of propositions with generic content, we might, within the bounds of credulity, be changing the content of each to something closer to:

(51) Some cats are black

[19] Boutilier offers the following counterexample to strengthening: If a match were struck, it would light, so if a match were struck and it was wet, it would light. Since this example involves abnormal conditions, though, it doesn't have as much strength against the defeasible form of the inference pattern.

This proposition is obviously true, and so the shift in initial plausibility from (48) to (49) may reflect a shift in content rather than a mere shift in context.

If this is what happens, then the inferential structure of (51) should be present for each individual conjunct of (49). It does not clearly appear to be. Should once in the history of the species a teal cat be born, never again to happen, it is thereby true that "some cats are teal," but one might think it quite strange to add "cats are teal" to the list in (49).

3.5 Normality and Defaults

Another mistaken tendency in the literature is to assume that any normality claim is automatically a default rule and vice versa. A default rule is a rule to the effect that if the antecedent condition—a kind or property—is satisfied, then the consequent of the rule is reasonably inferable, or can be concluded absent defeating information. Not all normality judgments correspond to default rules, nor is the reverse true. For instance, it's true that birds fly and it's normal for birds to fly, but it can't be inferred from the fact that Tweety is a bird that Tweety flies.[20] This is a controversial claim, since virtually every artificial intelligence and logic researcher for the past 30 years has assumed that one can.

Recall from earlier the distinction between a complete and incomplete normality. A complete normality is one that exhausts the normalities for the kind in the relevant respect. For instance, there is only one normal limb count for canines: four. To say that it is normal for dogs to have four legs is therefore to exhaust the field: there are no other rival normalities. To say that cats are black, though, far from tells us everything we need to know about the fur color of cats. Cats are black, brown, calico, orange tabby, white, and so on. The complete normality that corresponds to the incomplete "cats are black" is a quite long disjunctive list since cats have many normal fur colors.

[20] This example shows up repeatedly. Cf. for some instances Horty (2012), Strößner (2015), (Schurz 2009b), and Boutilier (1994).

This distinction is relevant to the current discussion in that only complete normalities can be reasonably inferred *on the basis of a normality*. That is, there are many grounds of reasonable inference, only some of which are normalities and normality judgments. We can infer that a human being lacks blonde hair simply because the numbers and therefore the statistical likelihood is overwhelmingly in favor of non-blonde hair. This is, it seems, a reasonable inference or at the very least a good bet. But it is not exclusively normal to have black or brown hair, even though the percentages are overwhelming, and so any inference from human to black or brown hair is not an inference on the basis of a normality—however reasonable that inference may be. We can therefore conclude, if the foregoing is correct, that default rules encompass more than normality judgments and not all normality judgments license default rules. Perhaps all complete normality judgments license default rules, but not *only* normality judgments ground the reasonableness thereof.

So the question at hand is whether or not "birds fly" expresses a complete normality judgment. Clearly, the answer is "no": penguins, emus, and dodos are all birds and perfectly normal birds at that. I'm at a loss for what it could mean—in the realm of non-statistical normality—for a penguin to be an *abnormal* bird. We should say, instead, that there are multiple normal mode of locomotion for birds, and they include flying, swimming, waddling, and diving.

The mistake in the literature is, therefore, to mix up normality judgments and default rules and therefore to treat them as extensionally equivalent—they are not. If "birds fly" expresses something to the effect of "it's reasonable to infer that a thing flies upon learning that it is a bird insofar as one lacks further information about its species," then "birds fly" is not a normality judgment but is instead the expression of a sort of statistical rule of inference. I don't think it expresses this judgment. Instead, I think it expresses something more akin to "there is at least one normal mode of locomotion for birds and that mode is flying."[21]

[21] Note that this is modeled loosely on the semantics that Nickel develops in (2016).

3.6 Overlooking normativity too quickly

One controversial aspect of the inferential structure of normality judgments is their *normative* inferential profile. Many complete normality judgments appear to license the claim that if the complete normal outcome doesn't obtain, then something has *gone wrong* or the subject of the normality *should* have the feature or one of the features involved in the normality.

Nickel thinks this isn't a general feature of normality judgments for the following reasons:

> "But there are many generics for which this inference intuitively fails. It is true that the lion has a mane, and yet it's false of any given female lion that she ought to have a mane. It is true that bees produce honey, and yet it's not true of any particular queen that she ought to produce honey. It's true that the human gives birth to live young, yet it's false of any given man that he ought to give birth to live young. It's true that the human is social, yet it's false of a particular hermit that she ought to be social: she is not to be criticized simply because she has chosen a hermit's existence."

Unless these inferences turn out to be licensed, the normativity of normality judgments in general is something we should question.

I am not, however convinced by Nickel's cases. It seems plausible that the subject of generics is not what it seems in certain cases, and that in other cases the inference to a normative judgment is relative to respects and is therefore content-sensitive. In the generic "bees produce honey," it seems like the claim is something like "bee colonies produce honey" or perhaps "bees work with those in their colony to produce honey." The normative judgment seems to attach to different subjects and so it's natural to think either that the subject of the normality judgment itself is different than 'bees' or that the inference to the normative judgment changes the subject.

Alternatively, when inferring from 'humans give birth to live young' to 'x human ought to give birth to live young,' we might wonder whether the respect itself demands that the value range be restricted to judgments about biologically female humans. We should furthermore take great care in our understanding of the claim involved. Is the claim that any woman who does not give birth is subject to criticism or is the claim instead that when human beings give birth, they either give birth to live young or something has gone deeply wrong? The latter seems like a quite plausible

interpretation. Humans laying eggs would be abnormal indeed. The same point goes for our hermit: the 'ought' simply can't be an 'ought' of obligation—after all, no one would claim that a stag is *obligated* to have horns. Instead, the claim seems to be that there will be some independent explanation for why the hermit is not social and that this explanation will cite factors that prevent her from taking part in an important aspect of human life. No criticism needed, but we *do* need some explanation of why the hermit doesn't have a social life and that explanation will almost certainly cite abnormal factors causing such an outcome.

What's the moral of this story? It is all too often that we move past normative claims, thinking that there's basically one sort of normativity—the kind that attaches to and governs rational agents. This is an overly narrow and problematic notion of normativity, and it is not the sort of normativity attaching to the vast majority of normality judgments.

3.7 Specificity Thesis

Finally, a few (Asher and Morreau 1991; Schurz 2009b, p. 190) have suggested (though not in discussions of normality judgments *per se*) that we may—at least as a generic rule if not as a universal principle—take the more specific normality judgment as true when it conflicts with a more general judgment. The commonly cited example is one I discussed above: "birds fly". These researchers have supposed that the fact that emus run conflicts with the fact that birds fly. The natural thought, then, is that the more specific judgment about emus takes precedence in deciding what one should conclude about the case. This 2nd-order judgment is enshrined in the general principle that the more specific of two conflicting normality judgments is the one that is true of the case.

At first, this seems to be a plausible solution to an apparent problem: two different judgments conflict and it appears that we systematically get the right result when we privilege one over the other. Upon reflection, though, it appears to be quite strange to think that, as the logic requires, emus are abnormal birds. Treating a species as an abnormal member of a genus seems to be clearly

problematic given the relation between a genus and its component species—a relation that is quite unlike the relation between individuals and their species. It is just as true, it seems, that birds run, since emus normally run. It would therefore be wrong to think that something flies simply because it is a bird; it instead either flies, or runs, or swims.

The upshot of this discussion is that there shouldn't be cases of more and less specific normality claims that nevertheless translate into default rules. Birds do not by default fly just as apes do not by default have prehensile tails. They do, however, by default present *either* feature. That is, birds by default fly, or swim, or run, or waddle. Apes by default either have prehensile tails or non-prehensile tails.

4. Final Thoughts on Inferential Structures

There are two central errors in dealing with normality judgments. First, many have assumed that normality judgments about higher taxa work essentially the same way as normality judgments about species or basic kinds. This is not correct, as I have argued. Second, many have rejected the possibility of incomplete normality judgments or simply ignored them in their discussions of the inferential structure of normality judgments. As a result, they've been mixed in with statistically grounded default judgments like "if it's a prime number, then it's odd" and nomically-grounded default judgments like "if it's a dog, then it's a mammal." It is, to the contrary, of utmost importance to separate these sorts of judgments.

Final Conclusion

Though this is a mere opening move in a long conversation to be had about normality judgments, the groundwork of a picture of what a normality judgment is in terms of content and structure has been laid. The notion of a normal feature or outcome can be explicated in terms of the explanatory link some kinds have with at least many of the features or outcomes that characterize them. A normality judgment is a generalization—an attribution of features to a set or kind of thing—that ascribes *normal* features to a kind. Since the normal features are the features that explicate the concept of the kind itself, a normality judgment is centrally a judgment that explicates the nature of a kind of thing.

A necessary precondition for such a view to be tenable is that a particular form of explanation—in which the specification of an individual of a kind explains that individual's possessing a particular feature—must be a sufficient form of explanation. If the account is to be that a normality judgment ascribes an explanatory link between a kind and a property or feature, then it must be sensical to suppose that there is such an explanatory link. The work of the second chapter was to defend such a view. An individual's possession of a property is explained by that individual's being the kind of thing it is. Flipper the dolphin has a blowhole because he's a dolphin. This only *appears* to be an insufficient or placeholder explanation because we are naturally interested in a further question: why do dolphins have blowholes? This is an importantly different question, and the original question has already been answered.

The explanatory link account itself finds inspiration in the Aristotelian tradition, according to which the fact of being a member of a kind serves as an explanation for a variety of facts about

oneself. The neo-Aristotelian tradition[1] has capitalized on this relatively simple notion found in Aristotle's texts to argue that a sophisticated and quite detailed picture of the good or moral life can be derived from a fuller understanding of the kind of thing the moral agent is. In our case, the kind humans has a whole multi-dimensional evaluative framework built into it. The idea that concepts have a normative or evaluation structure is central to a specified form of the explanatory link account of normality, according to which the explanatory link between a kind and its normal features is grounded in the fact that the kind serves as a standard of evaluation for success or perfection as a member of that kind. There's no further explanation for why a runner should be found at the track practicing sprints. They're a runner, practicing is a necessary component of the process of becoming an excellent runner. The evaluative nature of the kind concept gives the concept explanatory force.

In committing to a project akin to conceptual analysis rather than of commentary on the ways the concept of normal have been used or on the mistakes people make in asserting what is normal and what is abnormal, I have avoided most questions of deep social and ethical importance. Whether it be a strength or a weakness of the project, the central claims of this dissertation can be true almost regardless of the state of the world. There may very well be very few or no normalities in the world: the world might be either essentially random chaos governed not by principles but by individuals and their multifarious doings about which one cannot say much of scientific import, or the world may be governed by necessary and universal laws, or perhaps some mixture of the two. These are, by my lights, clearly bad descriptions of the world, but no argument to that effect is to be found in these pages. The project is, in this sense, wholly "second order" in that there aren't any first order claims about the state of the world that bear on the truth or falsity of the central account of normality judgments on offer.

[1] Cf. in particular Geach (1956), Anscombe (1956, 1957, 1958a, 1958b, 1979) Lawrence (2006).

In lieu of a more comprehensive summary of the results of this study, allow me to highlight some of what I take to be the core insights gained.

First, normality judgments are distinct from nomic and statistical judgments in the sense that these categories are mutually irreducible. One cannot infer from a judgment of the form "trees have leaves" a judgment of the form "all trees have leaves" or any judgment similar to "most trees have green leaves." Normality judgments are relatively insensitive to changes in the prevalence of a feature in the population of members of the kind in question, while statistical judgments are quite sensitive to such changes. Similarly, normality judgments are tolerant of apparent counterexamples, while nomic or universal judgments are constitutively intolerant of counterexamples.

Second, normality judgments sort into complete and incomplete normality judgments and this is often overlooked. Incomplete normality judgments attach to a sense of normality according to which some features or conditions are normal *but not every contrary feature or condition is abnormal*. It's perfectly normal for cats to be black, but it's also perfectly normal for them to be gray. Trees have leaves, but they also have needles. The sense of normality I call 'complete' is the sense according to which there is only one way to be normal in a given respect. It's normal for dogs to have four legs, *and no other limb count is normal*. This is a marked departure from many in the literature, who either focus exclusively on complete normality or don't acknowledge that there is such a phenomenon as incomplete normality.

Third, there is a form of explanation—one that turns out to play an important role in understanding what it is for a feature to be normal—in which one simply specifies an individual so as to *ipso facto* explain that individual's possession of a particular feature. Kind explanation likely works via a number of explanatory and pragmatic mechanisms, but one of the central pragmatic roles it plays is its shifting of the level of discourse (or attempting to, at any rate) "up" to the level of the kind. It's a sort of mistake in normal circumstances to ask why an individual raven is black, so in

specifying that individual as a raven, one is shifting the focus of the conversation from the individual raven to the kind of which it is a member—the kind which explains that individual's being black.

Fourth, according to the account I have developed and defended herein, what it means to be a normal feature is to be a feature the possession of which would be explained by the individual possessor's being a member of the kind for which one is normal. In simpler terms, a feature is normal for a kind if and only if there is an explanatory link between that kind and that feature.

Fifth, because of the conceptual link between normality and explanation, normality ends up playing an important role in the practices surrounding explanation. Normal features and particularly conditions are assumed to hold in the crafting of an explanatory narrative, and normal features serve as the default contrasts in calls for explanation. The normal, in short, populates the common ground and so serves as the default which must be explicitly overturned through updating the common ground.

Sixth, the extensional link between the normal and the frequent or regular is multifarious and complex. There is no simple principle like "the normal is always regular and the regular is always normal" governing this relationship and so, as with most things, one must take each case on its own. Nevertheless, as a general rule, normalities tend to be quite regular and regularities tend to be normal. If we find that something is no longer regular, we'll often question whether it is normal anymore. If we see that something is no longer normal, then we'll often expect its regularity to fade over time as well.

Seventh, in discussing the inferential structure of normality judgments, there are a few endemic tendencies each of which contribute to what I regard as mistakes in the literature. Researchers tend to ignore incomplete normalities, they tend to assume that all kinds are created equal and therefore that genuses work essentially the same way as species, and they tend to be insensitive to differences between universal judgments, default judgments grounded in probabilities or mere frequencies, and

default judgments grounded in normalities. Each of these tendencies leads to a corresponding tendency to focus on misleading examples in fields such as default logic and other nonmonotonic logics, in the inferential structure of generics, and in other fields pertaining to subject matters such as *ceteris paribus* laws.

Finally, normality is a rich and complex topic replete with philosophically urgent insights and questions. Most saliently given today's political climate is the potential for a study of normality judgments to illuminate a form of social judgment people often make in popular discourse. Conversations often take a similar form to the following:

> White people (in America) have many privileges [a normality judgment, as it turns out].
>
> That's not true, I grew up extremely poor and have never broken $40,000 income in a year.

This sort of apparent counterexample doesn't refute the original claim because the original claim has semantics along the lines of "when white people do well (financially, socially, politically, etc.), it is explained by their being white." It is not, nor does it imply, any statistical or universal claims, and so any apparent counterexample will be a *merely apparent* counterexample (unless it is a quite specific and often unnatural example). Let us review a set of similar social generalizations:

> Black men are targeted by the police
>
> Lawyers are cunning
>
> Men habitually interrupt women
>
> Cisgender folks have it easier than transgender folks

In short, these sorts of generalizations we make about social groups (along with a host of overtly racist, sexist, and in other respects problematic generalizations) are claims to the effect that *if and when* the ascribed outcome takes place, it is explained by the person's being a member of the social category. Each of these does not, however, have the implication that most or all members of the

social category will have the ascribed outcome. Not all white people will lead lives marked by extreme privilege, not all black men will find themselves the target of policing, not all men habitually interrupt women, and not all cisgender folks actually lead easier lives than all transgender folks. Understanding how these generalizations work and their function is in social discourse is central to understanding social discourse itself. The account on offer in this dissertation illuminates these judgments in a way that makes sense of the claims being made while also explaining why the criticisms of these judgments we rehearsed above do not hold water. This may be the most urgent practical payoff of the dissertation.

At the end of this study, it feels that there is much more to be said. For instance, how is it that there is an explanatory link between being a kind of thing and having the normal features characteristic of that kind of thing? What *grounds* this explanatory link? Normality intuitively has something to do with functionality, selection, adaptation, aptness, naturalness, and so on. I'll merely hint at what I take to be held in common by all normal properties: they are properties that individuals of a kind have because those properties play a particular role in helping those individuals meet a demand constitutive of the kind. A kind is akin to a role in a system in that it comes along with norms and standards. Just as one can play a role more or less well and just as one might have certain features because one plays a role (often because those features are good for an individual playing such a role), one can be a more or less perfect member of a kind and therefore one can have certain features because one is a member of the kind for which those features are normal.

There are questions of fineness of grain and accidentality that arise: does a Sumatran Rhino having one horn mean that Sumatran Rhinos with other counts of horns are abnormal? If so, what about the fact that they apparently have one horn only because of genetic drift and not because of the functional benefits of having one instead of two horns?[2] What about eye color? Are these mere

[2] Again, Gould and Lewontin (1994) are the *locus classicus* for this example in the philosophical world.

accidents so that there is no functional difference between them? If so, what in principle separates people born with normal eye colors from someone with genetically purple or indigo irises? Is it a matter of what has been historically established (as perhaps Millikan will say)? These are indeed interesting questions and not ones I intend to answer in these short concluding remarks. I will say two things, though: First, I don't think that it can be decided by statistics and so the statistically deviant nature of someone with purple or indigo eyes won't be the factor that distinguishes these from normal eye colors. Second, it's entirely possible that a feature like purple or indigo could *become* normal as it is bred into the human gene pool. Without a functional difference, any feature could in principle become normal once it has been present in a population for even a short time.

The normal is that which is explained by being the sort of thing one is. Being a sort of thing is like engaging in an activity or playing a role and in that vein having a normal feature is like having a feature that assists or enables engaging in an activity or playing a role—either as a necessary prerequisite or as a matter of so engaging or playing *successfully*. One might take a view that the normal and the normative are decisively separable. I take a different view: normal features are normative features. All of human thought, however, is suffused with normativity, and so one mustn't disparage normality for so being.

References

Ahn, W. kyoung, Kalish, C., Gelman, S. A., Medin, D. L., Luhmann, C., Atran, S., et al. (2001). Why essences are essential in the psychology of concepts. *Cognition*, *82*(1), 59–69.

Alexander, P. (1973). Normality. *Philosophy*, *48*(184), 137–151.

Anagnostopoulos, G. (1994). *Aristotle on the Goals and Exactness of Ethics* (Vol. 43). University of California Press.

Anscombe, G. E. M. (1956). Intention. In *Proceedings of the Aristotelian Society* (Vol. 57, pp. 321–332).

Anscombe, G. E. M. (1957). *Intention*. Harvard University Press.

Anscombe, G. E. M. (1958a). Modern moral philosophy. *Philosophy*, *33*(124), 1–19.

Anscombe, G. E. M. (1958b). On brute facts. *Analysis*, *18*(3), 69–72.

Anscombe, G. E. M. (1979). Under a Description. *Nous*, *13*(2), 219–233.

Asher, N., & Morreau, M. (1991). What Some Generic Sentences Mean. In G. N. Carlson & F. J. Pelletier (Eds.), *The Generic Book* (pp. 300–338). Chicago: University of Chicago Press.

Ball, B. (2013). Knowledge is normal belief. *Analysis*, *73*(1), 69–76.

Barrett, H. C. (2001). On the functional origins of essentialism. *Mind & Society*, *2*(1), 1–24.

Bastiaanse, H., & Veltman, F. (2016). Making the right exceptions. *Artificial Intelligence*, *238*, 96–116.

Bear, A., & Knobe, J. (2017). Normality: Part descriptive, part prescriptive. *Cognition*, *167*, 25–37.

Boorse, C. (1975). On the Distinction between Disease and Illness. *Philosophy and Public Affairs*, *5*(1), 49–68.

Boorse, C. (1977). Health as a Theoretical Concept. *Philosophy of Science*, *44*(4), 542–573.

Boorse, C. (2014). A second rebuttal on health. *Journal of Medicine and Philosophy*, *39*(6), 683–724.

Boutilier, C. (1994). Conditional logics of normality: A modal approach. *Artificial Intelligence*, *68*(1), 87–154.

Boyd, R. (1999a). Kinds, Complexity and Multiple realization. *Philosophical Studies*, *95*(1).

Boyd, R. (1999b). Homeostasis, species, and higher taxa. In R. A. Wilson (Ed.), *Species: New Interdisciplinary Essays*. MIT Press.

Burge, T. (2003). Perceptual entitlement. *Philosophy and phenomenological research*, *67*(3), 503–548.

Burge, T. (2005). Disjunctivism and perceptual psychology. *Philosophical Topics*, *33*(1), 1–78.

Burge, T. (2010). *Origins of objectivity*. Oxford University Press.

Calcott, B. (2009). Lineage Explanations : Explaining How Biological Mechanisms Change, *60*, 51–78.

Canguilhem, G. (1978). *The Normal and the Pathological*. Dordrecht, Holland: D. Reidel.

Cartwright, N. (1983). *How the Laws of Physics Lie*. Oxford University Press.

Cartwright, N. (1989). *Nature's Capacities and Their Measurement*. Oxford University Press.

Cartwright, N. (1999). *The Dappled World: A study of the boundaries of science*. Cambridge University Press.

Chierchia, G. (1995). Individual Level Predicates as Inherent Generics. In G. N. Carlson & F. J. Pelletier (Eds.), *The Generic Book* (pp. 176–224). Chicago: Chicago University Press.

Cimpian, A., Brandone, A. C., & Gelman, S. A. (2010). Generic Statements Require Little Evidence for Acceptance but Have Powerful Implications. *Cognitive Science*, *34*(8), 1452–1482. doi:x

Cimpian, A., & Salomon, E. (2014). The inherence heuristic: An intuitive means of making sense of the world, and a potential precursor to psychological essentialism. *Behavioral and Brain Sciences*, *37*(5), 461–480.

Clark, R. (1973). Prima facie generalizations. In G. Pearce & P. Maynard (Eds.), *Conceptual Change*. Dordrecht, Holland: D. Reidel.

Cohen, A. (1999). Generics, Frequency Adverbs, and Probability. *Linguistics and Philosophy*, *22*, 221–253.

Declerck, R. (1991). The origins of genericity. *Linguistics*.

Eckardt, R. (1999). Normal objects, normal worlds and the meaning of generic sentences. *Journal of Semantics*, *16*(3), 237–278.

Fenton-Glynn, L. (2016). Ceteris Paribus Laws and Minutis Rectis Laws. *Philosophy and Phenomenological Research*, *93*(2), 274–305.

Ferejohn, M. (1991). The Origins of Aristotelian Science. *Philosophical Review*, *102*(1).

Gardiner, G. (2015). Normalcy and the Contents of Philosophical Judgements. *Inquiry (United Kingdom)*, *58*(7–8), 700–740.

Gardiner, G. (2018). Legal Burdens of Proof and Statistical Evidence. In J. Chase & D. Coady (Eds.), *Routledge Handbook of Applied Epistemology*. Routledge.

Geach, P. T. (1956). Good and evil. *Analysis*, *17*(2), 33–42.

Gehrman, K. (n.d.). *Generics and Characteristic-based Evaluation (ms)*.

Gelman, S. A., & Bloom, P. (2007). Developmental changes in the understanding of generics. *Cognition*, *105*(1), 166–183.

Gelman, S. A., Cimpian, A., & Roberts, S. O. (2018). How deep do we dig? Formal explanations as placeholders for inherent explanations. *Cognitive Psychology*.

Gelman, S. A., & Hirschfeld, L. A. (1999). How Biological Is Essentialis.

Gelman, S. A., & Markman, E. M. (1986). Categories and induction in young children. *Cognition*, *23*(3), 183–209.

Gelman, S. A., Ware, E. A., & Kleinberg, F. (2010). Effects of generic language on category content and structure. *Cognitive Psychology*, *61*(3), 273–301.

Gould, S. J., & Lewontin, R. C. (1994). The Spandrels of San Marco and the Panglossian Paradigm: A Critique of the Adaptationist Programme. In E. Sober (Ed.), *Conceptual Issues in Evolutionary Biology*. The Mit Press. Bradford Books.

Greenberg, Y. (2007). Exceptions to generics: Where vagueness, context dependence and modality interact. *Journal of Semantics*, *24*(2), 131–167.

Greenberg, Y. (2012). Genericity and (Non) accidentalness. *Recherches linguistiques de Vincennes*, (41), 163–190.

Haukioja, J. (2007). How (not) to specify normal conditions for response-dependent concepts. *Australasian Journal of Philosophy*, *85*(2), 325–331.

Hauska, J. (2008). Dispositions and Normal Conditions. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, *139*(2), 219–232.

Hempel, C. G. (1965). Aspects of scientific explanation.

Hempel, C. G. (1988). Provisoes: A problem concerning the inferential function of scientific theories. *Erkenntnis*, *28*(2), 147–164.

Hempel, C. G., & Oppenheim, P. (1948). Studies in the logic of explanation. *Philosophy of Science*, *15*(2).

Henry, D. (2015). Holding For the Most Part : The Demonstrability of Moral Facts. In D. Henry & K. M. Nielsen (Eds.), *Bridging the Gap Between Aristotle's Science and Ethics* (pp. 169–189). Cambridge: Cambridge University Press.

Horty, J. F. (2012). *Reasons as Defaults. Reasons as Defaults*.

Huizen, J. (2015, July). Rare albino raven murdered. *Audubon*. https://www.audubon.org/news/rare-albino-raven-murdered

Joseph, G. (1980). The many sciences and the one world. *Journal of Philosophy*, *77*(12).

Kahneman, D., & Miller, D. T. (1986). Norm Theory. Comparing Reality to Its Alternatives. *Psychological Review*, *93*(2), 136–153.

Kaplan, A. (1967). A philosophical discussion of normality. *Archives of general psychiatry*, *17*(3), 325–330.

Khemlani, S., Leslie, S.-J., & Glucksberg, S. (2012). Inferences about members of kinds: The generics hypothesis. *Language and Cognitive Processes*, *27*(6), 887–900.

Khemlani, S., Leslie, S.-J., Glucksberg, S., & Fernandez, P. R. (2007). Do ducks lay eggs? How people interpret generic assertions. *Proceedings of the 29th annual conference of the cognitive science society*, 395–401.

Kingma, E. (2007). What is it to be Healthy? *Analysis*, *67*(2), 128–33.

Kingma, E. (2009). Paracetamol, poison, and polio: Why Boorse's account of function fails to distinguish health and disease. *The British Journal for the Philosophy of Science*, *61*(2), 241–264.

Kingma, E. (2016). Situation-Specific Disease and Dispositional Function. *British Journal for the Philosophy of Science*, *67*(2), 391–404.

Krifka, M., Pelletier, F., Carlson, G. N., Ter Meulen, A., Chierchia, G., & Link, G. (1995). *Genericity: An introduction*. (G. N. Carlson & F. J. Pelletier, Eds.)*The Generic Book*. Chicago: University of Chicago Press.

Kripke, S. (1980). *Naming and Necessity*. Oxford: Basil Blackwell.

Lawrence, G. (2006). Human good and human function.

Leplin, J. (2009). *A Theory of Epistemic Justification*. Springer Science & Business Media.

Leslie, S.-J. (2007). Generics and the Structure of the Mind. *Philosophical Perspectives*, *21*, 375–405.

Leslie, S.-J. (2008). Generics: Cognition and Acquisition. *Philosophical Review*, *117*, 1–49.

Leslie, S.-J. (2012a). Generics Articulate Default Generalizations. *Recherches linguistiques de Vincennes*, *41*, 25–44.

Leslie, S.-J. (2012b). Generics. In G. Russell & D. Fara (Eds.), *Routledge Handbook of Philosophy of Language*. Routledge.

Leslie, S.-J., Khemlani, S., Prasada, S., & Glucksberg, S. (2009). Conceptual and linguistic distinctions between singular and plural generics. *Proceedings of the 31st Annual Cognitive Science*

Society, 479–484.

Lewis, D. (1986). Causal Explanation. *Philosophical Papers, II*, 214–240.

Liebesman, D. (2017). Between Logic and the World, by Bernhard Nickel. *Mind*, *127*(505), 284–293.

Martin, C. (2003). An Amputee is Bipedal: The Role of the Categories in the Development of Abaelard's Theory of Possibility. In J. Biard & I. Catach-Rosier (Eds.), *La Tradition medievale des Categories (XIIe-XIVe siecles)* (pp. 225–42). Louvain and Paris: Peeters.

McCarthy, J. (1980). Circumscription—a form of non-monotonic reasoning. *Artificial intelligence*, *13*(1–2), 27–39.

McCarthy, J. (1986). Applications of circumscription to formalizing common-sense knowledge. *Artificial intelligence*, *28*(1), 89–116.

Millikan, R. G. (1984). *Language, Thought and Other Biological Categories*. Cambridge, MA: MIT Press.

Murphy, E. A. (1966). A Scientific Viewpoint on Normalcy. *Perspectives in Biology and Medicine*, *9*(3), 333–348.

Neumann, P. G. (1974). An attribute frequency model for the abstraction of prototypes. *Memory & Cognition*, *2*(2), 241–248.

Nickel, B. (2008). Generics and the ways of normality. *Linguistics and Philosophy*, *31*(6), 629–648.

Nickel, B. (2010a). Generically free choice. *Linguistics and Philosophy*, *33*(6), 479–512.

Nickel, B. (2010b). Ceteris Paribus Laws: Generics and Natural Kinds. *Philosophers' Imprint*, *10*(6), 1–25.

Nickel, B. (2014). The Role of Kinds in the Semantics of Ceteris Paribus Laws. *Erkenntnis*, *79*, 1729–1744.

Nickel, B. (2016). *Between Logic and the World: An Integrated Theory of Generics*. Oxford University Press UK.

Odling-Smee, F. J., Laland, K. N., & Feldman, M. W. (1996). Niche construction. *The American*

Naturalist, *147*(4), 641–648.

Odling-Smee, F. J., Laland, K. N., & Feldman, M. W. (2003). *Niche construction: the neglected process in evolution*. Princeton University Press.

Pearl, J. (1988). Probabilistic Reasoning in Intelligent Systems. Representation & Reasoning. Morgan Kaufmann Publishers.

Pelletier, F. J., & Asher, N. (1997). Generics and defaults. In *Handbook of logic and language* (pp. 1125–1177). Elsevier.

Pettit, P. (1999). A Theory of Normal and Ideal Conditions. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, *96*(1), 21–44.

Pietroski, P., & Rey, G. (1995). When other things aren't equal: Saving ceteris paribus laws from vacuity. *British Journal for the Philosophy of Science*, *46*(1), 81–110.

Prasada, S. (2000). Acquiring generic knowledge. *Trends in Cognitive Sciences*, *4*(2), 66–72.

Prasada, S. (2012). Type-token representations in conceptual representation, (2), 3235–3238.

Prasada, S., & Dillingham, E. M. (2006). Principled and statistical connections in common sense conception. *Cognition*, *99*(1), 73–112.

Prasada, S., & Dillingham, E. M. (2009). Representation of principled connections: A window onto the formal aspect of common sense conception. *Cognitive Science*, *33*(3), 401–448. doi:x

Prasada, S., Hennefield, L., & Otap, D. (2012). Conceptual and Linguistic Representations of Kinds and Classes. *Cognitive Science*, *36*(7), 1224–1250.

Prasada, S., Khemlani, S., Leslie, S.-J., & Glucksberg, S. (2013). Conceptual distinctions amongst generics. *Cognition*, *126*(3), 405–422.

Prentice, D. A., & Miller, D. T. (2007). Psychological essentialism of human categories. *Current Directions in Psychological Science*, *16*(4), 202–206.

Reeve, C. D. C. (1992). *Practices of Reason: Aristotle's Nicomachean Ethics* (Vol. 105). Oxford University

Press.

Reiter, R. (1988). Nonmonotonic reasoning. In *Exploring artificial intelligence* (pp. 439–481). Elsevier.

Rescher, N. (1994). *Philosophical standardism: an empiricist approach to philosophical methodology*. Pittsburgh: University of Pittsburgh Press.

Rhodes, M., & Gelman, S. A. (2009). A developmental examination of the conceptual structure of animal, artifact, and human social categories across two cultural contexts. *Cognitive Psychology*, *59*(3), 244–274.

Rol, M. (2012). On ceteris paribus laws in economics (and elsewhere): why do social sciences matter to each other? *Erasmus Journal for Philosophy and Economics*, *5*(2), 27.

Rosch, E. (1973). Natural Categories. *Cognitive Psychology*, *4*, 328–350.

Rosch, E. (1978). Principles of Categorization. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and Categorization* (pp. 27–48). Hillsdale: Lawrence Erlbaum Associates.

Rosch, E., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive psychology*, *7*(4), 573–605.

Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive psychology*, *8*(3), 382–439.

Rosch, E., Simpson, C., & Miller, R. S. (1976). Structural bases of typicality effects. *Journal of experimental psychology: human perception and performance*, *2*(4), 491.

Salmon, W. C. (1971). Statistical explanation. In W. C. Salmon, R. C. Jeffrey, & J. G. Greeno (Eds.), *Statistical Explanation and Statistical Relevance* (Vol. 69, pp. 29–88). Pittsburgh: University of Pittsburgh Press.

Salmon, W. C. (1977). Objectively Homogeneous Reference Classes. *Synthese*, *36*(4), 399–414.

Schiffer, S. (1991). Ceteris paribus laws. *Mind*, *100*(397), 1–17.

Schubert, L. K., & Pelletier, F. J. (1987). Problems in the representation of the logical form of

generics, plurals, and mass nouns. *New directions in semantics*, 385–451.

Schurz, G. (1997). Probabilistic Default Reasoning Based on Relevance and Irrelevance Assumptions. In D. Gabbay, R. Kruse, A. Nonnengart, & H. Ohlbach (Eds.), *Qualitative and Quantitative Practical Reasoning* (pp. 536–553). Berlin: Springer.

Schurz, G. (2001a). What Is `Normal'? An Evolution-Theoretic Foundation for Normic Laws and Their Relation to Statistical Normality. *Philosophy of Science*, *68*(4), 476–497.

Schurz, G. (2001b). Pietroski and Rey on ceteris paribus laws. *British Journal for the Philosophy of Science*, *52*(2), 359–370.

Schurz, G. (2009a). Ceteris Paribus Laws : Classification and Deconstruction. *Erkenntnis*, *57*(3), 351–372.

Schurz, G. (2009b). Normic Laws, Non-monotonic Reasoning, and the Unity of Science. In S. Rhaman, J. Symons, D. M. Gabbay, & J. P. van Bendegem (Eds.), *Logic, Epistemology, and the Unity of Science* (pp. 181–211). Dordrecht, Holland: Springer Science & Business Media.

Schurz, G. (2014). Ceteris Paribus and Ceteris Rectis Laws: Content and Causal Role. *Erkenntnis*, *79*, 1801–1817.

Smith, M. (n.d.). *Justification, Normalcy and Evidential Probability*.

Smith, M. (2007). Ceteris paribus conditionals and comparative normalcy. *Journal of Philosophical Logic*, *36*(1), 97–121.

Smith, M. (2010). What Else Justification Could Be. *Nous*, *44*(1), 10–31.

Smith, M. (2014). Knowledge, Justification and Normative Coincidence. *Philosophy and Phenomenological Research*, *89*(2), 273–295.

Smith, S. (2002). Violated laws, Ceteris Paribus Clauses, and Capacities. *Synthese*, *130*, 235–264.

Sosa, E. (1991). *Knowledge in perspective: Selected essays in epistemology*. Cambridge University Press.

Sosa, E. (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge (Vol. 1)*. Oxford: Oxford

University Press.

Sosa, E. (2009). *Reflective knowledge: Apt belief and reflective knowledge* (Vol. 2). Oxford University Press.

Sosa, E. (2010). How competence matters in epistemology. *Philosophical Perspectives*, *24*(1), 465–475.

Spohn, W. (2002). Ceteris paribus. *Erkenntnis*, *57*, 373–394.

Strahovnik, V. (2016). Defeasibility, norms and exceptions: normalcy model. *Revus*, (29), 61–76.

Strevens, M. (2000). The essentialist aspect of naive theories. *Cognition*, *74*(2), 149–175.

Strevens, M. (2012). Ceteris paribus hedges: Causal voodoo that works. *Journal of Philosophy*, *109*(April), 1–35.

Strevens, M. (2014). High-Level Exceptions Explained. *Erkenntnis*, *79*(10), 1819–1832.

Strößner, C. (2014). *Die Logik der Normalität: Untersuchungen Zur Semantik von Normalitätsurteilen*. De Gruyter.

Strößner, C. (2015). Normality and Majority: Towards a Statistical Understanding of Normality Statements. *Erkenntnis*, *80*(4), 793–809.

Thompson, M. (2008). *Life and action*. Harvard University Press.

Torrell, J.-P. (2005). *Saint Thomas Aquinas: The person and his work* (Vol. 1). CUA Press.

Toulmin, S. E. (1958). *The philosophy of science* (Vol. 14). Genesis Publishing Pvt Ltd.

Van Fraassen, B. (1977). The Pragmatics of Explanation. *American Philosophical Quarterly*, *14*(2), 143–150.

Van Fraassen, B. (1980). The Pragmatic Theory of Explanation. *The Scientific Image*, (5), 97–157.

Van Fraassen, B. C. (1980). The Pragmatic Theory of Explanation. *The Scientific Image*, (5), 97–157.

Veltman, F. (1996). Defaults in update semantics. *Journal of Philosophical Logic*, *25*(3), 221–261.

Wachbroit, R. (1994). Normality as a Biological Concept. *Philosophy of Science*, *61*(4), 579–591.

Wallace, W. A. (1974). Aquinas on the Temporal Relation Between Cause and Effect. *The Review of Metaphysics*, *27*(3), 569–584.

Winter, M. J. (1997). Are Fundamental Principles in Aristotle's Ethics Codifiable? *The Journal of Value Inquiry*, *31*(3), 31(3) 311-328.

Woodward, J. (2000). Explanation and Invariance in the Special Sciences. *British Journal for the Philosophy of Science*, *51*(2), 197–254.

Woodward, J. (2002). There is no such thing as a ceteris paribus law. *Erkenntnis*, *57*(3).

Woodward, J. (2003). *Making Things Happen: A Theory of Causal Explanation*. Oxford University Press.

Yalcin, S. (2014). Modalities of Normality. *Deontic Modality*, 1–30.