Constitutivism and Natural Normativity in Ethics

by

Samuel Parker Gavin

Bachelor of Arts, University of Auckland, 2009

Master of Arts, University of Auckland, 2011

Submitted to the Graduate Faculty of

the Dietrich School of Arts and Sciences in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2019

UNIVERSITY OF PITTSBURGH

DIETRICH SCHOOL OF ARTS AND SCIENCES

This dissertation was presented

by

Samuel Parker Gavin

It was defended on

July 30, 2019

and approved by

Stephen Engstrom, Professor of Philosophy

Jennifer Whiting, Professor of Philosophy

Karl Schafer, Professor of Philosophy, University of California, Irvine

Dissertation Director: Michael Thompson, Professor of Philosophy

Constitutivism and Natural Normativity in Ethics

Samuel Parker Gavin, PhD

University of Pittsburgh, 2019

This project reconciles two schools of thought about the foundations of morality and rationality: constitutivism and neo-Aristotelian ethical naturalism. Constitutivism takes inspiration from Kant and Hume and accounts for the demands of morality and reason by identifying them as constitutive of rational agency. Meanwhile, neo-Aristotelianism makes sense of morality and reason by identifying them as part of human nature. However, if rational agency is part of human nature then these two schools need not be divided. In fact, constitutivism in its strongest form can be shown to rely on the same foundational theory of normativity as neo-Aristotelianism. In this dissertation, I articulate and justify that common theory of normativity and the form of constitutivism based upon it.

I begin with constitutivism's basic concept: constitutive norms. These are norms that concern what it is to be the kind of thing they govern, such as an agent. The very idea of such norms is threatened by the problem of *violability*: if morality is constitutive of agency, for instance, how is it that agents can in fact violate moral norms? If this question cannot be answered, we cannot explain how constitutive norms are normative *in the broad sense* of providing any kind of standard of evaluation. I argue that a solution can be found in the neo-Aristotelian theory of *natural normativity*, which analyzes norms as violable generic generalizations about a species of living thing. I go on to address an important objection that not all generic generalizations are normative. I show that generics pertaining to living things can be formally distinguished from other generics in a way that shows why the former are normative. I conclude my argument by addressing the question of how constitutive norms on my account can be normative *in the strict sense* of providing an agent with

reasons. I show that well-known challenges for strict normativity faced by neo-Aristotelianism and constitutivism are versions of the same challenge and that solutions from both schools can be combined in defense of my combined approach.

Table of Contents

Preface ..... **ix**

1.0 Introduction **1**

2.0 Constitutivism and Constitutive Norms **7**

2.1 What is constitutivism? **7**

2.1.1 Intuitions **9**

2.1.2 Constitutive and regulative norms **11**

2.1.3 The promise of constitutivism **13**

2.1.4 Constitutivist requirements on an account of constitutive norms **15**

2.2 Accounts of constitutive norms **17**

2.2.1 Direct accounts **18**

2.2.1.1 Subjection accounts **19**

2.2.1.2 Compliance accounts **22**

2.2.2 Indirect accounts **34**

2.2.2.1 Constitutive aims **34**

2.2.2.2 Constitutive functions **37**

3.0 The Generic Account of Constitutive Norms **43**

3.1 Direct and indirect accounts: strengths and weaknesses **44**

3.2 Generic propositions **46**

3.3 The philosophical utility of generic propositions **49**

3.4 Generic propositions and constitutive norms **53**

3.5 The semantic objection **58**

3.6 Aristotelian essences without generic propositions **61**

3.7 Conclusion ... **63**

4.0 How to Derive Norms from Generics ... **65**

4.1 Standards of normality ... **67**

4.2 Relative generic propositions ... **68**

4.3 Absolute generic propositions and living things ... **71**

4.4 The normativity of absolute generic propositions ... **75**

4.5 Consequences for metaethics ... **80**

5.0 The Rational Authority of Constitutive Norms ... **83**

5.1 Challenges to strict normativity: shmagency, the rational wolf and biologism ... **86**

5.2 A joint solution ... **89**

5.2.1 The constitutivist-constructivist approach ... **90**

5.2.2 The neo-Aristotelian approach ... **94**

5.3 Against absolutist constitutivism ... **99**

5.3.1 NN-Constitutivism and analytic principles of reason ... **102**

5.3.1.1 First objection: Haase ... **103**

5.3.1.2 Second objection: Korsgaard ... **105**

5.3.1.3 Third objection: Lott ... **106**

5.4 A new biologism objection ... **108**

5.5 Conclusion ... **112**

6.0 Conclusion: Metaethical Constraints and Possibilities ... **114**

Bibliography ... **117**

List of Figures

Figure 1: Accounts of Constitutive Norms 18

Preface

Bringing this project to fruition has been a long and challenging process and many people have supported me in many different ways, be they intellectual, emotional or financial. I would like to briefly express my gratitude here.

I would like to thank my advisor Michael Thompson for years of enjoyable and instructive conversations and helping me maintain an optimistic outlook on the project. Many of the ideas presented here arose in discussions with him or are built on the foundation of his own published and unpublished thoughts, for which I am in his debt. Credit also goes to the other members of my dissertation committee Steve Engstrom, Jennifer Whiting and Karl Schafer, who gave me very useful comments and advice at several critical junctures.

I would like to thank my fellow graduate students in the Philosophy program at the University of Pittsburgh for their comments in successive dissertation seminars, along with commenters at the Saint Louis Annual Conference on Reasons and Rationality (SLACRR). I would also like to thank anonymous referees at *Res Philosophica* and *Philosophia* for their useful suggestions.

I'd like to give credit to Kathryn Lindeman for partly inspiring this project with her dissertation "Grounding Constitutivism", providing various comments and helping me realize that the defeasible compliance account is just wrong.

I would like to thank the administrators at the University of Pittsburgh Department of Philosophy and the previous department chair Bob Batterman for their invaluable support in navigating the various bureaucratic hurdles on the way to completion.

I would like to thank the Charlotte W. Newcombe Foundation for awarding me a Doctoral Dissertation Fellowship which supported me throughout one of the most important years of writing.

I would like to thank my teachers at the University of Auckland, where I completed my Bachelors and Masters. Thanks to them, I received a broad grounding in many areas of philosophy, setting me on the path to this attempt at a synthesis between philosophical schools. I'd like to make special mention of Rosalind Hursthouse and the late Jonathan McKeown-Green. Rosalind's classes on virtue ethics and ethical naturalism originally inspired my interest in this topic, and her advice and encouragement pushed me to develop my thoughts on it. This dissertation constitutes their most thorough and systematic presentation to date. Jonathan was a great friend and mentor, an excellent philosopher and a very generous sounding board for my nascent ideas about philosophy. Without his encouragement in further pursuing philosophy, not to mention his guidance in applying to overseas graduate programs, I may well not have come this far.

Lastly, I want to thank my family: my parents, for their unconditional support of my goals throughout my life, and my wife Vicky for standing by me on a daily basis and keeping up my morale through thick and thin. For that I am more grateful than I can express.

1.0 Introduction

This project is an attempt to reconcile two schools of thought about the foundations of morality. There is a long tradition, stemming from Aristotle, that makes sense of morality by considering it as a part of human nature. Another tradition, stemming from Immanuel Kant, tries to account for the demands of morality by identifying them as demands of reason. The first tradition considers us as human beings—a certain kind of animal with its own characteristic way of life—while the second considers us just as rational beings and abstracts away from our animal nature.

In contemporary philosophy, the Aristotelian tradition is represented by the neo-Aristotelian virtue ethics of Philippa Foot, Rosalind Hursthouse, Martha Nussbaum and various others, with support from Michael Thompson (e.g. Foot 2001, Hursthouse 1999, Nussbaum 1995, Thompson 2008). This family of theories attempts to find a foundation for moral facts in *natural normativity,* a distinctive kind of normativity that concerns the characteristic way of life of living things, exemplified in statements like “An oak tree ought to have sturdy roots”, "Cows need to be able to digest cellulose" and "Blindness is a defect in lions". When an individual exhibits the characteristic features of its kind, it exhibits what Foot calls "natural goodness”, and neo-Aristotelians hope to identify moral goodness as a special kind of this natural goodness. The view is well expressed by Peter Geach’s words, “Men need virtues as bees need stings.” (Geach 1977)

The Kantian tradition is presently represented by the neo-Kantian constitutivism of Christine Korsgaard and David Velleman along with other views that maintain the focus on reason, but give different accounts of it, such as Sharon Street's neo-Humean constitutivism and Paul Katsafanas' Nietzschean constitutivism (e.g. Korsgaard 2009, Velleman 1996, Street 2010, Katsafanas 2013). These theories take the norms of ethics to be constitutive principles or aims of practical rationality, agency or valuing (with these amounting to roughly the same thing). The principles are constitutive,

in that they simply specify norms that apply to, say, an agent merely in virtue of being an agent. They are norms that concern what it is to *be* an agent, and are taken to involve such things as taking the means to one's ends, or acting according to Kant's categorical imperative.

These two broad schools of thought have developed largely independently and yet bear some striking similarities. Firstly, they occupy a similar logical and dialectical space in the field of metaethics. Both attempt to present a third option distinct from approaches that either reject the attempt to understand morality as part of the natural world or reject the through-going reality of moral norms. Both are rooted in philosophical traditions that have been relatively recently revived in the hope of breaking up the stalemate of current metaethical debates.

The most significant similarity is that both approaches exhibit a concern with practical reason, even while the neo-Kantians give it more emphasis in their theories. The neo-Aristotelians are concerned with human nature, but they are particularly concerned with our nature as *rational animals*. It is true that unlike the neo-Kantians, they focus on natural standards of evaluation for non-rational species first, and then try to extend this account of normativity to humans. But the transition to humans involves a focus on practical reason: it is only when they concern this specific power that natural norms can have any authority over us. This view brings the neo-Aristotelian significantly close to the neo-Kantian view. Meanwhile, the neo-Aristotelian concern with human nature does not seem irrelevant to the neo-Kantian approach. The neo-Kantian school is concerned with the nature and function of practical reason as a faculty or power of a rational being. If human beings are rational beings and this power of reason can be a human power, it seems that considering its characteristic role in human life could be relevant to determining its purpose and mode of operation.

Finally, we should observe that both approaches are concerned to locate the source of normativity in the nature of the kind of thing in question, to put it roughly—in what it is to *be* an agent, or a member of a certain species, for instance. In doing so, both views mean to break down the

supposed barrier between facts and norms. Instead, the facts about the nature of agency, or about the life form, constitute the norms.

My long-standing suspicion, which I intend to substantiate in this dissertation, is that these two approaches are headed in the same direction and once closely examined and amended where necessary, they end up in the same place.[1] In particular, I want to show that through examining the basic structure and commitments of each school, we will see that each is to be founded upon the same basic theory of normativity. Neo-Kantian constitutive normativity and neo-Aristotelian natural normativity come to the same thing and may be given a common account. Upon this common ground, we may reconstruct neo-Aristotelian, neo-Kantian and even Humean, Nietzschean and other alternatives, which will belong within the same family of theories.

Both approaches have a lot to gain from this unification into a single school of thought. Constitutivism has for the most part been more clearly and popularly articulated than neo-Aristotelianism and the possibility of translating neo-Aristotelianism into constitutivist terms would be useful for broadening the understanding of this view. Constitutivism is also subject to a problem that neo-Aristotelianism has the resources to solve. This problem, which will figure prominently in my second chapter, is the challenge of violability. This challenge stems from the fact that it is difficult to see how constitutive norms can be violated—if it is constitutive of agency that I obey the categorical imperative and I fail to do so, it would seem that I cease to be an agent and thus cease to be governed by the norms of agency. And I cannot violate the norm if it does not apply to me at the time I am supposedly violating it. Natural normativity does not have this problem, since it is built into the account that the norms are generic statements that characterize the life form, but do not necessarily characterize every bearer or even most bearers of that life form. The possibility of defect

[1] In a recent work, Matthias Haase (2018) even classifies neo-Aristotelianism as a kind of constitutivism, as does Schafer (2018).

is acknowledged from the start. Now, my thought is that a combined theory of constitutive and natural normativity will be able to inherit this advantage of the neo-Aristotelian approach. In fact, it is the solution of the violability worry that will first point constitutivism in a neo-Aristotelian direction.

Another advantage for both approaches is this: it will be dialectically useful to present each as a variation on a single account of normativity, making it easier to compare the views to each other and to metaethical competitors that reject the appeal to either human nature or a substantive account of practical reason. There will also be some interesting results for both sides. We will see that constitutivism must take on some features of neo-Aristotelianism at the foundational level, like its focus on humanity as a species. Neo-Aristotelianism, in turn, will not be able to derive its typical virtue ethics directly from the foundational theory of normativity, as this theory will be shown to have Kantian or other applications as well.

So much for the aim of this project. I will now explain how I plan to execute it. To see where the constitutivist and neo-Aristotelian approaches converge, we need to return to their most abstract foundations: their accounts of normativity. Relatively little work has been done that focuses directly on the nature of constitutive or natural norms, without immediately considering their apparent implications for moral theory. We need to move slower than this if we are going to get the practical implications right further down the line. This is where my project comes in. In my second chapter, I begin with the constitutivist account of normativity, that is, the idea of "constitutive norms" that apply to something just in virtue of it being the kind of thing that it is. I review the field of constitutivist theories so as to isolate the fundamental motivations of the approach and present three conditions that any account of constitutive norms must meet if it is to remain true to those motivations. These are the constitutivity, broad normativity, and strict normativity conditions. In the remainder of chapter 2, I argue that all extant accounts of constitutive norms fail to meet more than one of the constitutivity and broad normativity conditions. The violability challenge will be central

to this chapter, as it presents the primary obstacle to broad normativity. In my third chapter, I will propose my own solution to the problem: an analysis of constitutive norms in terms of *generic propositions.* Here I draw on and extend the research of theorists like Julius Moravcsik, Michael Thompson and Sebastian Rödl into the logical form of generic generalizations (e.g. Moravcsik 1994; Thompson 2008; Rödl 2012). Constitutivism, if it is to succeed, must take seriously the generic formulation of constitutive norms, such as "Oaks shed their leaves in Fall", "Houses are habitable shelters" and "Agents take the means to their ends". These propositions are not to be reduced to any more familiar analysis such as statistical propositions like "Most oaks shed their leaves in fall" or universally quantified propositions like "For every x, if x is an agent, x takes the means to x's ends". The constitutivist must take these claims to be generic all the way down: the surface form of these sentences indicates that these propositions characterize the essence of a kind, like “oaks”, “agents”, “the domestic cat”, “the game of baseball”, or “human beings”. This is essential if the norms are to be constitutive of the members of a kind without implying that every instance of the kind must have the property attributed to it. This analysis will thus turn out to be the key commitment of constitutivism.

Having introduced the generic account of constitutive norms in Chapter 3 and shown that it can both meet the constitutivity condition and overcome the violability challenge, I proceed in Chapters 4 and 5 to show how it lives up to the broad normativity and strict normativity conditions respectively. Broad normativity is the property possessed by all standards of evaluation, "ought" statements, and instances of normative language in general. It is not a trivial matter to show that generic propositions can possess this property, since not every generic proposition establishes a norm in this sense (for instance, “Gold melts at 1064ºC”). However, in establishing this, we will see that the constitutivist must make moves that strictly parallel those made by the neo-Aristotelians, committing the approach to such things as organisms, faculties and functions, natural excellence and defect. To fulfill the ambitions of constitutivism, constitutive norms must also be norms in the strict sense of practical norms that give agents reasons to act. In Chapter 5, I will argue that constitutive

norms can possess this kind of normativity. In doing so, I will show how the constitutivist and neo-Aristotelian schools can pool their resources to deal with various standing objections to the practical import of the norms they posit (e.g. Enoch 2006, McDowell 1996).

By this point, a unified account of constitutive and natural norms will be before us. I conclude with some reflections on the metaethical implications of adopting this account. In particular I will consider what restrictions this unification will put on the neo-Aristotelian and neo-Kantian approaches, just how much room for variation in accounts of reason and morality this account of normativity allows, and what shape future debates between metaethical schools of thought should take if this common foundation is accepted.

2.0 Constitutivism and Constitutive Norms

In this chapter, I begin with an exploration of the intuitions and motivations behind the family of meta-normative theories collectively called "constitutivism". I present three conditions that any constitutivism must meet in its account of constitutive norms if it is to remain true to those motivations: the constitutivity, broad normativity, and strict normativity conditions. From there, I classify and evaluate the various accounts of constitutive norms currently on offer and argue that all of them fail to meet these conditions. This is preliminary to my own proposal for an adequate account of constitutive norms in Chapter 3.

2.1 What is constitutivism?

The term "constitutivism" is often used to refer to a family of theories about rational or moral standards that try to derive these from the constitutive features of agency. However, these theories are just one possible application of a broader approach that might be applied in any domain of normativity, just as one might find realism or anti-realism in any domain. Recognizing this, Kathryn Lindeman identifies as "constitutivism" any view that holds that "the normative standards governing [some] domain can be explained by appeal to constitutive features of the members of the domain" (Lindeman 2017, 231). These features may be the normative standards themselves, or something like an aim or function that generates them. Accordingly, I define *constitutive norms* as norms that are either identical with or grounded in constitutive features of the kind of thing to which the norms apply. Domains of norms about which one might be a constitutivist include morality, instrumental

rationality, epistemic norms, artifactual norms (e.g. "Sharpness is a measure of a fine kitchen knife") or biological standards (e.g. "A healthy dog needs a properly functioning respiratory system").

I would like to make it clear that in this chapter I am using "norms" and "normativity" in the broad sense that covers all kinds of directive or evaluative statements (using 'ought', 'right', 'wrong', 'good', 'bad', 'excellent', 'defective' etc.). Anywhere such language might be rightly applied is a domain of broad normativity. This is distinct from the strict sense of "normativity" according to which it is just the property of giving an agent reasons.[2] Explaining how normativity is possible in the broad sense is a task that is prior to explaining how it is possible in the strict sense. This goes unacknowledged in most discussions of constitutivism (e.g. Smith 2013, Enoch 2006) that not only focus exclusively on constitutivism in the domain of reasons, but also do not distinguish between the tasks of securing the broad normativity of rational or moral principles and securing their strict normativity. However, if we distinguish these tasks we will able to better tell the difference between objections to meta-ethical constitutivism, some of which are leveled against the broad normativity of constitutive norms (such as the "problem of bad action") and some against their strict normativity (such as Enoch's (2006) Shmagency objection and McDowell's (1996) rational wolf). Thus, I will only be discussing broad normativity in Chapters 2 through 4, that is, constitutivism as an account of broad normativity in various domains. In Chapter 5, I will address the question of how some constitutive norms can also possess strict normativity for rational agents.

Since there are many forms of constitutivism and many ways of articulating the view, I would like to begin by examining the fundamental intuitions that motivate constitutivism and proceed from there to build up a picture of its basic features. This account will provide a background against which

[2] Kantians like Korsgaard typically use the term in this way, along with many other philosophers, such as Scanlon (e.g. 2001). See J. J. Thomson (2008) for an account of the distinction between directive and evaluative norms — that said, nothing turns on it here.

I will evaluate different accounts of constitutive norms for their suitability for the purposes of constitutivism.

2.1.1 Intuitions

The basic intuition of constitutivism is that there is a distinction to be made between two kinds of norms or standards: some standards are somehow internal to the nature of the thing they apply to, and others are arbitrarily imposed from without. Certain things seem to be naturally subject to certain norms just in virtue of the kind of thing they are. Take artifacts, for example. It seems natural to expect that a pen can be used to write. If it cannot, it would be reasonable to judge it a defective pen, supposing there are no special extenuating circumstances. Moreover, this judgement seems more natural than judging a stick to be a bad stick because it is inadequate for writing. This is because the standard of being good for writing seems to have something to do with what a pen is, but nothing to do with what a stick is. Now, if some agent had a project of, say, writing in the sand, and was evaluating sticks for this purpose, it would make sense to judge some better than others. But even in this case, the standard of being good for writing still seems to be imposed from without, while there need be no special circumstances for it to apply to a pen. The standard seems to come along with the kind "pen" and determines whether individual pens are excellent or defective instances of their kind, even if no particular agent plans to write with them. Most artifacts are similar to pens in this way: they are subject to a set of standards just due to the kind of artifact they are.

The above example is not meant to prove that the standards for pens really are internal to the kind. And the task of articulating what it means for standards to be "internal" will come later. The example is just meant to illustrate the intuitions that motivate constitutivism. Let us review a few more examples. The sense that some standards or norms have a special status on account of being

internal to the nature of a kind can also be elicited in other domains. Various species of living things like trees, dogs or bacteria can be evaluated for their health and the proper functioning of their parts. The standards for evaluating these things seem to come along with the nature of the kind. In J. J. Thomson's (2008) terminology, the kinds are "goodness-fixing". For instance, while it seems to be an externally imposed standard that a horse be able to tolerate a rider, being able to run seems to be one of the marks of a healthy and thriving horse regardless of any purpose we might have for the animal. A horse that cannot run is not living up to a certain standard — a standard that applies to it just in virtue of its being a horse.

Mental states like beliefs are also often thought to have certain internal standards — or constitutive norms — that apply to them. Beliefs, it seems, ought to be true, or justified, or constitute knowledge, or some combination of these things. While certain beliefs also might be useful for various other purposes, regardless of whether they are true or false, these epistemic or truth norms seem to be more intimately related to the nature of belief than any other standards.

A final example where certain norms seem to be internal to the nature of a thing, or constitutive of it, can be found in the case of games. Games like chess or baseball are subject to various kinds of rules. Some rules have to do with executing a good strategy in pursuit of winning, or being respectful of your fellow players. Such rules can be violated without completely undermining the game. However, other rules are more central to what it is to play the game and there will be something wrong with the game as a whole if they are not observed. For instance, if a chess player outright refuses to defend their king when the king is in check and instead attempts to make some other move, they are breaking a fundamental rule of the game, not just making a poor strategic decision. Rules like this seem to be the constitutive rules of the game; they define what it is.

2.1.2 Constitutive and regulative norms

Each of the examples in the previous section has traded on a distinction between rules, norms or standards that are internal to a kind and those which are external or imposed from without. This can be expressed more rigorously as the distinction between constitutive and merely regulative norms. Using this distinction, we can express the basic intuition of constitutivism as the thesis that not all norms are regulative and constitutive norms are in fact instantiated in some domain. Examining this distinction will help us get a more precise picture of constitutivism.

The earliest version of the distinction seems to be Searle's distinction, adapted from Kant, between constitutive and regulative rules. Since then, the distinction has been broadened in its application to norms, standards and aims.[3] Christine Korsgaard, for instance, distinguishes between constitutive and external standards. J. J. Thomson articulates a version of the distinction when she compares goodness-fixing kinds with non-goodness-fixing kinds, and yet another version can be found in Philippa Foot's distinction between autonomous or natural goodness and secondary goodness (Thomson 2008, 21; Foot 2001, 26). Incidentally, I take this to be one sign of the similarity between neo-Aristotelianism and constitutivism.

It is Searle's original distinction, however, that gives us the closest we have to a canonical account of the concept "constitutive rule", from which the more modern variants are descended. It is worth quoting Searle here:

> As a start, we might say that regulative rules regulate antecedently or independently existing forms of behavior; for example, many rules of etiquette regulate interpersonal relationships which exist independently of the rules. But constitutive rules do not merely regulate, they create or define new

[3] Ishani Maitra, for instance, indicates that she uses the more general term "constitutive norms" rather than "constitutive rules" for this reason (Maitra 2011, 280n).

> forms of behavior. The rules of football or chess, for example, do not merely regulate playing football or chess, but as it were they create the very possibility of playing such games. The activities of playing football or chess are constituted by acting in accordance with (at least a large subset of) the appropriate rules. Regulative rules regulate a pre-existing activity, an activity whose existence is logically independent of the rules. Constitutive rules constitute (and also regulate) an activity the existence of which is logically dependent on the rules. (Searle 1979, 33)

Searle's purpose in introducing constitutive rules is to apply the concept to the rules of behavior, especially language use. However, the key idea has significance that goes beyond this context, that is, the idea of something the existence of which is logically dependent on the norms that govern it. This notion of logical dependence appears to be at work wherever something like constitutive norms crop up, setting them apart from merely regulative norms. That said, the idea of something logically depending on norms for its existence is not a precise one, and the account of constitutive norms we end up with will depend on exactly how this is spelled out.[4] For now, we can say this much: constitutive norms pertain to the essence of the thing in question, or what it is to be such a thing. They are logically prior to individual instances of the kind in some sense, with those instances depending on the norms for their existence as things of that kind. This distinguishes them from merely regulative norms, the application of which is more or less arbitrary, as they depend on our interest in the thing in question and are logically posterior to the existence of the governed individuals.

[4] To give an example of one complication, it might be observed that a horse's existence doesn't typically depend on its actually being able to run. But this norm does seem to bear some connection with what it takes to count as a horse. This is just one thing that an account of the logical dependence between norms and existence of individuals must accommodate, as will be seen below.

2.1.3 The promise of constitutivism

I will review different accounts of constitutive norms in section 2.2, along with their associated versions of the logical dependence just introduced. However, not every account of constitutive norms is going to be adequate to the purposes of constitutivism and I intend to identify those accounts that are. So, to this end, we need to first consider what the purposes of constitutivism are. Why be a constitutivist about anything? What is the appeal of constitutivism as a meta-normative theory? Why would we be interested in norms that stand in this special relation to the existence of the things they apply to, however it is to be spelled out? There are three reasons that we can point to.

Firstly, constitutive norms seem to be a basis for *objective* normativity in the world independent of our subjective interests. Here I mean normativity in the broad sense, rather than the strict sense mentioned above. Norms that are merely regulative only apply to an object in virtue of special circumstances or our interests in the object and may be called into question or overridden in various ways. For instance, to explain why a Rottweiler should have a docked tail, we need to look beyond the dog itself to the aesthetic preferences of dog shows and breeders, or the requirements of the task the dog was historically trained for. People can and do contest whether we should evaluate Rottweilers by these standards, which are in no sense necessary or normatively inescapable. Constitutive norms, on the other hand, need no special explanation for why they should apply. It is not a matter of choice whether a pen is to be evaluated by its ability to write. If this is not a legitimate standard for evaluating a particular item, then that item is not a pen at all. We could doubt whether being usable for writing is in fact a constitutive norm for pens, but in that case we are doubting our received account of what a pen is, not the force of the norm.

Secondly, identifying certain norms as constitutive seems like it might account for their *rational authority* — their normativity in the strict sense of the term. That is, we might be able to appeal to them to explain not only why an agent is evaluable by a certain standard but also why they have

reason to self-consciously comply with that standard. If this standard is constitutive of agency or rationality or something along these lines, this may render it categorical or inescapable. Now, establishing the strict normativity of some standards for rational agents is a more difficult task than showing that there are broadly normative standards of agency. The latter requires showing that an agent's existence *qua* agent depends on them meeting or being evaluable by a certain standard, but the former requires showing that the same standard is a rule they must consciously follow in their practical reasoning or identify as a reason for acting. Broad norms of agency might not have any rational authority over an agent's decision making, even while they might be inescapable in a certain sense. For instance, it may be that whether an individual follows the categorical imperative is a standard for determining whether they count as an agent. But, as David Enoch (2006) has objected, it remains open to an agent to query why they should care about counting as an agent, especially if the standards of counting as such are particularly substantive. Overcoming this problem goes beyond establishing that constitutive norms are a source of objective normativity. This is why I have postponed a discussion of the strict normativity of constitutive norms to Chapter 5.

Thirdly, constitutive norms are meant to be metaphysically unproblematic. Constitutivism rejects the positing of non-natural or primitive normative facts as too dogmatic, but also attempts to retain realism about the norms of the domain in which it is brought to bear. Constitutive norms promise to balance these two desiderata by being identical with or grounded in descriptive facts about kinds. For instance, a constitutivist about practices will think the norms appropriate to an individual game of baseball are given just in describing the general practice of baseball. A constitutivist about biological standards will take the standards of health in an animal to be nothing over and above what is given in a full description of the life cycle of that animal.

If norms are identified with or grounded in descriptions of kinds, they can be grounded in our metaphysics of kinds and objects, without having to add norms as a separate category in our ontology. The existence of the norms would be metaphysically nothing over and above the existence

of the kinds in question. This would constitute bridging the is-ought, descriptive-evaluative gap and bear some of the standard advantages of naturalism, like explaining the supervenience of normative properties on physical and mental properties. That said, it's important to note that constitutivism may not give us anything so strong as a reduction of the normative to the non-normative. Some constitutivists might want that, but as I have said, some accounts of constitutive norms will *identify* norms with descriptive facts about the kind. Instead of reducing one metaphysical category to the other, this breaks down the distinction between the two. So, constitutivism may end up only being consistent with a non-reductionist form of naturalism. If we have a conception of nature that excludes normative phenomena, constitutivism may end up expanding that conception of nature rather than explaining the normative in terms of what that conception already contains.

2.1.4 Constitutivist requirements on an account of constitutive norms

In the light of these constitutivist motivations, we can formulate several rough requirements an account of constitutive norms must fulfill if it is going to do what constitutivists want it to do.[5] These criteria are somewhat imprecise and don't determine exactly how they are to be fulfilled – but to get any more precise than this would be to overly regiment the idea of constitutivism, excluding too many of the diverse views that come under that title.

> **1. Constitutivity:** An account of constitutive norms must explain how the norms are grounded in or identical with the facts about a kind or its nature and through this, how individual members of the kind

[5] Note that these do not have a one-to-one correspondence with the three motivations.

depend on the norms for their existence as things of that kind.[6] This provides for both the objectivity of the norms and their metaphysical simplicity. The main challenge to constitutivity we might call "externality". This is the threat that the posited norms will be merely external standards grounded in something outside the nature of the individual or its kind.

2. Broad normativity: An account of constitutive norms must explain how the norms can function as evaluative standards at all. This is a rather obvious condition, but not an easy one to fulfill for constitutivism. The most immediate challenge to broad normativity is the problem of *violability*. It can be hard to see how constitutive norms could ever be violated. For instance, if what it is to be a good house is based on what it is to be a house at all, it is hard to see how a bad house could exist. If it is constitutive of houses to be habitable shelters, then this seems to imply that a house that is not habitable is no house at all. And this would mean that the constitutive norm of being a habitable shelter is never violated by *any* house. However, various theorists have observed that it is a condition on the normativity of a principle that it be possible to violate it.[7] To have any normative force, a principle or standard must at least require something that is not already necessarily and trivially provided.

3. Strict normativity: An account of constitutive norms must explain how they can have *rational authority*. This only applies to rational or moral norms, and it is a separate issue from their broad normativity. It is one thing to say that instances of rationality or agency are evaluable by certain standards, but another to say that rational agents are given reasons by those standards. More has to be said to explain how there can be standards by which an individual thing is required to evaluate *itself*. The other two criteria must be met before strict normativity even has a chance. I will observe in passing

---

[6] This brings together in one formula the features of constitutive norms mentioned in Korsgaard (2009), Lavin (2017) and Searle (1979).

[7] This problem has been called the problem of "bad action" for constitutive norms of agency, originating with Peter Railton (1997)**.** It has been presented more generally as an issue for any account of norms by Doug Lavin (2004).

that the main challenge to strict normativity is David Enoch's (2006) "Agency-Shmagency" objection and its variants. One of the benefits of distinguishing the issues of broad and strict normativity is that it allows us to see that even if the Shmagency challenge rules out strict normativity, constitutivism might still have a role to play in the philosophy of normativity beyond metaethics.

These requirements, though not introduced so explicitly in the past, find expression in Korsgaard's statement that "every object and activity is defined by certain standards that are both constitutive of it and normative for it" and Searle's statement that constitutive rules regulate as well as constitute. However, even the more explicit criteria above are not fully precise. They serve to remind us of the aspirations of constitutivism and there may be more than one way for an account of constitutive norms to live up to these. However, they are determinate enough for us to judge accounts as better or worse with respect to them. I will now proceed to present the main accounts of constitutive norms on offer, so as to examine how they stand with respect to these conditions — that is, whether they can live up to the promise of constitutivism.

2.2 Accounts of constitutive norms

We can begin to distinguish accounts by their way of fulfilling the constitutivity condition. Specifically, we can divide them up by how they account for the way in which the existence of an individual member of a kind F *qua* F depends on the norms for Fs. Fs might be baseball games, cats, houses, agents — anything which we might think is subject to constitutive norms.

The first division to be made is between accounts where the existence of the Fs directly depends on the F-norms and accounts where the relation is more indirect. In the former cases, which I will call *direct* accounts of constitutive norms, the existence of Fs depends on a relation they bear to the F-norms, such as compliance with the norms or simply being subject to them. In the latter cases,

which I will call *indirect* accounts of constitutive norms, the existence of Fs depends on their possession of a constitutive feature that makes them subject to the F-norms, such as a constitutive aim or function. A further important division is to be made amongst direct accounts of constitutive norms that make the existence of Fs depend on compliance with the F-norms. The degree of compliance required may be either complete, partial or complete with certain excusable exceptions. I call this last case *defeasible* compliance. To help visualize these distinctions, I provide Figure 1 below.

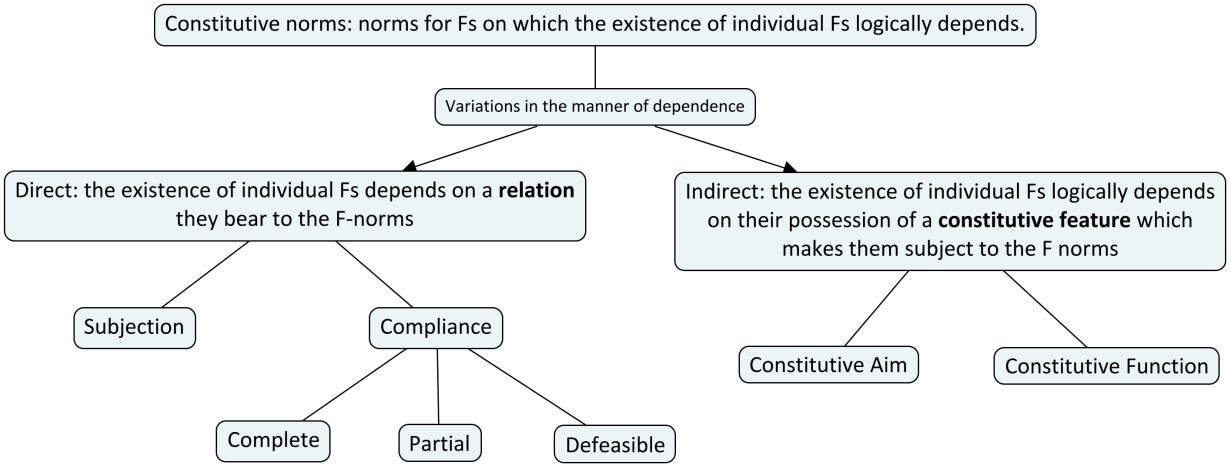

Figure 1: Accounts of Constitutive Norms

2.2.1 Direct accounts

Direct accounts of constitutive norms are perhaps the most natural approach to take, especially if you conceive of constitutive norms as parts of a definition of the thing they govern. Searle gives the following example of one of the constitutive rules of chess: "A checkmate is made when the king is attacked in such a way that no move will leave it unattacked." This is nothing more than a tautology that states what constitutes this particular move in the game. Ishani Maitra, who has a similar

conception of constitutive norms in the context of games, also presents the following as one of the constitutive norms of baseball:

(Three strikes norm) A batter is out when he has earned three strikes. (Maitra 2011, 280)

Definitional rules like this are normative for chess and baseball because they regulate what players are to do in various ways. For instance, a batter must perform the actions appropriate to being out if he earns three strikes, on account of the three strikes norm. These actions will be specified in other rules of the game. In the chess case, the rule makes it the case that a player must accept checkmate and cease play if they cannot free their king from check. Again, the full consequences of this are spelled out in the full set of rules. In Searle and Maitra's view, these rules are not only normative but also constitutive of chess and baseball, in that they specify what it is to play these games, rather than a different game or no game at all (Maitra 2011, 280). The existence of chess and baseball games logically depends on these rules; the games will cease to exist or turn into different games, if they do not bear the right relation to the rules.

2.2.1.1 Subjection accounts

As I indicated above, there are two possible candidates for this "right relation". Firstly, it could be that for a game to be a game of chess or baseball it must be *subject* to the constitutive norms of chess or baseball. These norms must apply to it; it must be appropriate to evaluate it as a poor instance of the game if the rules are broken, or excellent if they are followed. John MacFarlane (2014) presents a subjection account of constitutive norms in a contribution to the literature on assertion. In stating his Truth Rule for assertion ("At a context c, assert that p only if p is true at c") he establishes the

following kind of logical dependence of an assertion on the rule: "To say that the Truth Rule is constitutive of assertion is to say that nothing that is not subject to this rule can count as an assertion" (MacFarlane 2014, 101). MacFarlane gives the same sort of account for moves in games, such as castling in chess: "the rule of chess that says you can't castle if the king is in check is partially constitutive of the move of castling. A move that was not subject to this rule would not be castling" (2014, 101). MacFarlane draws a clear distinction between being subject to a rule and complying with the rule. A game of chess or an instance of castling might not live up to the constitutive norms that define it. "A move may be subject to a rule either by obeying it or by being in violation of it. One can castle incorrectly" (MacFarlane, 101n). This means that constitutive norms, for MacFarlane, need not describe what actually happens in a game of chess, just what ought to happen.

Subjection accounts of constitutive norms are unsuitable for the purposes of constitutivism. This is not a problem for MacFarlane, since he is not trying to defend constitutivism. Nevertheless, for our purposes, it is important to see that on an account like MacFarlane's, the norms fail to meet the constitutivity condition. This condition requires that the norms for something need to be part of its nature in some way — they need to be identical with or somehow grounded in descriptive facts about the kind. However, MacFarlane's account doesn't require that there be any such connection with the facts about Fs, for an individual F being subject to the F-norms. The norms are not required to have any relation to the kind beyond the fact that individuals of the kind are necessarily subject to them. So they need not have anything to do with the descriptive facts about Fs. As far as this account is concerned, there need be no connection between the facts about how people actually play chess, and the fact that the no-castling-in-check-rule is normative for chess games. It could just be a brute normative truth about chess, in which case the metaphysical ambitions of constitutivism have been thwarted. If the norms are purely normative facts associated with the kind without themselves describing the kind or being grounded in such descriptions, then the theory bears no advantage over non-naturalist realism.

Subjection accounts also leave it open that the norms might be external standards, imposed for different reasons in each instance of chess, as long as instances of chess cannot fail to be subject to them. For instance, perhaps in some games players are paid not to castle-in-check, in others they are threatened with violence if they attempt to do so, and in yet others, castling-in-check offends the aesthetic sensibilities of the audience so players prefer to avoid it. And to cover further cases, perhaps we only use the word "chess" to refer to games when this rule gets imposed by some external cause. In all these cases, the rule has nothing to do with the facts about chess or its character as a practice. This makes it impossible for a constitutivist to maintain the distinction between constitutive rules that pertain to the nature of a practice, and merely regulative rules that might be imposed on its participants. If a rule is to be constitutive of chess, there must be something that explains why it is associated with chess more closely than other rules. There must be some fact about chess that accounts for why the existence of a chess game depends on that rule, rather than e.g. the rule of sportsmanship that one ought to be respectful of one's opponent. But MacFarlane's account does not give any reason why it is the no-castling-in-check rule is essential for playing chess, and not the sportsmanship rule. The sort of reason a constitutivist would want is something like this description of the practice of chess: "Chess games are subject to the no-castling-in-check rule because it is part of how chess players currently and traditionally practice chess and part of how they teach the practice to others. Castling-in-check is something players try to avoid as an error in their own play and correct as an error in the play of others".

MacFarlane's account is not incompatible with an explanation like this, but it leaves it open that there is no such explanation and the rule is grounded in something external to the practice of chess. It also leaves open that the rule is just a brute normative fact about chess, which is not the sort of thing constitutivists want to posit. They don't want the rules of chess to be simply associated with it, independently of descriptive facts about the kind of game it is and what people do when playing the game. After all, part of the point of constitutivism is to explain *why* members of a kind are necessarily

subject to a rule, not just to reiterate that they are. Thus, identifying the constitutive norms of Fs as just those to which all Fs are necessarily subject is not sufficient to meet the constitutivity condition.

Again, these criticisms are not directly aimed at MacFarlane, since he is not attempting to defend any kind of constitutivism, but just happens to have a definition of constitutive norms. This definition is perfectly adequate for his own more limited purposes in the philosophy of language. However, it is worth bringing it up to see why constitutivists have not adopted a similar account. To meet the constitutivity condition, it is not enough to observe that constitutive norms are necessary for the kind to which they apply. An account must also show a connection between those norms and the descriptive facts that constitute the nature of the kind. Even if we say that the existence of Fs logically depends on their being subject to certain norms, this does not establish that those norms are part of the nature of Fs. And without that, we cannot keep the promises of objective normativity and metaphysical simplicity.

2.2.1.2 Compliance accounts

In the last section, we considered accounts of constitutive norms that attempt to meet the constitutivity condition by making the individual Fs in question dependent on being *subject* to the F-norms for their existence *qua* Fs. But we might think that a more natural way for a kind to be dependent on norms for its existence would be for its existence to depend on whether it *complies* with the norms. In other words, the constitutive norms for Fs would set a standard such that if an individual F failed to meet that standard, it would be cease to be an F. Accounts like this have a certain intuitive appeal, since we often speak in a way that implies that a game that doesn't follow the rules of baseball isn't baseball at all, or an animal that doesn't go through the normal life cycle of bats isn't

a bat, or a cup of coffee that is sufficiently bad ceases to deserve the name 'coffee'.[8] This also makes sense of the common formulation of constitutive norms as definitions or descriptions of the kind (e.g. "A batter is out after three strikes", "Bats are nocturnal and give birth to live young", "Coffee is dark in color"). If the norm just *is* a description of the kind, an individual both lives up to the norm and counts as one of its kind, by matching the description.

A paradigmatic compliance account of constitutive norms is Christine Korsgaard's. Korsgaard formulates constitutive norms, or principles, as she calls them in the case of action, as descriptions of the kind of object or activity they apply to. For instance, "[I]t is a constitutive principle of walking that you put one foot in front of the other" (2009, 29). "A house is habitable shelter." "The walls are joined at the corners, the insulation goes into the walls, the roof is placed on the top, and so on, so that the weather is kept out, and a comfortable environment is created within. That is the form of a house" (2009, 27).[9] As we should expect, she asserts a logical dependence of the object or activity on the standard or principle. Regarding activities like walking, she says, "in all these cases, we can say that unless you are guided by the principle in question, you are not performing that activity at all" (2009, 29). And in general: "every object and activity is defined by certain standards that are both constitutive of it and normative for it. These standards are ones that the object or activity must at least try to meet, insofar as it is to be that object or activity at all." (Korsgaard 2009, 32)

Such an account has no trouble meeting the constitutivity condition. Korsgaard distinguishes constitutive standards from external (merely regulative) standards by the fact that constitutive standards "apply to the thing simply in virtue of its being the kind of thing that it is", while external

[8] I take this last example from Lindeman (2017).

[9] These constitutive standards for houses are naturally meant as an example of the form of description only. It could be that upon closer examination, houses are too diverse for a single description of them all and we would need to differentiate several types of house and describe the constitutive standards of each of them in turn.

standards do not. A house is subject to standards of providing shelter in virtue of what it is, making these standards constitutive. Other standards like avoiding the obstruction of any neighbor's view do not apply to the house merely in virtue of what it is and so are not constitutive. Now, the formulation here may seem to reverse the order of logical dependence, making the norms dependent on the being of the thing, rather than making its being dependent on the norms, as we have seen so far. However, we have also seen that Korsgaard appeals to dependence in the other direction as well, when she says that an object must comply with or try to meet the relevant standards to be what it is. Not only is something an F in virtue of following the norms of Fs, but the norms of Fs apply to Fs in virtue of Fs being what they are. Since this "in virtue of" relation goes in both directions, it would better to express it as an *identity* between the norms and the facts about what it is to be an F. The norms themselves state what it is to be an F - that is, they are *constitutive* of Fs. This is what makes compliance accounts a straightforward way of accounting for norms that are both normative and constitutive - they simply identify the norms with facts about the kind.

It's important to note that Korsgaard's language is not as precise as it could be. We might ask, which of these really determines whether something is an F: meeting the constitutive standard, *trying* to meet the constitutive standard, or being *guided* by the standard. Korsgaard does not sharply distinguish between them. Some of the reason for this is probably an attempt to avoid implying that in order to count as an F, something must actually meet all the standards for Fs and thus be a *perfect* F. As an account of constitutive norms, this is not only overly demanding, but also runs afoul of the broad normativity condition. To see why this is the case, let us further examine this account, along with some of the alternative accounts Korsgaard could have in mind here.

Complete compliance: the naive account

The complete compliance account is a natural first pass at explaining constitutive norms but is not popular because it fairly obviously fails to meet the broad normativity condition. This account says that in order to count as an F, an individual must comply with all the constitutive norms associated with Fs, thus making the norms conditions of kindhood — that is, being a member of the kind. So, for example, if it is a constitutive norm of baseball that a batter is out after three strikes, a game that fails to comply with this rule ceases to be a game of baseball. Similarly, if it is a constitutive norm of agency that I follow the categorical imperative and I fail to do so, I fail to act at all. This account immediately runs into problems.

To begin with, the complete compliance account creates an excessively high standard for what counts as (e.g.) baseball or agency. We might not think that the claim that in baseball a batter is out after three strikes implies the claim that in all baseball games all batters that earn three strikes are considered out. After all, a mistake in counting could be made, or someone might cheat. And presumably an agent does not fail at being an agent entirely due to a single mistake. However, this issue can be set aside for now, since the account suffers from a larger problem.

The main problem with this "naïve" account is that it seems to make constitutive norms impossible to violate. This has been called the violability worry, or the "problem of bad action" in the context of agency (Lindeman, 2017, 232). To return to the example of the categorical imperative, if it is constitutive of acting that I *always* obey this imperative, then I do so on any occasion that I exercise my agency. So, if on some occasion I fail to act in accordance with the categorical imperative, I do not exercise my agency on that occasion; I fail to act. However, the norms of agency apply to actions, and if I am not acting on that occasion, the norms do not apply to what is going on with me, any more than they apply to what is going on in my digestive system. And I cannot violate a norm if it does not apply to me at the time I am supposedly violating it.

Now, this is a serious issue because normativity — and thus the broad normativity requirement — requires the possibility of not complying with a norm. If it is possible to get something right, it must be possible to get it wrong. Following Doug Lavin (2004), we may call this the *error constraint* on normativity. Lavin distinguishes several versions of this constraint, the first being the requirement of the mere logical possibility of error. That is, there must be no contradiction in the thought of something failing to meet the relevant standard. I take this constraint to apply to norms of all kinds and it is the one that concerns us here.[10]

It needs to be *logically* possible to violate a norm, just because a norm has to require something substantive of the individual to which it applies. A principle that is logically impossible to violate is one that requires something that has no contrary. It is a trivial truth that nothing ever falls short of such a standard. For instance, suppose there was a principle of agency that said, "Either act in accordance with virtue or don't act in accordance with virtue". This is an empty principle because it doesn't really require anything of me. Since I will be complying with it no matter what happens, it has no normative force. A principle cannot guide my actions if it does not somehow limit what actions are open to me. Similarly, "Act as you please" is an empty principle according to Bishop Butler, just because everyone necessarily does this whenever they act.[11] Assuming an account of action where action is a matter of following your strongest desire, telling someone to act as they please does not constrain their actions in any way, since doing otherwise isn't even conceivable.

This problem has been felt most acutely by constitutivism about agency, since the principles that such theories propose begin to look like they can only tell an agent to do things that it is logically

---

[10] Lavin also identifies a more robust version of the constraint, which Korsgaard applies in the case of agency, that requires the real psychological possibility of someone choosing to violate a norm, if it is to be a real norm. He goes on to argue that such a constraint is not feasible, since it would rule out a conception of ideal agency, which theories like constitutivism rely on. But this is beside the point, since the issue for constitutivism I am describing here is a problem with meeting just the mere logical error constraint.

[11] Example taken from Lavin (2004).

impossible for them to avoid. The injunction to fulfill the constitutive principles of agency and thus count as acting seems like nothing more than the injunction "Whenever you act, act." This is logically impossible to violate, and so does not seem to provide any guidance for *the way in which* I should act. If I ever cease to comply with it, I'm not acting at all and so not doing anything wrong by its lights. It does not constrain my possibilities of action in any way and so is an empty principle, not a substantive norm.

That said, this is also a problem across the board. If a pen that cannot be used to write is not a pen at all, there is no sense in which it is a bad pen. This would mean the standard of writing utility is not a norm for pens, since every pen that exists already necessarily complies with it. For another example, let's assume it is a constitutive norm of chess that the players take turns and no one takes two turns in a row - this norm articulates part of what it is to play chess. If I then proceed to take two turns in a row, whether accidentally or with the intention of cheating, I may claim that what I did didn't count as playing chess at all, so couldn't have violated any of its rules. This is surely a bad result to get out of an account of the norms of chess.

To summarize, the complete compliance account does not allow constitutive norms to be violated, and thus does not allow them to be normative. However, this account seemed like a perfect way to ensure the constitutivity of the norms (this is, the logical dependence of the individuals on the norms) by making compliance with the norms the direct determinant of kind-membership. The subsequent challenge is to find a way to make the norms violable without giving up on their constitutivity.

Partial compliance: threshold accounts

Here's one possible response to the violability problem:

> Certainly, asking for complete compliance is excessive. But we can solve the problem just by stipulating that constitutive norms only determine kind-membership up to a certain threshold of compliance. An action, or a cup of coffee, must meet a certain minimum standard in order to count as one of those things, but it need not be perfect. Bad action and bad cups of coffee can exist - they only cease to be if they become too bad. That is how we should spell out the logical dependence of the individuals on the constitutive norms: the individuals must comply with the norms to a sufficient degree.

This response attempts to solve the violability problem by requiring only partial compliance with the constitutive norms rather than complete compliance. We saw from our examination of the naive account that we have a problem whenever the conditions for being subject to a standard are the same as the conditions for meeting the standard. In such cases, nothing can fail to meet the standard without ceasing to be subject to it, and thus it is trivially true that everything meets the standard. This "threshold" account partially separates the conditions of being subject to the constitutive norms (i.e. passing the threshold) and the conditions of meeting them, thus making it possible for there to be defective instances of the kind.[12] This seems to allow for the norms to be both constitutive and normative. As in the naive account, they operate as descriptions of the kind that individuals cannot violate on penalty of ceasing to exist as that kind of thing, but they only do so up to a certain threshold. Once that threshold of compliance is met, they operate as violable norms for evaluating individuals as better or worse. This is not just an ad hoc modification to the naive account — it has its own

[12] Views that take this approach endorse what Lindeman calls the Threshold Commitment (2017, 232).

intuitive appeal. It makes sense that a baseball game might fail to respect a rule here and there, but will cease to be baseball if too many are violated. An animal may suffer from all sorts of diseases and abnormalities before it ceases to count as a member of its kind (for instance, by dying). And sometimes coffee is so bad, it's no longer coffee. In general, the view is that if an individual meets enough of the F-norms or meets them to a sufficient degree, it counts as an F. If it falls below this threshold, it ceases to be an F.

The partial compliance or threshold account provides a popular way of thinking about constitutive norms that is often tacitly assumed by theorists.[13] However, I will argue that in trying to fulfill both the constitutivity and normativity conditions at once, threshold accounts fail to meet either. Let's consider an example. Let's say kitchen knives are subject to a constitutive standard of sharpness, such that being sharper, up to a certain degree, makes a kitchen knife better. There is a constitutive norm of the form, "Kitchen knives ought to be *this* sharp." On a threshold account, a kitchen knife that falls too far short of that standard may cease to be a kitchen knife at all. We should be able to identify a threshold of sharpness below which something is simply not a kitchen knife, and above which something becomes a better kitchen knife. To count as a kitchen knife, perhaps something only has to be sharp enough to cut butter, but not a tomato or a mango. A kitchen knife that cuts butter is good enough to at least continue to count as a kitchen knife, but it still has a long way to go to achieve excellence.

Let's consider different degrees of the sharpness standard for kitchen knives. The standard of being sharp enough to cut butter is very important for kitchen knives, and we might consider it separately from the standard of being sharp enough to cut the leathery skin of a mango. The butter standard lies on one side of the threshold of kindhood. Kitchen knives cannot fall below this standard without ceasing to exist. The mango standard lies above the threshold and can be violated. For kind

[13] Lindeman (2017) finds more explicit commitments to it in Korsgaard (2008a) and Katsafanas (2013).

subject to constitutive norms, we could draw up a similar division between the norms that fall at or below the threshold and must be met for an individual to exist, and those which fall above the threshold and are required for excellence but have no bearing on the existence of individuals.

Now, it could be that the norms do not neatly fall into those above and below a consistent threshold. For instance, perhaps the norms have relative weightings or conditional relationships, such that something has to meet any 5 out of 10 norms, or if sharpness increases, safety becomes more important. Even so, it is still the case for a given individual at a given moment, there are some norms such that if they were violated the thing would cease to exist, and others that if violated, the thing would become worse. So a division of norms is always possible, even if the threshold is complex and variable.

The problem with this division is that it also divides normativity and constitutivity. The norms above the threshold are violable and thus can be normative. However, there is no sense in which the being of individuals is dependent on them, so they do not meet the constitutivity condition. Meanwhile, the norms below the threshold do meet the constitutivity condition, since individuals must comply with them to exist, but since they cannot be violated by existing individuals, they fail to meet the normativity condition. Thus, on the partial compliance account, a given norm for a given individual at a given time is thus either constitutive or normative, but not both. However, we want to be able to assert that a norm is both constitutive and normative for the same individual at the same time. Otherwise, the norms cease to be able to be used to evaluate an individual as soon as they become crucial for that individual's existence.

In summary, the threshold account of constitutive norms, in attempting to meet both the constitution and normativity conditions, ends up splitting the norms into two, one set that is constitutive but not normative, and another set that is normative but not constitutive.

The defeasible compliance account

Despite the failure of the threshold account, it does seem to be on to something important. Constitutive norms must be able to be violated, in order to be norms, but there also must be a limit to this. It shouldn't be the case that a pen can completely fail to meet every standard associated with pens and still count as a pen, or that a baseball game can ignore every rule of baseball and still count as baseball. At the least, such a limitation seems plausible, and it also seems necessary for maintaining the dependence of Fs on the F-norms for their existence qua Fs. However, if this is so, it seems like we should want the norms to be violable, to maintain normativity, but also not violable, to maintain constitutivity. One way of avoiding a complete contradiction here would be to distinguish two kinds of violation, one that leads to ceasing to be a member of the kind and the other to merely becoming a worse member of the kind. This avoids following the partial compliance account in setting a quantitative threshold of violation that would divide the norms into two sets. Instead, it might allow a single norm to play the two functions of determining kindhood and evaluating individuals without these two functions interfering with each other.

Ishani Maitra (2011) provides such an account in a discussion of the constitutive norms of games. Maitra considers the three strikes rule mentioned above: A batter is out when he has earned three strikes. Maitra allows that not every game of baseball actually complies with this norm. This norm might be violated by accident. For instance, the umpire might fail to pay attention and allow a fourth strike before declaring the batter out. This makes the game a poorer instance of baseball and the actions of those involved can be rightly seen as errors, but it does not prevent the game from continuing and so is a properly violable norm. Allowing this leaves room for the norm to be properly normative for baseball, since existing games can fail to live up to it.

Maitra also recognizes the possibility of *flagrant* violations of norms. In the case of games, such violations are those that are intentional and sufficiently marked by everyone concerned. For instance,

a flagrant violation of the three strikes rule would occur if a batter refused to step down after three strikes, demanding an additional attempt. Unlike an accidental violation of the rule, this kind of violation actually stops the game. It is not just that an error has occurred — there is now a failure to play baseball at all. The only way to continue is to start the game anew, without the rogue batter, or to change the rules to accommodate a fourth strike, in which case a new version of the game is underway. If everyone acquiesces without any rule change, that indicates that it was a game of "four-strikes baseball" all along. So, either the game ceases for everyone, the game ceases for the rogue batter, or the game is transformed into a different kind of baseball or re-identified as a member of that kind. Thus, a flagrant violation of a constitutive norm does not make an activity worse but causes it to cease to be as a member of the kind, or else indicates that it is being evaluated by the standards of the wrong kind.

If this distinction between accidental and flagrant violation makes sense, it allows the three strikes rule to determine what counts as baseball while also being violable. It is a *constitutive* norm because it cannot be flagrantly violated without destroying the game. This is the sense in which the game logically depends on the norm. And it still counts as *normative*, because it can be used to judge a game better or worse or the actions of players correct or incorrect in the case of accidental violations. In this way, the account attempts to fulfill both conditions, by replacing the quantitative threshold of the partial compliance account with a qualitative difference in violations. This is still a compliance account of constitutive norms because it requires that members of a kind comply with its constitutive norms in order to count as members of that kind. But this requirement on kindhood is *defeasible* in cases where the failure to comply was accidental.

Unfortunately, the defeasible compliance account suffers from a difficulty of its own. This difficulty concerns the concept of a flagrant violation. When exactly does a violation count as flagrant? The marks of a flagrant violation in a game were that it is intentional and sufficiently marked by everyone concerned. But these do not clearly extend to biological, artifactual or psychological cases

of constitutive norms. A more general account would be required. A first pass at an answer might say that a violation is flagrant when the individual ceases to be a member of the kind as a result. But this definition would be circular, since we want to appeal to the notion of flagrant violation to explain how compliance with a norm can affect kindhood. A second attempt might note that being intentional and marked by all concerned indicates that the violation is no accident, and so define a flagrant violation as a violation that is *non-accidental*, or *purposeful*. But firstly, it's not clear that purposeful violation of the rules *must* destroy a game — everyone might simply tolerate it.[14] And secondly, if we ask what makes a violation non-accidental or purposeful, the only general answer to be given seems to be: it's when an event does not in fact violate any norms, but changes or destroys the thing in question. And this once again gives us a circular definition, since we are trying to find out how flagrant violation can do this. A third attempt might try to define a flagrant violation as one that is sufficiently bad. But this would collapse the defeasible compliance account back into the partial compliance account by using a threshold of compliance to determine kindhood.

So it seems that while initially appealing, the defeasible compliance account is unsuccessful. It understands constitutivity in the same way as the other compliance accounts, making the norms conditions of kindhood. It then tries to make this consistent with violability by having this feature only come into play in cases of flagrant violation. But we have seen no way to properly define these cases without making this account a mere repackaging of the partial compliance account.

[14] Thanks to Kathryn Lindeman for pointing this out.

2.2.2 Indirect accounts

We've seen that direct accounts face a number of problems. Subjection accounts, where each individual must be subject to the norms, do not do enough to establish that the norms are grounded in the descriptive facts about the kind, so they cannot meet the constitutivity condition. Meanwhile, compliance accounts have problems with violability, making it difficult for them to meet the broad normativity condition, or both conditions at the same time. Indirect accounts, however, promise to avoid all of these problems by specifying a criterion of kindhood independently of the constitutive norms, while still grounding the norms in that criterion. This criterion is a *constitutive feature* of the kind. There are two main candidates to be found in the literature: constitutive aims and constitutive functions. I will now argue that despite their promise, both aims-constitutivism and functions-constitutivism once again end up separating the constitutive and the normative.

2.2.2.1 Constitutive aims

On the aims-constitutivist approach, an individual is a member of a kind F based on its possessing the constitutive aim of Fs. It is then subject to the F-norms in virtue of possessing that aim. Since this approach regards aims as psychological items, it can only be applied in domains involving conscious agents, such as the domains of practical and epistemic rationality. This limits the use of constitutivism in domains like the biological, for instance, though meta-ethical constitutivists may have no trouble accepting that.

Two particularly noteworthy theorists that take this approach are David Velleman and Paul Katsafanas. Velleman (2000) introduces a constitutive aim of truth for belief. This aim is meant to ground the truth norm that belief is subject to. In later work, Velleman (2009) also introduces a

constitutive aim of self-knowledge for action, which in turn grounds reasons for an agent to act in a way explicable by their beliefs and desires. He goes on to argue that this self-consistency norm has "Kinda-Kantian" consequences. Paul Katsafanas makes use of the notion of constitutive aims of action in support of his Nietzschean constitutivism (presented in full in Katsafanas 2013). For Katsafanas, these aims are agential activity and 'power' (roughly, the overcoming of resistance). These aims generate norms that make ideal a life of continuous progress and attainment and rule out extreme egalitarianism, alongside other consequences.

Indirect accounts in general have no trouble meeting the broad normativity condition, since it is perfectly possible to have an aim or a function and not fulfill it, thus ensuring violability. At the same time, aims are meant to generate norms. Katsafanas presents this feature as a principle:

> *Success*: If X aims at G, then G is a standard of success for X.

X is here any object, action or activity, and Katsafanas takes the principle to be an uncontroversial one, following from the concept of an aim. In the case of belief, the aim of truth makes a propositional attitude a belief, while actually believing the truth is required for a belief to meet its constitutive norms. In accordance with Katsafanas' principle, if belief aims at truth, truth is a standard of success for beliefs. The constitutive norm thus comes along with the constitutive aim without being identical with it. Now, if the aim of truth determines whether something is a belief without any need for compliance with norms, this allows for its violability. And if the truth norm is grounded in this constitutive aim, that takes care of constitutivity. It is not clear, however, if Success establishes this kind of grounding.

Katsafanas equates Success with the claim that aims *generate* standards of success. This latter claim would be sufficient for the kind of grounding needed to meet the constitutivity condition. But

it not so clear that it is the same claim as Success. Andrew Huddleston (2016) argues that Katsafanas seems to be relying on a more substantial principle, which Huddleston formulates as follows:

> *Generation*: A's constitutively aiming at G grounds the fact that the standards of success for A-ing are given by G and grounds the fact that G has the reason-giving character that it does.

To illustrate that this does not follow from Success, Huddleston points out that Katsafanas' principle does not specify the direction of grounding. In other words, Success allows for either the standard of success to be grounded in the aim, or the aim to be grounded in the standard of success. In fact, Huddleston thinks that the latter might even be more plausible. Taking chess as his example, he notes that if there is a norm that specifies the conditions under which one wins, that could serve to explain why chess players have the aim of achieving that condition. This seems to make more sense than the players' aims to achieve checkmate grounding the fact that checkmate is the winning condition. But regardless of which direction the grounding relation should go, the fact remains that Success does not favor one over the other, so Katsafanas must rely on the more substantive Generation claim, which does not so clearly and uncontroversially follow from the concept of an aim.

Katsafanas addresses an objection similar to Huddleston's in his book, the objection that the mere fact that we have an aim does not generate normative conclusions (2013, 57). This is roughly the objection that Huddleston's Generation does not hold. In his reply, Katsafanas reformulates Success as a slightly more substantive principle that is closer to Generation: "if you aim at A, you have reason to take steps toward realizing A". Katsafanas argues that this claim is unobjectionable and is all the constitutivist needs; "So long as we accept Success, the mere fact that we have an aim does entail that we have reasons to fulfill it" (2013, 61). However, this reply misses the true force of the objection. It is not that this principle is incorrect. The problem is that this principle no longer follows from the mere concept of an aim but amounts to an independent normative principle of instrumental

rationality. If aims only generate norms "so long as we accept Success", it is not the case that reasons are entailed by "the mere fact that we have an aim", since they are entailed by the aim *plus* this instrumentalist principle.

Even if we accept the instrumental principle about aims and reasons, we might inquire into its status as a norm — what grounds it? It cannot be generated itself by some constitutive aim plus the instrumental principle, since then it would rely on itself. So the aims-constitutivist account of constitutive norms ends up appealing to a normative principle that is not itself a case of constitutive normativity. Accordingly, the norms supposedly generated by constitutive aims do not meet the constitutivity condition because they are not in fact generated by or grounded in those aims alone but only come about via the intercession of an independent normative principle.

2.2.2.2 Constitutive functions

Another recent indirect account of constitutive norms is the proper function approach of Lindeman (2017). Instead of a constitutive aim, this approach posits a constitutive *function* of a kind F on which the existence of F *qua* Fs depends and which generates the F-norms.[15] Functions can be found in all kinds of organisms and artifacts and so have the potential to expand the domain of constitutive norms beyond the psychological realm to which constitutive aims were restricted.

Lindeman arrives at the proper function approach at the conclusion of an argument against threshold accounts of constitutive norms. In fact, her argument stands against compliance accounts in general, since it maintains that compliance with constitutive norms is not required at all for counting as a member of a kind. Lindeman observes that the attempt to make kindhood depend

[15] Introduced by Lindeman in her dissertation (2014) and mentioned in her (2017).

directly on the constitutive norms must lead to inviolability in one form or another. However, she says there are other ways to link constitutive features of kindhood and the constitutive norms "other than making the norms themselves somehow constitutive of kind-hood" (2017, 23). This leads her to adopt an indirect account where kind-membership is determined by a non-normative constitutive feature, which in turn grounds the norms. Instead of constitutive aims, Lindeman suggests that the constitutive feature should be the proper function of the kind in question. Each member of the kind has this proper function and is subject to the norms of the kind in virtue of it. To give some examples, it is a fact about the human heart that its proper function is to pump blood and if it did not have this function it would not be a heart. This makes it subject to certain norms of strength, responsiveness, structural integrity, and so on. Similarly, to count as a kitchen knife an item must have the proper function of cutting common food items. This function makes it subject to standards of sharpness and being easy to handle.

I have no objection to the idea that proper functions are constitutive features of a kind. However, the challenge for this view, just as with the constitutive aims approach, is to show how the norms are grounded in these features without appealing to some other source of normativity. This is what is required to meet the constitutivity condition. In this case, everything hangs on the given account of proper function. I think it is plausible that on our intuitive everyday concept of function, the possession of such a function establishes a standard of evaluation — the function of a tool or a bodily organ is nothing more than what it *ought* to do. We might formulate this in the style of Katsafanas' Success principle:

> **Function-Success:** if a kind F has a proper function, this establishes a standard of success for individual Fs of performing that function.

If this is just part of the concept of a proper function, we will not have the same problem we had with the constitutive aims approach. The problem there was that any substantive version of the Success principle seemed like an independent normative principle, separate from the concept of what it was to have an aim. However, Function-Success may stand a better chance of being part of the concept of a proper function. If so, we can get norms out of constitutive facts about the proper function of a kind. However, we have only said that Function-Success is a part of our intuitive concept of proper function. An account of function does not come close enough to our everyday concept will not benefit from this intuitive entailment and an additional normative principle will be needed to generate any norms.

Lindeman herself adopts the *etiological* account of proper function, a common account of function in the philosophy of biology. On this account, a function is the causal role played by a trait that historically promoted its own reproduction in later members of the kind in a lineage. The lineage might be one of reproducing organisms or repeatedly produced artifacts. Etiological proper function is thus determined by the causal history of the kind rather than the present functioning of individuals. Let us return to the example of the human heart and its function of pumping blood. According to this view, a particular heart has this function not because it is currently pumping blood — indeed, it might not be. Instead, it is because earlier iterations of the heart pumped blood, which caused organisms with this feature to reproduce and pass it on in accordance with natural selection, eventually producing the heart in question. Since an individual is not required to perform the proper function of its kind to count as a member of the kind, but only to have a certain kind of history, the norms of performing the function can be violated without ceasing to be a member of the kind, thus avoiding the violability problem.

One of the problems with etiological accounts, however, is that in depending on the idea of a past or ongoing process of selection, they do not have much to do with our everyday notion of function. Ruth Millikan, one of the original formulators of the etiological (or historical 'selected effect') account,

says her account is not intended to be an analysis of our concept of function but a theoretical definition which, if fruitful, might be used to revise that concept (Millikan 1989). After all, whatever the function of say, the appendix, might have been that caused it to evolve, on our ordinary notion of function, the function it has now might be something else, or it might not have one. Similarly, while the genes that produce sickle cell disease might have a causal role that caused them to evolve (they prevented malaria), we would ordinarily consider these genes and the cell structure they produce to be a defect, rather than a functional trait.

Another way to put this objection is that it's not clear that the etiological account of proper function entails Function-Success. Suppose that given a kind F, we have the following fact about Fs: earlier members of the F lineage performed a function k that caused k to be replicated in later generations. This is supposed to ground the following standard of success: Presently existing Fs should perform k. The connection between the two is certainly not obvious — it is valuable, or correct, or excellent for Fs to do something simply whenever a previous generation's doing it caused later generations to do it as well? This seems open to question. At the least, it does not seem to be a definitional truth that those facts about the history of a kind establish its present norms. And if we need some sort of principle to support Function-Success asserting the value of doing whatever historically caused replication of itself, then the account finds itself in the same predicament as the constitutive aims approach and its instrumental principle.

The problem here may stem from the assumption that etiological proper functions are suitable to be a constitutive feature of a kind at all. Constitutive features constitute the nature of a kind and an etiological proper function is a feature of the history of a kind. These are two different things: the nature of a kind is not constituted by its history, or at least not its history alone. Its history brought about its nature, but cannot be identical with it — an individual or a kind cannot be a history alone, since it must also have some constitution in the present. That something had a certain history may entail certain things about how it is now, but is not a feature of how it is now. And it is going to be

difficult to ground presently existing norms regarding how a kind should be in facts about the kind that do not concern its present nature.

In any case, it doesn't seem like there is anything in the definition of etiological proper function that such functions generate norms. We could choose to evaluate current members of a kind against the historical requirements of natural selection, and this might prove useful for the purposes of biology, but that would be our choice and a result of our interests, rather than a standard internal to the kind itself.

How do other accounts of biological function fair? For instance, there are accounts that make function a matter of a trait's present causal role in a system, such as the 'dispositional' account. For example, on this account the heart's function is to pump blood because that is what it is presently disposed to do. It is more intuitive to suppose that what an organ is disposed to do is constitutive of it than the history of what previous copies of it were disposed to do. However, such accounts leave little or no room between what something ought to be doing and what it does — a complete compliance account vulnerable to the violability challenge. Allowing for exceptions will yield a threshold account or a defeasible compliance account with their attendant problems.

The 'organizational' account, a hybrid between the dispositional and etiological accounts, defines a biological function as the role something plays in *presently* keeping itself in being within a self-maintaining system (Mossio, Saborido and Moreno 2009). However, this account has difficulty accounting for the function of traits (such as sex cells) that involve the production and maintenance of traits across generations (Artiga and Martinez 2016). One possible solution would be to regard the *kind* rather than the individual as the self-maintaining system. This thought of a self-maintaining kind corresponds with the neo-Aristotelian conception of a species or life form. Such a kind has traits that when instantiated, serve to ensure the kind remains instantiated, which may include organs, sex cells and so on. This brings us around to the neo-Aristotelian account of proper function (established by Philippa Foot (2001) and Michael Thompson (2008) and defended at length by Micah Lott (2012a).

On the neo-Aristotelian account, the proper function of part of an organism is the role it plays in the life of that organism that allows it to survive and reproduce in the manner characteristic of the species. For instance, the function of the human heart is to pump blood, because this is how it characteristically contributes to the survival of the human organism and continued instantiation of the kind. However, if it's "characteristic of the kind" for its members to have parts that perform certain functions, it's no longer obvious that something can be a member of the kind without performing its appropriate functions. The original benefit of having an indirect account is lost, which was its ability to easily avoid the violability challenge. Now, the neo-Aristotelian solution to this problem is the characterization of the kind with generic generalizations, and it is to a version of that solution that we will turn in Chapter 3.

To conclude this section and this chapter, we have seen that indirect accounts of constitutive norms do not fare much better than direct accounts. They either fail to meet the constitutivity condition because they cannot come up with a constitutive feature of a kind that grounds norms by itself, or they threaten to slide back into compliance accounts in which the constitutive facts about the kind are identical with the norms. Such accounts, along with direct accounts in general, had little trouble with establishing constitutivity by identifying the norms with facts about the kind, but could not reliably overcome the violability challenge to broad normativity. Thus, I conclude that no account yet meets the constitutivity and broad normativity conditions and the promise of constitutivism goes unmet. In the next chapter, I will propose a new account of constitutive norms that is better suited to meet these conditions.

3.0 The Generic Account of Constitutive Norms

In Chapter 2, we reviewed the currently available accounts of constitutive norms and found them all inadequate for the purposes of constitutivism. We saw that constitutivism requires that constitutive norms exhibit both normativity and constitutivity. A necessary condition on the former was violability. Compliance accounts had trouble providing this, as on these accounts the constitutivity of the norms consists in the necessity of compliance for members of a kind, whether complete, partial or defeasible. Meanwhile, subjection accounts and indirect accounts had no trouble with violability, but could not establish constitutivity: the grounding of the norms in descriptive facts about the kind to which they apply, or their identity with such facts. These accounts could not rule out the need for additional normative principles to generate the norms in combination with the facts about the kind.

In this chapter, I want to see how we can salvage the best features of the direct and indirect approaches while avoiding their weaknesses. To this end, I will propose a new account of constitutive norms. I call this the *generic account,* and its foundation introduced in this chapter is the analysis of constitutive norms as a kind of *generic generalization* about a kind. Adopting this analysis can reconcile constitutivity and violability. However, it will take more work to prove that we can get broad normativity out of generics, so apart from establishing violability, that task will have to wait until the next chapter.

First, a few words about the structure of the chapter. In section 3.1, I will review the strengths and weaknesses of the accounts of constitutive norms we rejected in Chapter 2. In sections 3.2 and 3.3, I will introduce *generic propositions* as a possible analysis of constitutive norms and examine their philosophical appeal. In section 3.4, I will argue that as an account of constitutive norms they maintain the strengths and eliminate the weaknesses of other accounts. In sections 3.5 and 3.6 I will

respond to two possible objections to the proposed account on semantic and metaphysical grounds. I will conclude with some remarks on what work remains to be done to defend the account.

3.1 Direct and indirect accounts: strengths and weaknesses

Chapter 2 divided accounts of constitutive norms into direct and indirect accounts. Direct accounts took the constitutivity of the norms to reside in a relation between the norms and individuals of the kind to which they apply. The most popular candidate for this relation was *compliance*. Compliance accounts took constitutive norms to be constitutive due to the necessity for individuals to comply with them in order to count as members of their kind. The degree of compliance varied between accounts — the naive account required complete compliance, the threshold account required partial compliance and a third alternative, the defeasible compliance account, required compliance except in cases of accident.

These accounts struggled to explain how the norms could be violable, a precondition for broad normativity. The complete compliance account was a non-starter since the norms were completely inviolable. The partial compliance account divided the norms into two groups: those above the threshold required to count as a member of the kind and those below the threshold. But this resulted in the norms below the threshold being inviolable and the norms above the threshold having no basis for their supposed constitutivity. Finally, the defeasible compliance account failed to come up with a satisfying account of its attempted reconciliation of violability and constitutivity, relying on the concept of flagrant violation which suffered from unclear application and circular definition.

Meanwhile, the indirect accounts had little trouble with violability but struggled to establish the constitutivity of the norms. They appealed to *constitutive features*, which are straightforwardly constitutive by being necessary properties of a kind. However, they could not show how the norms

were grounded in those features in a way that would make the norms constitutive. The candidate features, constitutive aims and constitutive functions, did not clearly ground the norms on their own without the addition of an independent normative principle. These accounts had no trouble with violability since no level of compliance with the norms was required to be a member of the kind — the presence of the constitutive aim or function was the only requirement for that.

A successful account of constitutive norms should try to acquire the strengths of both approaches while avoiding their weaknesses. From the indirect accounts, it would be good to appropriate the idea of grounding the norms in a necessary, universal constitutive feature of the kind, and making this the basis of their constitutivity rather than a compliance requirement on kindhood. It seems that compliance accounts of constitutivity will inevitably create problems for the violability of the norms. However, compliance accounts had less trouble establishing a grounding relation between the norms and the constitutive features of the kind by making the norms themselves constitutive of the kind. This was done by making them identical with certain *constitutive facts* about a kind that describe its nature. Is it possible to combine all these features? Could the norms be constitutive facts that describe the kind *and* be grounded in a necessary universal feature of the kind *and* still not be criteria that must be complied with in order to count as a member of the kind? Perhaps this could be the case if the norms were descriptions of a nature or essence that is necessarily possessed by all members of a kind but may or may not be "actualized" or "manifested" in a given member of that kind. Now, this nature would have to comprise the actual present features of the kind, not just what its members have done in the past as in the proper function account. But if this was possible, the norms would describe the constitutive features of the kind by describing its nature, and they would be grounded in this nature in just the way that a fact is grounded in the state of affairs it articulates, and there would be no compliance condition on kindhood insofar as the features might not be "manifested". Of course, we have not yet seen any kind of norm or any kind of nature that can meet all these conditions. How can a norm describe a necessary feature of all members of a kind without all those members

necessarily complying with it? I think the answer to this question is: *when the norm is a generic generalization about the kind*. I will introduce these generalizations in the next section.

3.2 Generic propositions

In this section, I will introduce propositions of a particular logical form which I call *generic propositions.* They are generalizations of the form "Fs are Gs" which characterize all Fs including those which are not Gs, a feature which may initially seem paradoxical. They can be simultaneously descriptive and normative, and are thus what we need to combine the strengths of the compliance and indirect accounts of constitutive norms. They are both norms for members of a kind and describe the constitutive features of those members, which we may call their *essence*. So, how are such extraordinary propositions possible?

Julius Moravcsik is one of the first to explicitly draw attention to them in his paper "Essences, Powers and Generic Propositions" (Moravcsik 1994). In this paper, he starts by examining the linguistic category of generic sentences, with the following examples:

(1) Beavers build dams.

(2) Humans have eyesight.

(3) Germans drink beer.

(4) Cougars cross this range.

All four of these examples are syntactically the same: they are bare plural generics. However, Moravcsik notes that (3) and (4) are roughly equivalent in meaning with the statistical generalizations "Germans usually drink beer" and "Cougars frequently cross this range", whereas (1)

and (2) are not equivalent in meaning with "Beavers frequently build dams" and "Humans usually have eyesight".

Intuitively, we might say that (1) and (2) say something about the nature of beavers and humans, while (3) and (4) do not necessarily say anything about the nature of Germans or cougars. Furthermore, the truth of (1) and (2) is insensitive to whether the corresponding statistical generalizations are true: if environmental damage rendered most beavers deformed and most humans blind, it would still be true that beavers build dams and humans have eyesight.[16]

Moravcsik calls (1) and (2) *essence-revealing generics*. Moravcsik takes essence-revealing generics to characterize essences on Aristotle's concept of essence and to express a primitive kind of fact in terms of which we might define such essences. That is to say, Aristotelian essences are just the set of properties attributed to a kind by the true essence-revealing generics about it. This view of essence is contrary to the standard contemporary view on which essences are understood to involve necessary properties. Aristotelian essences can be blocked in their expression by abnormal environmental conditions, so while a member of a certain kind necessarily possesses the essence of that kind, the properties associated with it may not be expressed if the right conditions do not obtain. This is one reason they are apt to be expressed with generics of this form. Instead of adopting Moravcsik's terminology, I will refer to essence-revealing generics as *generic propositions*, or GPs. This partly for the sake of simplicity, and partly to emphasize the fact that these are not the same as generic *sentences*, which is what "generics" typically refers to.

Apart from allowing for exceptions in abnormal conditions, Moravcsik observes that generic propositions possess a lawlike force in characterizing a particular kind. They have a modal character such that they say more than that a certain distribution of traits in a population is actual, but that it

[16] To address a potential objection to Moravcsik here, I think 1 and 2 could be interpreted as making the statistical claim, in which case they would be false in the stated circumstances. However, I also think that another interpretation is clearly possible that is not equivalent to this claim.

must obtain insofar as conditions are normal. This means they also support counterfactual conditionals about members of the kind. "Beavers build dams" says something not just about actual beavers do but about what possible beavers would do.

It's important to distinguish generic propositions from *characterizing generics.* Characterizing generics are the kind of generic sentence usually used to express generic propositions. Characterizing generics have been receiving an increasing amount of attention in their own right in philosophical semantics.[17] However, GPs are not always what is expressed by characterizing generics. The following characterizing generics express GPs:

(5) Ravens are black.

(6) Wood burns.

(7) Buffalo form protective circles.

The following characterizing generics do *not* express GPs.

(8) Cats with blue eyes are blind.

(9) Cars with an expired registration sticker are in violation of the law.

(10) Sharks kill people.

(11) Mosquitos carry the West Nile virus.

Keep in mind that GPs make a lawlike statement about a kind in which a property is attributed to the kind that is instantiated by members of the kind, individually or collectively, in normal

[17] For instance, Matt Teichman (2015, 2016) and David Liebesman (2011) provide accounts of characterizing generics as directly characterizing kinds, which is quite germane to their being interpreted as GPs. However, more popular accounts, such as that of Sarah-Jane Leslie (e.g. 2007), take generics to involve a form of quantification.

circumstances. (8) and (9) characterize groupings that seem too arbitrary to be considered proper kinds. (10) and (11) don't seem to say anything about what is *normal* for mosquitos or sharks, but rather point out properties that we find particularly striking or relevant to our practical interests.[18] (8) through (11) don't seem to say anything about the essence or nature of a kind.

As well as being expressed by only a limited range of characterizing generics, GPs are not defined by their expression by these sentences. They can also be expressed by locutions like "Ravens are a sort of animal that has black feathers", "Ravens normally have black feathers", "Ravens characteristically have black feathers", "Ravens have black feathers, all else being equal" or, in an archaic turn of phrase, "It belongs to a raven to have black feathers." Michael Thompson argues that all these are simply ways of making explicit the same content we might express with a standard generic (Thompson 2008).[19] So, the semantics of characterizing generics is not especially relevant to the analysis of GPs, unless it provides some threat to the possibility that generic propositions are *ever* expressed by characterizing generics. Bernhard Nickel (2016) does actually make an objection along these lines, which I will address below in section 3.5.

3.3 The philosophical utility of generic propositions

The interesting properties of generic propositions have attracted attention from other philosophers as well, especially in the neo-Aristotelian tradition. Michael Thompson (2008) takes them to be both crucial for any cognition of living things and to ground norms for them, calling them

[18] The connection between characterizing generics and our psychological tendencies has been theorized by Leslie (e.g. 2008).

[19] Furthermore, Thompson, Moravcsik and Sebastian Rödl (2012) all argue that Aristotle's frequent phrase "*hôs epi to polu*", or "for the most part", is used to flag that a predication is to be read as a GP.

"natural-historical judgements" and the sentences that express them "Aristotelian categoricals". Here he has in mind the kind of general observations that a field naturalist might make about a species, e.g. "Bobcats breed in spring", "*Pinus radiata* grows to full height over 40 years". He takes the primary object that such judgements characterize to be *life forms*, a term which refers to a pre-theoretical conception of species prior to the more precise notions found in biology and corresponding with the Aristotelian notion of a species' essence or form. The life form, characterized by natural-historical judgements, provides a standard by which it is natural to measure the performance of individuals of the species. For instance, if a *Pinus radiata* tree does not grow to full height over 40 years we can infer that there must be something wrong with the specimen or its environment.

Philippa Foot (2001) makes Aristotelian categoricals (henceforth, "ACs") the cornerstone of her account of natural normativity. However, she thinks that it is only a subclass of such sentences, the "teleological" ACs, that really characterize a species' life cycle and provide norms for it. For instance, "The blue tit has a round blue patch on its head" might be a true AC, but it isn't a true *teleological* AC about the blue tit since it appears that the blue patch plays no role in the bird's life cycle; no further feature of the species depends on its presence and thus it does not count as a norm. Thompson does not really disagree with Foot here, since he would say that the teleological ACs are just the ACs that express natural-historical judgements. The others are characterizing generics of some other form, or statistical claims.

Quite apart from their normative import, Thompson and Foot both take generic propositions (under the name of natural-historical judgements or teleological ACs) to be an elementary and ineliminable form of description of living things. Thompson argues at length that they cannot be reduced to statistical claims, claims that are merely normative like "Bobcats ought to breed in spring" or ceteris paribus claims, if the "ceteris paribus" is doing anything more than just flagging that the statement expresses a GP. More crucially, Thompson and Foot both hold that vital processes like locomotion, nutrition and reproduction along with parts like limbs and organs cannot be identified

in individuals without some conception of the general features of the species and that this conception is constituted by a set of GPs.

I will briefly review Thompson's argument here (which Foot endorses). One useful example is mitosis in cells. A case of mitosis in single-celled organisms counts as reproduction, whereas in more complex organisms it counts as growth and self-maintenance. The only way we are able to tell that a creature is reproducing rather than growing is by looking to the wider context in which the process takes place (Thompson 2004, 64). Even if the two cells go their separate ways after mitosis, this isn't enough to tell, since it could just be a failed instance of growth. The wider context against which we need to see the process is the life form, as Thompson says. We need to know what kind of creature the thing is and how such creatures normally live before we can identify the process it is going through. If we learn that organisms of this kind undergo mitosis to reproduce then we can conclude that it's reproducing, not growing. If we learn that the cell is not meant to go its own way afterwards, even though it did, we can conclude that it was growing, but the process failed. This knowledge of how things ordinarily are and how they are "meant" to go for a kind takes the form of GPs, e.g. "Organisms of this kind are constituted by a single cell", "Organisms of this kind undergo mitosis to reproduce." Similarly, when faced with a novel organism, by pointing to various parts and identifying them as its limbs or its eyes or so on, we are committing ourselves to GPs about that kind of organism. If I identify this appendage of this jellyfish as one of its tentacles, I am not just asserting that this one jellyfish has a tentacle, but also implying that tentacles are things that jellyfish have. I commit myself to this being a feature of the life form, something that is normal to existing and possible jellyfish, and this commitment is best expressed with the GP, "Jellyfish have tentacles" (Thompson 2004, 52).

In this fashion, Thompson and Foot argue for the necessity of GPs as a form of description. They are simply those propositions that make explicit the understanding of life forms and their generic features that is presupposed by our talk of biological parts and processes. The notion of a life form is also needed for this reason, as the object characterized by these general features. It is prior to any

more determinate biological conception of species, as it simply designates a kind that is subject to description by GPs.

Sebastian Rödl thinks that GPs have an important role to play even outside the biological domain. He takes them not just to characterize vital processes, but processes in general, like wood burning or stones falling. In *Categories of the Temporal*, he extends Thompson's argument that GPs express our implicit understanding of the broader context in which a vital process takes place (Rödl 2012, 180). *Generic thoughts,* or *time-general laws* as he calls them, characterize not just life forms but *substance* forms, like "wood" or "stones". In order to know that a particular piece of wood is burning up and turning into ash, we must have some *general* knowledge of the process of wood burning, since in this case the process hasn't completed, and indeed, it might be interrupted by a sudden downpour and never be completed. Thus, our knowledge that the wood was undergoing a process with a certain characteristic end point is based on our knowledge of what happens with wood in general. This knowledge can be expressed, once again, with GPs, like "Wood turns to ash when it burns." Similarly, "Stones fall to the ground", "Glass shatters when struck sufficiently hard" or "Gases expand to fill an enclosed space". While GPs like this have no normative properties (a fact that we will return to in the final section of this chapter) they still exhibit the features Moravcsik points out, like lawlike force while allowing exceptions in abnormal conditions.

To give a few more examples to motivate the introduction of GPs, Alice Drewery thinks that lawlike characterizing generics—a category that appears to correspond with GPs—are what can best account for laws of nature in general (Drewery 2005). A sentence like "All copper salts burn with a green flame" has modal and explanatory properties that make it unfit to be analyzed as a mere universal generalization. Instead, Drewery argues it should be analyzed as a version of the generic claim "Copper salts burn with a green flame", which has the lawlike properties needed, with the addition of "All" as a marker that in this case, the exception class associated with the generic is empty. As well as laws of nature, Matthias Haase (2009) has noted that the laws that Kant takes to govern

the use of our faculties, such as the laws of the understanding studied in Kant's logic, both describe the operation of our faculties when nothing interferes and provide a standard that such operation ought to meet. This seems to make generic propositions the best analysis. In the realm of metaethics, Foot and Thompson are not the only theorists to suggest that generics might form the basis of a metaethical theory.[20] Finally, Matt Boyle (2012) argues, citing Moravcsik, that if we are to understand doctrines like the claim that humans are "essentially" rational, we must understand essences in terms of generic generalizations.

3.4 Generic propositions and constitutive norms

In the previous section we reviewed some reasons to think that generic propositions are a distinct form of generalization that plays an important role in our thought. Now, let us see how they can help with the analysis of constitutive norms. We said above that constitutivism needs a way of thinking about constitutive norms on which the norms describe a necessary constitutive feature of all members of a kind while not being met by all members of that kind. This is the only way to combine the strengths of the direct and indirect accounts of constitutive norms and ensure both constitutivity and violability, the prerequisite for normativity. From the direct accounts, we take the principle that the norms themselves must directly articulate the nature of the kind, to ensure constitutivity. From the indirect accounts, we take the principle that this constitutivity cannot require any level of

[20] Kieran Setiya (2012, 142) considers the metaethical position of Natural Constructivism: "For a trait to be a virtue is for creatures of one's life form to believe that it is a virtue." The clause "Creatures of one's life form believe that (some trait) is a virtue" is understood to be a generic proposition. Karl Schafer (2015) considers "Generic Rationalist Constructivism" as a possible constitutivist account of rationality in terms of what considerations a rational being would be moved by (generically).

compliance with the norms but only that the norms be grounded in a constitutive feature, so as to ensure violability.

We are now in a position to see how this is possible, if we take the norms to be GPs. GPs alone can characterize a kind's constitutive features, and be grounded in these features as their description, without implying that any particular member of the kind possesses these features. Let us consider how they are both constitutive and violable. GPs describe the essence or form of a kind, something that all members of a kind possess *qua* members of that kind. Consider some of our examples of constitutive norms from Chapter 2:

(12) Cats have four legs.

(13) Houses are habitable shelters.

(14) A batter is out after three strikes.

(15) In walking, you put one foot in front of the other.

These are already generic sentences. If we take them to be GPs, (12) doesn't just tell us something about just the cats with four legs, or just the normal cats, or the distribution of the trait of having four legs in the cat population. It attributes "being four-legged" to each cat as part of its essence or form, including to those in which it is not currently instantiated, due to injury, deformity, or being in an earlier stage of life in which legs haven't grown yet. Thus, we can say that while *actually* having four legs (i.e. complying with the four legs norm) is not constitutive of being a cat, it is constitutive of being a cat to be a member of a kind that *essentially* has four legs, or in other words, can be generically characterized as having four legs with GPs. It is part of the essence of a house to be a habitable shelter, and something that lacks that essence by failing to fall within the scope of a GP (i.e. (13)) that attributes to it that generic property, is not a house. It is part of what it is to be a batter in baseball that one is out after three strikes. We can now say that this means that (14) is a true GP that characterizes the kind to which all batters belong. This GP characterizes every batter's essence (or at

least, the essence of baseball games, which includes the constitutive features of batters). This is the case even though that essence is not expressed in some batters due to the failure of the agents on which its expression depends. (15) describes what every instance of walking *essentially* is at all times, and what it *actually* is when everything proceeds normally.

From these examples, we can see that there must be two ways of characterizing an individual. There is on the one hand characterizing what we might call the essential, potential or generic properties of an individual. Having these properties is nothing more than being a member of a kind with certain *generic* features, and thus they determine kind-membership. On the other hand, there is characterizing what we might call the occurrent, actual or actualized properties of an individual.[21] Generic propositions do the former kind of characterization, because they attribute generic features to a kind. This is the kind of essence that generic propositions reveal, that leads Moravcsik to call them "essence-revealing generics": not the properties that are *actualized* in all members of a kind, but the properties of the kind itself. Now, the property of being a member of a kind with certain features is a necessary and constitutive feature of all members of that kind. Thus generic propositions are *constitutive facts* about a kind by describing its constitutive features. At the same time, because GPs do not entail that the properties of the kind are actualized in individuals, this makes them violable. The GP analysis of constitutive norms is thus an improvement over the compliance accounts considered above. This is because GPs describe features that are constitutive of a kind and characterize every member of a kind (with respect to its essence), yet do not entail every member's inviolable compliance (with respect to their actual condition). It is also an improvement over indirect accounts because the violability of the norms does not make it hard to ground them in those constitutive features — they simply describe them.

[21] We might just as well speak of Aristotelian potency and actuality.

Now, let us review why this analysis is superior to the subjection account from Chapter 2. The subjection account asserted that a norm is constitutive just when members of the kind it governs are necessarily subject to it. The GP analysis too appears to say that norms are constitutive when members of a kind are necessarily subject to them, in virtue of having them articulate their essence. However, it is just this "in virtue of having them articulate their essence" that makes the GP analysis superior. In section 2.2, we said that the subjection account was correct, but incomplete, as it had no explanation of why any kind was necessarily subject to the norms associated with it. We can now say that a kind is subject to norms just when those norms are true GPs that characterize the kind. The statement that a kind is characterized by a GP is more substantive than the statement of the brute fact that it is subject to a norm. For instance, it has implications about what individuals are like in normal circumstances and influences our expectations about actual and possible individuals of the kind. It is a kind of fact that we have independent philosophical reasons to countenance and that does more than simply constitute a norm. In this way, the generic account goes well beyond the subjection account.

It would also be worthwhile to compare the generic account with specific indirect accounts, to forestall certain objections. One such objection comes from a proponent of the etiological proper function account. On this account, the causal history of an individual makes it a member of a functional kind and the function associated with this kind makes it subject to the norms. The proper functionalist might say:

> It is all very well that GPs can characterize a kind without implying anything about the actual features of the individual. But proper functions can do this too. Furthermore, the proper functionalist realizes that it's important to be able to say in virtue of what something is a member of a kind possessing that function or described by those GPs. This is especially pressing since it cannot be the actual performance of the function or the actualization of the GPs that makes something a member of the kind. The etiological proper function account can answer this question: it is by being the product of a lineage in

which a certain function was preserved by natural or artificial selection that something is a member of a kind with that function. But there is no corresponding story about what makes an individual subject to a certain set of GPs, that is, how an individual comes to possess the essence that makes it a member of the relevant kind.

How, then, does something get to fall within the scope of a set of GPs? How is a generic essence acquired? The GP analysis can provide such a story. In fact, the GPs for a kind themselves articulate the conditions for falling within their scope as generalizations. Amongst the GPs about a kind, we can find characteristic *creation* conditions for members of that kind. Whatever productive or reproductive process is characteristically responsible for producing members of the kind, this will be described by some of its GPs. Some will characterize the creation conditions directly, for example:

Kittens are born after going through a process of fertilization and gestation.

Coal is produced by the carbonization of dead vegetation.

Other GPs may just identify the kind as a member of a general class with common creation conditions.

Walking is an intentional action. (And so comes about when decided upon.)

Baseball games are group activities. (And so come about at the decision of a group.)

Houses are buildings. (And so come into being when built.)

Pens are crafted tools. (And so come into being when crafted.)

Thus, to fall within the scope of a set of GPs and thus be a member of the relevant kind, the individual has to be the product of the creation process characteristic of that kind. In conclusion, the etiological function account is right in thinking that the history of an individual is what determines

its kindhood. But the history that makes something an F is just the form of (re)production characteristic of Fs, as described by GPs.[22]

In the next two sections, I will address further possible challenges for the generic account of constitutive norms.

3.5 The semantic objection

A further objection needs to be dealt with that threatens to rule out any appeal to generic propositions for any purpose. In a discussion of alternative approaches to generics, Bernhard Nickel presents an objection to Thompson's natural-historical judgements, stating that in fact generics never express such judgements (Nickel 2016, Chapter 4). This objection, if it is successful, creates just as much of a problem for GPs in general. Nickel is concerned that if Thompson is right about the existence of natural-historical judgements, "then certain generics express metaphysically basic facts and that just means that no informative characterization of genericity is possible for the propositions expressed by these generics" (Nickel 2016, 107). Since Nickel wants to provide such a theory of

[22] Because they both rely on individuals having a certain history, both the etiological proper function account and the generic account are subject to Swampman-like objections. That is, they are bound to say that freaks of nature like Davidson's (1987) Swampman that do not come about in the normal way but pop into existence, physically identical to members of a kind like human beings, are nonetheless not real members of the kind. In fact, they must say they have none of the organs, functions, or vital processes characteristic of the kind they resemble. Lindeman addresses these "swampobjections" in her defense of the proper functionalist account by explaining away the appearance of normativity in terms of Swampman's interests (2014, 75-81). Even though his organs don't have proper functions, their operations may serve his interests and thus may be judged better or worse for that. I, on the other hand, am inclined to say that on our intuitive understanding of living kinds, a creature that comes about in a completely different way from normal is a borderline case of belonging to the kind, and the norms for it do not clearly apply or fail to apply.

genericity (a metaphysics of generics), he finds it necessary to provide an argument against Thompson. The problem is that any general metaphysics of generics would conflict with some generics, the GPs, having their own primitivist metaphysics. Now, I don't dispute that a GP theorist is committed to true GPs being metaphysically basic facts. However, I'm not so sure that the fact that we cannot analyze them in terms of concepts like processes, mechanisms, explanatory strategies or psychological tendencies means that we can't meaningfully explore the relations they bear to these things. Rödl, as we have seen, thinks that GPs are crucial for understanding processes, for example. And it's not clear to me that a primitivist metaphysics of GPs could not find a systematic place in a broader metaphysics for generics - we would not know until it was attempted. Regardless, let us proceed to Nickel's argument, since it stands independently of its motivations.

Nickel interprets Thompson as holding that natural-historical judgements of the form "Fs are Gs" ground claims of the form "All Fs ought to be Gs", and that this is what their normativity consists in. However, Nickel argues that no form of generic, even natural-historical judgements, could ground such normative statements, because many generics do not state a feature that should be fulfilled by every individual of the kind. For instance, any generic that characterizes a subgroup of a species like a sex is not something that should be fulfilled by all species members. "The lion has a mane" is true, but it is not true that "All lions ought to have manes", since only male lions should (Nickel, 108). "Bees produce honey" is true but it is not true for each individual bee that it ought to produce honey—that's not the job of queens or drones, for instance.

Nickel goes on to consider ways in which Thompson might be able to deny that these are really natural-historical judgements, but we need not consider these, since it's clear that Thompson would want these to be natural-historical judgements, as they are important parts of their species' natural history and still have normative force. If a male lion does not have a mane or a beehive is not collectively producing honey, then something is amiss, even if these norms do not apply equally to every individual of the species. Indeed, we might well add more examples to Nickel's arsenal. "Buffalo

form protective circles" has already been flagged as a case where it is groups of buffalo that ought to form protective circles, and not each individual. Furthermore, most natural-historical judgements only pertain to a certain phase of the life cycle: "Monarch butterflies form chrysalises" only applies as a norm to individuals at the caterpillar stage, not to eggs or adults. And to give one further case, "Aquatic frog eggs hatch into tadpoles which metamorphose into adult frogs" is a true natural-historical judgement but it's certainly not the case that all frog eggs ought to do this. The failure of most eggs is part of the species' characteristic way of perpetuating itself in its characteristic environment.

All these examples are natural-historical judgements, and they are all just as normative as any natural-historical judgement that applies strictly to each and every individual. This latter observation is the key to resisting Nickel's argument. It's clear that the normativity of these judgements cannot reside in their grounding a statement of the form "All Fs ought to be Gs". Their normativity must be more complicated than that. Thompson nowhere commits himself to the norms taking this simplistic form, though admittedly he does say that natural-historical judgements are normative for individuals by way of this inference pattern: from "The F is G" and "This F is not G" we infer "This F is defective in that it is not G", according to a very abstract sense of defect (Thompson 2008, 80). This does not account for different subgroups within the species which are the targets of particular norms, but since Thompson also gives examples of natural-historical judgements with more specific subjects than the species, such as "the female bobcat" or "the mother", he would presumably have no problem with admitting more specific kinds of natural-historical judgement beyond the paradigm case that takes the name of the species as subject. After all, if we consider the changes that every organism goes through in its life cycle, almost no natural historical judgements will attribute properties that all species members should have at all times. Thus, it does not make any trouble for the idea that generic propositions are norms for a kind, that they are not always to be fulfilled by the traits of individuals, but also by groups, phases, sexes or other divisions of the kind.

We should also observe that no one GP gives the full story of how it is to be fulfilled. All you can glean from "The lion has a mane", without any further information about lions, is that some class of groupings of lions ought to have manes and all lions ought to contribute to this in their normal way. Further GPs can be provided to fill in the gaps in this story: adult male lions contribute to the fulfillment of this norm by growing manes, juvenile male lions contribute to it by growing into adult lions, female lions contribute to it by giving birth to healthy male offspring and so on. In general, all the possession of the kind's essence requires for an individual is contribution in the way characteristic of the kind to the expression of its essential traits by itself, other individuals, groups, its own future phases and so on. What this way of contributing consists in will be specified by further more specific GPs. Interestingly, this is not far from Nickel's own account of characterizing generics, which posits various different ways and respects of being normal with respect to a generic trait, so as to make sense of cases like "Lions have manes." To conclude this section, there is no threat to Thompson's natural-historical judgements being expressed by generics. There would be a problem if natural-historical judgements had to always ground norms for individuals only, but we have seen that this is not the case.

3.6 Aristotelian essences without generic propositions

Another question that may arise in response to the GP analysis of constitutive norms is why this form of description is so important once we have Aristotelian essences or forms on the table. Are these not sufficient to ground constitutive norms? Why not consider this just from the perspective of analytic metaphysics rather than coming at them by way of a form of description? Korsgaard might be put forward as an example of someone who appeals to the Aristotelian concept of form as a ground for constitutive standards ("it is the teleological organization or form of the object that supports

normative judgments about it"), but does not appeal to generic propositions to introduce this concept (Korsgaard 2009, 28). On the other hand, she does often formulate constitutive norms as GPs ("Houses are habitable shelters"), and yet runs into trouble showing just how the form of something can be both normative and constitutive for it by relying on a threshold account of constitutive norms, as we saw in Chapter 2. So, I think the answer to these questions is that generic propositions are essential to specifying essences or forms of a kind suitable for grounding constitutive norms.

A good illustration of this can be found in the case of Jeremy Fix's account of constitutive norms in his "Two Sorts of Constitutivism" (forthcoming). Fix's view is very similar to the GP analysis except that he articulates it explicitly as a metaphysics of constitutive norms, intended to identify the *properties* in which they are grounded, rather than their logical form. Fix distinguishes between "norms first" and "nature first" constitutivisms. These roughly correspond to my direct and indirect accounts of constitutive norms. Fix's account is of the latter sort as it attempts to identify properties of an individual that constitute its nature, some of which can then be shown to ground the norms. Some of these properties, according to Fix, are *contingent essential properties* (henceforth, CEPs). Taken collectively, they are basically generic Aristotelian essences as I have described them. CEPs are properties like those that an artifact needs to perform its function (e.g. the properties of having clear lenses for glasses), or what a capacity needs to do what it does (e.g. maintaining balance for our capacity to walk) (Fix, 6). Such properties may not be actualized in every member of a kind, but they still create a standard of success for members of that kind, all of which possess them as part of their essence.

If we grant Fix's claims about CEPs, the fact remains that if an individual has a contingent essential property, this is partly in virtue of certain facts about its kind. The property of having clear lenses is a contingent essential property of a pair of glasses if and only if it is a fact about its kind that "Glasses have clear lenses". Indeed, the fact that something can have an essence that might not be actualized is possible because the properties are primarily attributed to the kind, not the individual. So there is

nothing more to an object X having a contingent essential property than a certain proposition being true of the kind X belongs to. Such propositions are GPs, and as a result, CEPs of individuals are dependent on the truth of GPs about kinds. An individual may possess certain CEPs, but this is nothing over and above its kind having certain features, which amounts to a set of GPs being true. So while we can say, with Fix, that the norms that apply to an individual do so in virtue of the nature of the individual, we may also say the following: the nature of an individual is just the nature of its kind and this nature consists of properties best defined in terms of GPs. In this way, it seems that neither nature nor norms can be said to be prior, and the grounding relation bottoms out in a kind of identity between the nature of a kind and the norms that characterize it.

3.7 Conclusion

Though I have defended the GP analysis of constitutive norms from several objections, there is one issue that has not been addressed. This is the fact that I have not yet provided a complete argument that the account can meet the broad normativity requirement on accounts of constitutive norms. I have only defended the claim that the GP analysis meets the constitutivity requirement and does not fail the broad normativity requirement due to inviolability. Now, the theorists I've referenced do think that generic propositions are or can be normative standards. However, to prove this, more work needs to be done. We've seen enough to establish that generic propositions are standards of *normality*, but we don't yet have an explanation of how a standard of normality can also be a standard of evaluation. While some GPs certainly seem to operate like this, some do not, indicating that we haven't yet struck upon the source of their normativity. For instance, consider GPs about natural inanimate things such as "Snow is white", "Wood burns", "Glass shatters under pressure" or "Gases expand to fill a closed space". Each of these statements tells us what happens

normally and in general but carries no implication that glass that does not shatter or gas that behaves abnormally because of abnormal conditions is *defective*, or that glass that did shatter would be *better* as glass. Yet it is just as clear that other examples *do* have normative implications — if I fail to put one leg in front of the other when walking, this is a *failure*, something has gone *wrong*. Either I've made a mistake or there is some obstacle in my environment causing me to fail, and this just follows from the fact that I'm engaged in walking while not complying with the generic proposition describing what walking involves. So it would seem that GPs always seem to establish what is *normal*, but only sometimes what is *normative,* and we do not yet have an explanation of what makes the difference between the cases. Since the best hope for constitutivism seems to be the GP analysis, more work needs to be done to show exactly when and why these generalizations function as norms. I turn to this task in the next chapter.

4.0 How to Derive Norms from Generics

In Chapter 3, I presented an account of constitutive norms as generic propositions (GPs). I argued that this account could overcome the problems facing all other accounts of constitutive norms and fulfill the conditions laid down in Chapter 2 for any account of constitutive norms suitable for the purposes of constitutivism. However, one condition remained only partially fulfilled: the *broad normativity* condition. This condition requires simply that constitutive norms operate as standards of evaluation of some kind, rather than mere non-normative descriptions. The primary challenge to this in previous chapters was the issue of *violability*, but we have by now given ample proof that GPs are violable, since they admit of exceptions. The challenge now is to show that when exceptions to GPs do occur, these are at least sometimes cases of failure, defect, privation or some other normative status, and when they do not occur, this is good.

Now, neo-Aristotelians like Foot and Thompson have already fixed upon generic generalizations and made them the basis of an account of normativity (Thompson 2008; Foot 2001). Thompson (2003) calls the account *logical Footianism*. This account takes the normative force of GPs to be clear. It is plausible to think that they simply *are* normative when the examples used are about living things, like "Cats have four legs". A cat that does not have four legs is not just abnormal, but suffers from an injury or a defect — this condition is a bad one; something has gone wrong. However, it becomes clear that we do not have the whole story when we look at GPs about naturally occurring inanimate things, as in Rödl's examples from the last chapter "Wood burns" and "Stones fall to the ground" (2012, Chapter 6). These GPs still have the lawlike character and exception toleration of "Bobcats breed in spring" but have no normative implications. Wood that does not burn or stones that do not fall are abnormal cases, but they are nothing more than abnormal. They are not defective, deprived or deficient in any way. There is a standard of *normality* here but no standard of evaluation.

In the light of these examples, any theorist who wants to assert that generic propositions are normative must account for the difference between GPs about animate and natural inanimate subjects. What is it that makes the former normative and the latter non-normative? It is not enough to observe that their subject matter is different, since that would suggest that it is nothing about this form of proposition that grounds normativity but some independent feature of living things. Thus, if we are to defend the thesis that a primitive kind of normativity arises from GPs, we must identify a *formal* difference between GPs used to characterize animate subjects and those used to characterize inanimate subjects that explains why the former are normative and the latter are not.

As a side note, from here on I will refer to naturally occurring inanimate things just as "inanimate things". I do not mean for this category to include artificial inanimate things. Generics pertaining to these things, like "Pens are usable for writing" and "Nests protect the eggs" do not seem to be non-normative, but since these are produced and used by living things, I think the question of their normativity is best dealt with downstream of the topic of the norms pertaining to living things.

Returning to the problem at hand, the task of this chapter will be to give a full account of how the GP analysis of constitutive norms can meet the broad normativity condition. This account will also have an important upshot for constitutivism: it will show that the neo-Aristotelian notion of a *life form* is essential for any theory that relies on normative GPs. As a result, constitutivism will end up being just as committed as neo-Aristotelianism to the idea of life forms as the source of normativity. On this foundational thesis, the two accounts of normativity converge.

The chapter will proceed by drawing a distinction between "relative" and "absolute" GPs, showing how the latter are normative, and then examining the similarities between the resulting account of normativity and neo-Aristotelian natural normativity and what these similarities mean for the further task of extending this account of constitutivism to rational and moral norms.

4.1 Standards of normality

In order to draw a formal distinction between GPs about living and non-living things, we will need a closer analysis of both kinds. First, let us consider a feature common to both: their establishment of a standard of normality. Even non-normative generics, like "Wood burns" or "Snow is white", establish what might called a "norm" in the most ordinary, colloquial sense. They state what happens in general, what happens normally or typically, how things normally are. Each states how things are in *normal conditions*, and leaves open the possibility of permissible exceptions to the rule that could take place in abnormal conditions. Each articulates a *standard of normality*, against which individual cases might be measured in determining whether they are normal and to be expected, or abnormal and in need of a special explanation. Such a standard is clearly not a standard of evaluation, of better or worse, excellence or defect, right or wrong, and so is not normative in the philosophical sense. However, since it will emerge that these standards are close relatives of full-blooded norms, so it would not be wrong to consider them a kind of "proto-norms".

Properly normative generics also establish standards of normality. Consider propositions like "Cats have four legs" or "Kitchen knives are sharp". As well as providing a standard for evaluating excellence and defect, these statements influence our expectations of what actual and possible individuals of the kind will be like and indicate when a special explanation of their condition is required. This is something that comes along with their form as GPs: they have a lawlike quality in characterizing a kind, telling us what it is like. Thus, all normative GPs will also be standards of normality. But, as we have seen, not all standards of normality are norms. Our question, then, is this: when is a standard of normality also a norm?

4.2 Relative generic propositions

Let us now consider the features of non-normative standards of normality. First, some examples:

Wood burns.[23]

Snow is white.

Gold melts at 1064ºC.

Glass fractures under pressure.

Chicago winters are cold.

Submerged objects displace a volume of liquid equal to their volume.

These GPs characterize what an individual of the kind is like, *should it chance to find itself in certain conditions*. Each has associated with it a set of *normal conditions* in which the property mentioned is instantiated in an individual. We can make these conditions explicit with an additional conditional qualifying clause:

Wood burns in the presence of sufficient oxygen and heat.

Snow is white when the air is clean and there is no local algae or sulfurous liquids.

Gold melts at 1064ºC in terrestrial pressure conditions.

Glass fractures under pressure when it is not braced or reinforced and the pressure is sufficient.

Chicago winters are cold in early 21st century climate conditions.

Submerged objects displace a volume of liquid equal to their volume if the liquid has sufficient room to move.

[23] Here we are considering wood just as a material, i.e. dead wood, not a component of a living plant.

These conditional claims are themselves conditional on the individual finding itself in certain conditions in which it is normal for the conditional to be realized. For instance, wood burns in the presence of sufficient oxygen and heat when it is sufficiently dry. We could iterate the process of explicit qualification through adding conditions indefinitely to achieve more and more specific statements of the conditions in which the individuals normally have the stated reaction. But in general, for any such GP, no matter how specific, there remains a background assumption of normal conditions that are not made explicit in the sentence.

We might ask, in just what sense are these background conditions *normal*? They are not themselves generically attributed to the kind, but are just the conditions of a given individual, produced by some set of accidental circumstances. It is not part of the nature of wood to be in the presence of sufficient oxygen and heat to burn. There is nothing abnormal about gold that is not in terrestrial pressure conditions. Dead wood and gold may be more prevalent in certain parts of the universe, statistically speaking, but such facts do not have the lawlike force of a generic proposition. This is because there is no characteristic environment for dead wood or gold to be in, unless we are thinking of our experience of dead wood or gold and the situations in which we are accustomed to encountering them. Thus, when we say that "Gold normally melts at 1064°C", the normal conditions in which this occurs must be those that are "normal" relative to us. Now, it is an objective fact about gold that it responds this way in those conditions, but the choice of those conditions as the characteristic ones which may be left implicit relies on our perspective and interests.

Before we continue, it is worth noting that for some GPs, the standard of normality must be exceptionless because all abnormal conditions that would create an exception are physically impossible. For instance, Drewery's example "Copper salts burn with a green flame", which she takes to be a generic with an empty exception class (Drewery 2008, 374). In such a case, we need not make an arbitrary decision on the environmental conditions to hold fixed in describing it. Since all physically possible conditions will be suitable for the instantiation of the property, all count as

normal. Furthermore, since these GPs that describe physically necessary states of affairs cannot be violated, they are clearly out of the running for being potential norms. For this reason, I will leave this kind of GP aside and focus on those that have physically possible exceptions.

So, with exception of physically necessary GPs, it would seem that the properties generically attributed to inanimate things in GPs are only *essential* to them relative to us, and only constitute what we might call a relative standard of normality. Their exhibiting the traits that we consider normal is ultimately an accident of the circumstances in which a given individual is located. However, since we only find some of those circumstances unusual or in need of explanation from our perspective, we designate the rest as normal. We use GPs to do so since these identify properties as needing no special explanation in terms of external causes. If some gold is not melting at 1064°C, we want to know why, because the cause must be an environmental condition that is very unusual from our point of view. If it is melting at 1064°C, this is unremarkable to us — and so we say that this is just a feature of the kind: "Gold melts at 1064°C". By establishing a relative standard of normality, we describe the kind at the same time as expressing norms for our explanatory practices. However, it will often be appropriate not to suppress the conditions we take to be normal, but rather make them explicit in a conditional clause as presented above, for instance, if we want to compare different environments.

Now, I think the fact that generic propositions about inanimate things only establish *relative* standards of normality is a clue to why these standards do not function as norms for the things they describe but only for our explanatory practices. And it also suggests another possibility in logical space: *non-relative* standards of normality. I turn to this possibility in the next section.

4.3 Absolute generic propositions and living things

We have said that a standard of normality states what members of a kind are like in normal conditions, but we found that in inanimate cases, the conditions designated as normal were only those that were normal relative to us. However, could there also be an objective way of determining the normal conditions for a GP's instantiation, making the standard of normality not relative but absolute? That is to say, what if there were normal environmental conditions that characterized a kind itself, independently of our explanatory interests? All that would be needed for this to be the case would be for the normal instantiation conditions of GPs describing the kind to be *themselves* attributed to the kind with further GPs.

For instance, let us imagine what would follow if gold could be characterized in this way (counterfactually, of course). We could say that "Gold melts at 1064°C" and individual instances of gold melting at this temperature would not have to be explained in terms of a chain of causes that got them into the environment we consider normal for gold. They could be explained just in terms of general facts about gold and the circumstances gold is characteristically in, expressed with GPs, e.g. "Gold is found in terrestrial pressure conditions", "Gold is not alloyed with other metals". In these examples, the conditions that are normal for gold relative to us are instead features of the kind on their own. And there would be further generic facts about gold that could be referenced to explain these facts, without having to reference the particular conditions of individual samples. Of course, gold cannot be characterized in this way, but if it could, it would be a kind for which there is an *absolute* standard of normality, rather than a relative one.

The GPs describing such an absolute standard of normality would be formally different from those describing relative standards of normality. While both sorts of GPs have an associated set of normal conditions in which a property is instantiated, for *relative* GPs, as we might call them, the

normal conditions are merely possible conditions of an individual and are expressed with a conditional clause, as in these familiar examples:

> Gold melts at 1064°C when it is in terrestrial pressure conditions.
>
> If it is in terrestrial pressure conditions, gold melts at 1064°C.

However, if terrestrial pressure conditions were themselves normal for gold, we could eliminate the hypothetical character of the condition, creating a GP like "Gold melts at 1064°C *because* it is found in terrestrial pressure conditions". This statement entails what is the case when both the property of gold and its conditioning environment are normal: that both normally obtain and that there is a normal causal relationship between them. Let us use the term "absolute generic propositions" to refer to GPs that allow for this kind of conditioning relation between different GPs.

So, we have before us the idea of a kind of entity which is characterized with GPs that constitute absolute standards of normality, rather than relative ones. These GPs are formally different from relative GPs in that they can be connected with other absolute GPs with "because". Putting aside for now whether there are any such entities, let us consider, a priori, what such an entity would have to be like. We will find that such consideration can take us quite far. To begin with, such a being would have normal conditions for coming into existence, staying in existence and perhaps even going out of existence, and these conditions would come about in a standard way. In normal conditions, the continued existence of the thing would be no accident and would not require any special case-by-case explanation. And since those normal conditions would normally obtain, it would be normal for an individual of the kind to stay in existence. For this reason, its continued existence could be attributed to the nature of thing itself. In this sense, it would be appropriate to call such an entity *self-maintaining*.

The coming-into-being of such entities would not be an accident either, since as stated, this would also be characterizable with absolute GPs. It would be normal for them to come into being through a certain process, absent interfering accidents. For it to be normal for them to come into being in a certain way, the conditions for their coming into being would have to *normally obtain*, instead of obtaining by accident. It would have to be a feature *of the kind* that its creation conditions come about from time to time — and this amounts to saying that the kind would have to somehow arrange for its own instantiation. Either the individual entities would have to somehow bring themselves into being, which would seem to require a miraculous creation *ex nihilo,* or each individual or group of individuals would have to bring others into being. Thus, our normally existing entities must engage in self-production or reproduction. Let us set the former option aside, since it seems to be paradoxical.

So far, we have found that an entity characterized by absolute GPs must be a *self-maintaining* and *reproducing* kind of being, having defined these concepts in a very abstract way. Having seen this, I think it would not be wrong to call such entities *living things.* Life, in this bare logico-metaphysical sense, is just existing in a fashion characterizable with absolute GPs. This concept of life is in fact the same as the concept of life that Michael Thompson introduced in "The Representation of Life" (1995) which has since been taken up by neo-Aristotelians, starting with Foot (2001).[24] Thompson defines living things as entities subject to certain forms of description, chief among them being descriptions of vital processes ("This bobcat is finding a mate") and natural-historical judgements which express our general understanding of the life form implicit in the descriptions of vital processes ("The bobcat breeds in spring"). Natural-historical judgements are generic propositions and now we are in a position to see that they are *absolute* generic propositions, to be specific. This is due to the fact that Thompson takes living things to be characterizable by GPs not just with respect to their internal

[24] A newer version of Thompson's paper is presented in *Life and Action* (2008).

properties but also with respect to their environment. It is not just characteristic of bobcats to have certain physical features and behaviors, but also to be in an environment with certain characteristics (sources of nutrition, shelter, other bobcats etc.). The normal environment for the instantiation of the bobcat GPs like "Bobcats breed in spring" is itself the subject of characterization by bobcat GPs like "Bobcats breed in spring, in a well-defined home range." Now, a generic proposition whose instantiation conditions are themselves generic propositions is what an absolute GP is, so these GPs about bobcats express absolute standards of normality, rather than merely relative standards.[25]

Now, apart from showing us that we're on track towards reconstructing logical Footianism from within constitutivism, the concept of life does not play a crucial role in my argument. For now, I will refer to the subjects of absolute GPs using the abstract term *normally existing entities* (i.e. entities such that if conditions are normal, they exist!) Although the concept of a living thing is a lot more intuitive, and although the neo-Aristotelian concept of a living thing just introduced is also entirely abstract and free from any empirical content, it is difficult to suppress intuitions arising from our experience of living things or knowledge of biology. My argument is not meant to rely on any of these intuitions — it is meant to flow just from the logical form of absolute generic propositions. If we rely

[25] At this point, one might wonder about the status of autocatalyzing processes or chemical chain reactions, like fire or crystallization. We wouldn't call these processes living things, but don't these processes normally persist and create their own persistence conditions? And doesn't that make them the sort of thing that absolute GPs describe? The answer is no: fire and crystallization are not properly described with absolute GPs and are not alive, precisely because they are not *normally* in the conditions for the autocatalysis to occur. A living thing like a mangrove tree is not only normally near water, but grows in soil that is normally near water, and is normally in geographical conditions where such soil is generated and persists, and is normally in a situation in which those geographical conditions obtain and so on. But while fire is sometimes in conditions for the fire to grow and spread, there is no sense in which fire is normally in those conditions. So the environment selection and maintenance attributed to living things with absolute GPs does not apply to chain reactions, even though their accidental influence over their environment superficially resembles life.

on the idea of a living thing and the many things we may associate with it then it will not be clear that absolute GPs can ground normativity by themselves.

So, we've managed to determine the difference between relative GPs, which are definitely non-normative, and the sort of GPs to which neo-Aristotelians appeal. The task remains to show that these absolute GPs are normative. It is not just neo-Aristotelianism that depends on the completion of this task, but also the generic account of constitutive norms. If the task is successfully completed, then constitutivism will be vindicated in its account of broad normativity, while also sharing this account with neo-Aristotelianism.

4.4 The normativity of absolute generic propositions

So far we have considered two kinds of entities — entities characterized only with relative GPs, and entities characterized with absolute GPs. With that said, I now intend to further examine the concept of a normally existing entity, in order to argue that the absolute standards of normality that apply to normally existing entities must also count as *standards of goodness*, standards by which the individuals might be evaluated. That is to say, they are *norms*, in the broad sense.

The fact that the conditions obtain for the continued existence of a normally existing entity cannot be an accident, attributable to an external cause. Instead, members of the kind must characteristically produce those conditions. This is what it means for something to be self-maintaining: to have its continued existence as a feature of its own nature, attributed with absolute GPs. Now, maybe such an entity could be entirely featureless, apart from doing one thing: placing itself in its conditions of continued existence. However, in a material universe, it would have to do that in a particular way and undergo a process of acting on an environment containing more than one item. For this reason, it would have to have some degree of ontological complexity — it would need more than one part, so

as to be able to perform the distinct acts of initiating the environment managing process, carrying it out, checking when it was complete, and starting it again as necessary. Each of these parts would have to normally perform these acts, and do so in conjunction with the other parts. In summary, a normally existing entity needs to have a normal *way* of existing, separated into different roles played by its different characteristic parts. Now, if this kind of entity has parts that serve to play different roles in keeping it in being, we may introduce an a priori concept of *purpose, function, or end* that applies to it. Each part will have a function, the performance of which is conditionally necessary for the continued existence of the entity.

Given that it has functional parts, it will become part of the entity's task of self-maintenance to maintain these parts, each of which will require different conditions of continued existence to be met. For this reason, the parts of a normally existing being will have to serve functions of maintaining each other as part of their function of maintaining the whole. Each part will be simultaneously operating for the sake of maintaining the operation of the others and the whole being. Now, it seems perfectly legitimate to call a being made up of functional parts that form a self-maintaining system an *organized being*. (Indeed, this corresponds with Kant's definition).[26] The result of these considerations is that normally existing beings are not only alive, in an abstract sense, but they must also be *organisms.* We may also call them *teleological systems*, to emphasize the fact that all the various functions of their parts are coordinated with each other with purpose of bringing about their mutual success.

Thompson and Foot have already observed that GPs about living things can be joined by teleological connectives like "in order to" (Foot 2001, 30-32, Thompson 2008, 77-79.) For instance, "Oaks shed their leaves in fall in order to avoid frost damage" connects "Oaks shed their leaves in

[26] See Immanuel Kant, *Critique of Judgement* §65. "In such a product of nature, every part not only exists *by means of* the other parts, but is thought as existing *for the sake of* the others and the whole— that is as an (organic) instrument. … Only a product of such a kind can be called a *natural purpose*, and this because it is an *organized* and *self-organizing being*."

Fall" with "Oaks avoid frost damage". Both Foot and Thompson see this is as one of the distinguishing marks of the GPs they are interested in. We can now explain why they have this feature. We noted above that the form of absolute GPs allows them to enter into explanatory connections with "because" that entail the truth of both connected GPs. This is because the normal conditions that bring about any feature of the kind are themselves normal for the kind. So we can say that in general, the feature is present because of those conditions. For instance, we can connect "Hormonal changes occur in bobcats in spring" with "Bobcats breed in spring" in this way: "Bobcats breed in spring because hormonal changes occur in them at that time." However, we now have identified normally existing beings as organized beings, in which the normal persistence and operation of a feature serves to promote the normal persistence and operation of the others. So, for every normal condition like hormonal changes that produces a normal feature like breeding in spring, we can now see that the condition is normally *necessary* for the production of that feature. So wherever we can assert a dependence between two GPs with "because", we can assert a dependence in the other direction as well using a teleological connective like "in order to". For example, if we take, "Bobcats breed in spring because hormonal changes occur in them at that time" as an absolute GP, we can infer the converse relation: "Hormonal changes occur in Bobcats in spring in order to induce individuals to breed." In general, we can infer from the fact that trait A is normally caused by condition B, that it is the characteristic role of B to cause A as part of the system, which is what it means to say that condition B obtains in order to produce A. To give another example,

> Kudzu reaches sunlight because it climbs surfaces.

If this is absolutely normal for kudzu, then it follows:

> Kudzu climbs surfaces in order to reach sunlight.

Consider also:

> In humans, nutritional substances get conveyed to the parts of the body that need them because they are broken down into molecules.

If this is absolutely normal for humans, then it follows:

> In humans, nutritional substances are broken down into molecules in order to convey them to the parts of the body that need them.

In this way, the possibility of teleological connection is inherent in the form of absolute generic propositions.

To continue, we have seen that normally existing things depend on the normal operation of their parts for maintaining their normal way of existing, which in turn is nothing more than the normal operation of the parts as a teleological system. We have seen that the use of absolute GPs is what requires and enables the description of an entity as an organism. An interesting consequence follows from the interdependence of the parts of an organism. For inanimate things characterized with relative GPs, the explanation of a property could either end with the thing itself, thus making it a normal characteristic of the kind, or end with external causes, making it an abnormal exception to the normal characteristics. In the organized being, where we can have a teleological chain of organs and operations that depend on each other, we can have a chain of explanations that begins with an abnormality in one part which is traced to the lack of a normal property in another part, and perhaps only then to the absence of the normal external conditions for that part's operation. All abnormality is still ultimately to be considered an accident, coming about from causes external to the organism,

but the organism's articulation as a teleological system allows abnormalities first to be given a proximate explanation in terms of other internal abnormalities and thus attributed in a limited fashion to the organism itself. Thus we can distinguish between internal accidents and external accidents — in the former, the proximate explanation of an abnormality will be internal, and in the latter, we immediately look to external circumstances to explain the absence of normal operating conditions. The distinction between these two cases allows us to define for ourselves, in an a priori fashion, concepts of *internal defect* and *environmental privation*.

We may now rely on the notion of defect to describe abnormality in the parts and operations of an organism. If a part is not making its normal contribution to the self-maintenance of the system of parts, and this is not immediately due to abnormal conditions external to the organism, we may call that part defective. This performance of its normal role may of course be a matter of degree — a part could be more or less defective at various times. If it exhibits no defect, we may describe it as *excellent* — an excellent specimen of the sort of organ it is, doing well at what it is meant to do. Taking this contrast of excellence and defect as our starting point, we may go on to introduce all kinds of evaluative and directive language to describe the parts and operations of living things, with respect to how normal they are in their contribution to the maintenance of the teleological system. For example:

> Blood vessels work *better* when they are not clogged with fatty deposits.
>
> Astigmatism is a frequent cause of *poor* eyesight.
>
> Cells *ought* to stop dividing when the body no longer needs more.

We can also use this language to evaluate the environment of an organism for sufficiency or privation, e.g. "The soil for this species of succulent *ought* to be sandy and have *good* drainage."

I submit that the introduction of this broadly normative language is perfectly appropriate, in accordance with ordinary usage, and is not being used in a special technical sense. This is the language we use to describe the parts and operations of living things with respect to how well they perform their normal function, and their environment, with respect to how well it enables these normal functions. However, we can now explain why this language is appropriate for describing standards of normality that apply to living things, but not for describing standards of normality that apply to inanimate substances. This is because standards of normality for living things are also evaluative standards, and this in turn is because they are described with absolute GPs. The use of absolute GPs implies that an entity is a teleological system. In a teleological system, each part has a function consisting of what it normally does to produce and maintain the other parts. And a part's function, which is normally necessary for the existence of the system and ultimately its own existence, represents a standard for it to live up to that is not merely a standard of normality — it is something that it can do well or poorly. The proper use of broadly normative language, or at least one proper use, is to mark this transition between relative standards of normality and the absolute standards on which the normal existence of an organized being depends. Thus, we can assert, quite literally, that systematic normality for a normally existing being is *the good* for that being.

4.5 Consequences for metaethics

We now have an account of how some generic propositions are normative on account of their logical form. This has finally made good on the promise of chapters 2 and 3, of finding an account of constitutive norms that can meet the constitutivity and broad normativity conditions. We can now say that the generic analysis of constitutive norms provides such an account. We've already seen in the Chapter 3 that the generic account meets the constitutivity condition and avoids the violability

challenge to broad normativity, and we have just established in full the broad normativity of absolute generic propositions. This section will examine several consequences for any metaethics founded on this basis and for constitutivism in particular.

To begin with, we can now better justify logical Footianism on behalf of the neo-Aristotelians. We have just articulated a rationale for the basic claims of this approach. Recall that logical Footianism holds that generic propositions about living things establish a norm with respect to which the parts, operations and environment of members of the species may be evaluated. This is a standard of what Foot calls 'natural goodness'. Foot does note that some generic propositions ('Aristotelian categoricals' in her terminology) are 'teleological' and that these are the normative ones. But she does not give an account of why some generics are teleological and how this grounds normativity without importing norms from elsewhere. The account given in this chapter fills in that story.

Having filled in this missing argument for neo-Aristotelianism, we also have reason to think that any attempt to derive normativity from generics will have to do it by way of logical Footianism. This includes our account of constitutive norms in terms of generics. This means constitutivism will also be committed to a controversial feature of logical Footianism that the argument of this paper has shown to be necessary, namely, its *life form relativity*. This account of normativity makes norms life form relative, as there can be no norms outside the facts about what is normal for a particular life form. Absolute GPs possess normativity in virtue of their integration into a teleological system. This means there cannot be any "free-standing" norms outside of such a system and for any particular norm, it must be possible to specify the teleological system of which it is a part. To put it another way, normative language cannot properly be applied to a feature of a kind that does not play a normal role in the self-maintenance of that kind. Such a feature can be described with relative GPs at best, which we have seen are not normative. So, a relevant self-maintaining kind must be able to be specified if there is to be normativity and a self-maintaining kind is a life form, as we have seen.

Let us call a constitutivism founded on natural normativity via the generic account of constitutive norms *NN-Constitutivism.* Something that NN-Constitutivism will have to face is the various objections leveled at neo-Aristotelianism. For instance, William Fitzpatrick (2000) objects that closely associating biological functions with what is good, as neo-Aristotelianism does, is scientifically and intuitively untenable. Constitutivism will now be open to this charge too. On the other side of the equation, constitutivism will also stand to benefit from arguments made in favor of neo-Aristotelianism. For instance, Micah Lott provides what I take to be a sufficient rebuttal of Fitzpatrick (Lott 2012a).

NN-Constitutivism is so far limited only to the biological realm. It does not yet give us a constitutivist account of the norms of morality or rationality, which is our ultimate target. In other words, we have shown how constitutive norms can be normative in the broad sense of providing some evaluative standards and directive norms, but not in the narrow sense of providing reasons, or being authoritative for an agent. This is one of the main goals of constitutivism: to provide an account of rational or moral norms. We are also missing an account of constitutive norms of artifacts and social practices. In Chapter 5, I will show how the account of constitutive norms given in this chapter can be developed into a properly practical constitutivism which accounts for norms in the narrow sense. We will also be able to explore the consequences of life form relativity for norms of rational agency. As for artifacts and practices, since these are things that rational agents make or do an account of their norms will once again have to postponed as something downstream of an account of rational norms in general.

In conclusion, this chapter has presented a crucial step in an argument for either constitutivism or neo-Aristotelianism: the derivation of broad normativity from generic propositions. While the task of establishing their strict normativity is a separate step, it depends on the success of this one, and each of these schools of thought will now be able to avail themselves of the innovations of the other and proceed in tandem.

5.0 The Rational Authority of Constitutive Norms

In this chapter, I extend my account of constitutive normativity into the realm of rationality and ethics. Before we begin, let us take stock of the argument so far. In chapter 2, I introduced three conditions which an account of constitutive norms must meet if constitutivism is to be successful as an account of normativity. The first two apply to normativity in any domain, and the third only to practical norms for rational agents. These were the constitutivity, broad normativity, and strict normativity conditions. Constitutivity requires that a set of putative norms be grounded in or identical with facts about a kind. Broad normativity requires that the norms do not merely describe but establish a standard of evaluation. Strict normativity requires that the norms have authority over rational agents. In chapters 2 and 3, I argued that the only hope for reconciling the demands of constitutivity and broad normativity was an analysis of constitutive norms as *generic propositions* (GPs) and established that such an account would meet the constitutivity condition. In chapter 4, I provided a full account of constitutive norms in terms of GPs and showed that it also meets the broad normativity condition. I argued that broad normativity could be possessed only by "absolute" GPs, whose normal conditions of instantiation are themselves expressed as GPs. I went on to argue that norms of this sort form parts of a teleological system and as a result the fundamental kind of constitutive norms are those that characterize the life cycle, traits and environment of a *life form*. If there are to be other constitutive norms, such as rational, social, or artifactual norms, these must be somehow derived from the natural norms pertaining to a species of living thing. This placed constitutivism on common ground with neo-Aristotelian ethical naturalism, with respect to its most fundamental commitments. As a result, we can identify constitutive normativity on my account with neo-Aristotelian natural normativity.

With the broad normativity condition out of the way, the purpose of this chapter is to show how a constitutivism that identifies constitutive norms with neo-Aristotelian natural norms can meet the strict normativity condition. For brevity, I call such a view *NN-Constitutivism* ("NN" for "Natural Normativity"). The task is to show how NN-constitutive norms can be binding for rational agents and what kind of metaethics might be founded on this basis. This is necessary to fulfill the aim of constitutivism and neo-Aristotelianism alike to account for the authority of rational or moral norms.

Now, there are many existing strategies for establishing strict normativity both amongst constitutivists and neo-Aristotelians and I do not intend to add any new arguments to this inventory.[27] Instead, I will proceed by reviewing the main challenges to strict normativity for both constitutivism and neo-Aristotelianism, and identify the best line of response amongst constitutivists and neo-Aristotelians for the purposes of NN-Constitutivism. This will be the main task of sections 5.1 and 5.2.

In section 5.1, I will argue that the main challenge that each side faces is the same. On the Kantian side it appears as David Enoch's "Shmagency" objection and on the neo-Aristotelian side it appears as John McDowell's rational wolf example and the threat of "biologism", but all present the same challenge to the rational authority of the constitutive norms.

An adequate answer to this challenge can also fruitfully draw on points made by both schools of thought. In this way, the combination of both into NN-constitutivism works out to the position's advantage. In section 5.2 I will show that the best approach shared by many of the Kantian and neo-Aristotelian treatments of the problem is an appeal to practical reason and our knowledge of its norms "from within". Following the neo-Aristotelian Micah Lott (2014), I call this the *practical reason response*. Our knowledge of these norms is described in various ways such as "internal" or

[27] A note on terminology: it is worth noting that these strategies are sometimes presented as defences of *constructivism* rather than constitutivism, as in Street (2010). Furthermore, in some cases, constitutivism is used to refer exclusively to answers to this problem of strict normativity, as in Michael Smith (2015).

"spontaneous", "autonomous" or "self-conscious", and this property underwrites the authority of the norms known for the rational agents who know them. However, this approach also faces several further challenges which I will devote the rest of the chapter to addressing.

In section 5.3, I observe that NN-Constitutivism must be committed to a certain version of the practical reason response on which the authority of rational norms do not depend on their being derivable from an analysis of the concept of a rational agent or their applicability to all possible rational agents. Instead, this authority derives from the norms characterizing a certain species to which the agent belongs, along with their practical knowledge of these norms. However, some constitutivists hold that it is necessary and sufficient for strict normativity that the norms be analytically derivable from the mere concept of a rational agent, regardless of species. This view is incompatible with NN-Constitutivism and I will appeal to Doug Lavin's argument against it. I will further defend this argument from three challenges presented by Lott, Matthias Haase and Christine Korsgaard.

In section 5.4, I reconsider the objection that any facts about a life form are empirical and so any identification of rational norms with facts about a life form would give the empirical study of life forms (such as biology) an undeserved power to establish or refute those norms. This is the threat of "biologism" mentioned above. Michael Thompson has provided a decent response to this threat — however, a more sophisticated version of the biologism objection is possible that will require more work to rebut.

Having providing these defenses of NN-Constitutivism and its suitability for accounting for the authority of rational or moral norms, I will conclude with some reflections on one of its interesting features. This feature is that despite its Aristotelian foundations, the account might yet be used in support of a broadly Humean or Kantian metaethics.

5.1 Challenges to strict normativity: shmagency, the rational wolf and biologism

The best-known challenge to the rational authority of constitutive norms comes from David Enoch's (2006) paper "Agency, Shmagency: Why Normativity Won't Come From What is Constitutive of Action". In this paper, Enoch takes to task the theories of Korsgaard, Velleman and Rosati. These theories posit constitutive norms of action or agency as the root of rational or moral norms. For instance, on Korsgaard's theory, if one habitually fails to follow the categorical imperative, one fails to exhibit the autonomy constitutive of agency and ceases to count as an agent. Enoch's objection is that while the constitutive norms of agency might be *metaphysically* inescapable, they do not seem to be *normatively* inescapable. Even if one cannot avoid being an agent and exhibiting those properties that are constitutive of agents to some degree, this alone does not entail that I rationally ought to exhibit those properties to the highest degree. In my terminology, these norms might be broadly normative and provide a standard of evaluation for being a good or full-blooded agent, but they are not clearly normative in the strict sense. Someone might yet ask, "Why ought I strive to be an agent?" — this seems to be an open question. If they must comply with the standards of agency in order not to be downgraded to a state less than agency — call this "shmagency" — it's not clear why they ought not to prefer this lesser status, especially if the conditions of agency are very exacting. Or, if agency to some degree really is inescapable, it's not clear why they ought not to comply with its standards as reluctantly and half-heartedly as possible.

John McDowell (1998) poses a parallel problem for neo-Aristotelianism in his much discussed "Two Sorts of Naturalism". In this paper, McDowell responds to Philippa Foot's early work on natural norms. His argument amounts to a warning against neglecting to establish the strict normativity of natural norms alongside their broad normativity. He imagines a wolf that has somehow acquired the power of rationality and observes that such a creature would not be thereby rationally bound to pursue every behavior characteristic of its species such as cooperative hunting, even though this

would fulfill its natural norms. By dint of his rationality, he would be able to call the many normal behaviors of his species into question and consider for each whether there was a *reason* to engage in it. That is not to say that there would be no reason for him to engage in these behaviors after full consideration, but simply that they would not automatically have rational authority over him, just as the biological norms of human digestion do not automatically have rational authority over us.

Another version of this worry has arisen since the full presentation of Foot's views in *Natural Goodness*, expressed by Copp and Sobel (2004), among others. In the light of studies in evolutionary biology and anthropology that may paint an unpleasant picture of human nature, one might ask why we ought to care so much about instantiating that nature. Such things are just what a rational being might choose to rise above — it might even be rationally required of us to do so. And yet on the natural normativity view, it seems like any empirical findings about typical human behavior will immediately become normative for us, effectively granting biology authority over ethics. This would seem to violate the "autonomy" of ethics, and put us at risk of sanctioning all sorts of odious behaviors to which we might have a natural predisposition. Michael Thompson (2004) calls this the threat of 'biologism'. The thought is much the same as McDowell's, which is that no mere empirical fact should have rational authority over us unless it receives an independent endorsement from reason.

In all these cases, the rational authority of our constitution as agents or human beings is called into question. A good summary of the common problem can be found already formulated in Korsgaard, ten years prior to Enoch's paper. In *Sources of Normativity*, she observes that humans are self-conscious and this "sets us a problem no other animal has" (Korsgaard 1996, 93-94). She continues: "It is the problem of the normative. For our capacity to turn our attention on to our own mental activities is also a capacity to distance ourselves from them, and to call them into question […] *"*Is this perception/desire really a *reason* to act?" The reflective mind cannot settle for perception and desire, not just as such. It needs a *reason."* To this we might add that the reflective mind also cannot

settle for facts about agency or the human life form "just as such"; these are not automatically reasons either.

Korsgaard notes that this problem creates the possibility of a kind of moral skepticism, which she characterizes as the view that reflection on our reasons will never terminate and nothing can be verified as truly giving us reasons via reflection alone. The concerns raised by Enoch, McDowell and others are more localized versions of this skepticism, suggesting that mere facts about agency or human life will not survive the kind of rational reflection that could vindicate them as reasons. However, whatever the scope of the skepticism, the solution, thinks Korsgaard, must come from within reflection itself: "If the problem springs from reflection then the solution must do so as well. If the problem is that our perceptions and desires might not withstand reflective scrutiny, then the solution is that they might." Thus, in establishing the strict normativity of any considerations, and in the case of constitutive norms as well, we must show how it is possible to reflectively endorse those considerations as authoritative for reason. Reflection must be able to come to an end.

Now, because of the unlimited power of rational reflection to call the authority of desires, beliefs, facts and putative norms alike into question, it may seem that coming to such an end is impossible. And this will be the case unless there is some standard *internal* to rational reflection by which it judges all of these things. Korsgaard makes this point in the Kantian way, observing that in order to obligate us to do anything reason must provide its own internal principle or law—in other words, it must be *autonomous*. For Kant, this internal principle is his categorical imperative, but we need not take on this commitment at this juncture. Kant's key point about strict normativity, which many constitutivists and Aristotelians alike have recognized in their own way, is that it is only possible if there are norms that are internal to reason's reflective activity.

What it is for a norm to be internal to reason's reflective activity is not immediately clear, of course, and there have been various attempts to account for it, to which we proceed in the next section. Before that, however, I want to note that various philosophers from both schools have tried

to curtail the advance of skepticism at an earlier point by denying the substance or coherence of the "why should I care" question, citing the *inescapability* of agency, or humanity. It would seem to be a moot point just how the norms of agency get a grip on us if we cannot avoid them anyhow. Enoch's reply to this, already mentioned above, is that even if agency is inescapable, one might nevertheless carry on with it reluctantly and without really endorsing the norms that one has no choice but to comply with. Now, Luca Ferrero (2009) argues that in the case of agency, this kind of "alienated participation" that Enoch presents is in fact impossible, since there is no standpoint external to agency from which the agent can make a judgement of the norms of agency and form their attitude of reluctance. But this does not seem to improve matters, since it only seems to diminish the autonomy of reason. If the norms of agency are such that we are forced to comply with them without the possibility of rational reflection, their rational authority seems to come into question. What really needs to be shown to establish the rational authority of practical norms is that we in fact reflectively endorse them.[28] Any inescapability is relevant only as a consequence of this. That is, it would be a practical contradiction for us to effect indifference to principles of action that we simultaneously endorse and in that sense we cannot escape from them. The key point, however, is that we cannot escape from them because we have committed ourselves to them.

5.2 A joint solution

The challenge we have just reviewed has been taken very seriously by both constitutivists and neo-Aristotelians and each have provided a range of responses. However, as I have just explained,

[28] John Hacker-Wright (2012) makes such a point, saying that the inescapability of an activity does not imply that the norms of that activity have any claim over us. Instead, we must have practical knowledge of agency as a normal part of our nature, not just something we cannot escape from.

the best way of framing the problem is as the task of explaining how constitutive norms can be vindicated from within practical reason, whatever that comes to. Thus, I will focus on those responses from both schools that aim at addressing this question. On the whole, they are complementary and can be jointly put to work in the defense of NN-Constitutivism.

5.2.1 The constitutivist-constructivist approach

One solution to Korsgaard's problem of how certain norms can be internal to the activity of reason can be found in her own work. This is the claim that the authoritative principles are those which constitute the reflective agent as such. To have any rational authority for us, constitutive norms must not only pertain to our *constitution* as agents, but also to our *self-constitution* (Korsgaard 2019). They must not be merely norms that make us agents (this much would only establish what I call their broad normativity) but they must also be norms we self-consciously follow and thereby *make ourselves* into agents. Thus for Korsgaard the constitutive norms of agency are "internal" to reason because they are the standard that we ourselves choose to judge our maxims by when we come to make practical decisions. Enoch's would-be shmagent thus already endorses the principles of agency they are considering rejecting.

This recalls Kant's view that reason is subject to categorical obligations insofar as it is autonomous or practical "of itself" — which is to say that reason itself issues the imperatives that it is to follow and does not take its cue from any contingent desire or empirical circumstance.[29] Now,

[29] For one of Kant's clearest accounts of the distinction between a reason that is practical of itself and a reason that is not, see *Religion within the Limits of Mere Reason* (Akademie edition, 6:26n). An excerpt: "For from the fact that a being has reason does not at all follow that, simply by virtue of representing its maxims as

why are we entitled to say that reason is in fact practical-of-itself and brings itself under such norms? Or, to direct the question at Korsgaard's view, why are we entitled to the claim that these norms of agency, characterized in such-and-such a way, are those that are internal to the operation of reason? For Kant, it is because we are in fact aware of these rational obligations when we come to act.

> We can become aware of practical laws just as we are aware of pure theoretical principles, by attending to the necessity with which reason prescribes them to us and to the setting aside of all empirical conditions to which reason directs us.[30]

This basic consciousness of reason's own norms is what Kant calls a "fact of reason":

> Consciousness of this fundamental law may be called a fact of reason because one cannot reason it out from antecedent data of reason, for example, from consciousness of freedom (since this is not antecedently given to us) and because it instead forces itself upon us as a synthetic *a priori* proposition that is not based any intuition, either pure or empirical.[31]

A similar thought finds expression in the work of Sharon Street, even while she only takes us to be bound by Humean norms of instrumental rationality. Street (2010) asserts that as agents each of us occupies a *practical standpoint*, consisting of a certain set of values to which we are self-consciously committed. "Values" for Street is a broad category including such things as "ought" claims, evaluations in terms of good and bad, judgements that there is reason to do something and so on. In combination with the non-normative facts, these values may entail that an agent has reason to

---

suited to universal legislation, this reason contains a faculty of determining the power of choice unconditionally, and hence to be "practical" on its own (*für sich*); at least, not so far as we can see."

[30] Kant, *Critique of Practical Reason*, Akademie edition, 5:30-31.

[31] Kant, ibid.

do or value various further things. This is what Street calls entailment from within a practical standpoint. Now, a particular normative principle will be a source of reasons for an agent only if it is entailed from within their practical standpoint. And finally, a principle will only be universally authoritative for rational beings if it is entailed from within *any* practical standpoint, that is, it follows from any set of values and non-normative facts. We find ourselves committed to such principles as soon as we deliberate about what to do, no matter our starting point.

Street is skeptical about most would-be universal principles and rejects the idea that something like the Kantian categorical imperative would be entailed from within any practical standpoint. The point which brings her in line with Korsgaard and Kant, however, is that there may be certain principles that bind us because we are already first-personally committed to them from within our own standpoint. Kant's fact of reason may be understood as the fact that we possess certain values as rational beings that we are already committed to bringing to bear when we come to reflect on what to do.

Let us again consider how these views constitute a response to Enoch's challenge to the strict normativity of rational or moral norms. Street considers herself a constructivist, which she takes to be the view that "the truth of a normative claim consists in that claim's being entailed from within the practical point of view". Since I use the term "normative" in a different way, in my language this claim would translate to "the *strict normativity* of a normative claim consists in that claim's being entailed from within the practical point of view." Agreeing with Street, Karl Schafer (2015) gives a similar formulation of the central claim of metaethical constructivism as "the norms are grounded in metaethically authoritative points of view". In Kantian or Humean constructivism, the most authoritative point of view is that which is constitutive of agency, since is the point of view that is valid for every agent, since every agent adopts it. So, the answer to Enoch here is that the point of view that he is imagining retreating to so as to call into question the rational authority of the norms of agency is just the very same point of view that is constitutive of agency.

To expand on this, Enoch imagines an agent stepping back to a detached rational perspective to consider whether agency or shmagency is preferable, but this perspective is itself just the perspective of agency. Furthermore, by presenting it as the position from which normativity can validated or invalidated, even Enoch seems to recognize its authority. It is not so much that we cannot escape it, but that we already do spontaneously adopt it and recognize it as the source of authoritative practical judgements. Thus, whatever norms we endorse just in taking up that position can be used to resolve the "Why should I?" question. When directed at particular imperatives like "Console my distraught friend", a particular reason can be given, e.g. "Because that would be sympathetic thing to do". If such a principle is not to be found already within the perspective of practical reason, it may be rejected — if otherwise, it is validated. If we instead ask whether or not to be a rational agent in general, the answer is trivial, since being these things entails a commitment to all of the principles that practical reason can validate. Thus to ask these questions is to ask of one's own practical reason whether it requires what it requires, to which the answer is going to be an enthusiastic "Yes!" With such answers provided by reason itself, the "Why should I?" question is no longer open.

Now, it still remains a matter for debate what exact norms we find present in our practical consciousness. But this does not undermine this approach to the grounding of strict normativity. All that is required is that there be some "fact of reason", that is, some set of norms that we already bring to bear from within that consciousness. Or in Street's terminology, from within any practical standpoint. Enoch may yet have a valid complaint if his claim is only that there are no robust norms that arise from any practical standpoint. But if it can be shown that some do, their authority is assured.

NN-Constitutivism can adopt most of these claims wholesale. The main addition to be made would be that the norms to be grasped from within the practical standpoint would have to not just be constitutive principles of agency but facts about the *human life form*. For a better understanding of how this is possible, we should turn to the neo-Aristotelians.

5.2.2 The neo-Aristotelian approach

Micah Lott (2014) gives a name to the problem as faced by neo-Aristotelianism: the *authority-of-nature challenge*.

> Given that we are rational creatures who can 'step back' from our nature, why should we see human nature as authoritative for us? Granting that moral goodness is natural goodness for human beings, why should we care about being good human beings? (Lott 2014, 762)

He argues that the best response to this challenge is what he calls the *practical reason response*. This response involves identifying practical reason as itself a natural faculty of human beings, and thus part of human nature according to the neo-Aristotelian conception of a life form. Since NN-Constitutivism is also based on a version of this conception (reconstructed in terms of absolute generic propositions in Chapter 4) this approach ought to be applicable to NN-Constitutivism too. On the neo-Aristotelian conception of a life form, each faculty, organ or trait of a living thing is oriented towards that creature's good, which is self-maintenance and reproduction in the manner characteristic of the species. Each part will have its own characteristic contribution to make towards these ends, whether direct or indirect, and it is to be evaluated according to how well it performs this function. This is the source of the broadly normative standards that apply to each living thing and its parts.

Now, if practical reason is considered a natural power of human beings, then it too is directed at the human good and its standards must be the norms of its natural goodness in making its contribution to the human life form. For this reason, our practical rationality does not allow us to step back from the norms of human nature and evaluate them from a position of indifference. This is because some of those norms are the norms of practical reason itself which govern its proper use,

and the others are the norms of other parts of the human organism with which practical reason is integrated. Thus, McDowell's rational wolf scenario would not really be possible as it is presented. We would have to insist that if the wolf's reason is to be evaluated, it needs to be as a natural faculty of his. This means he would have to be a member of a species of wolves that possesses rationality. But if this was granted, there would be no rational grounds for the wolf to step back from his nature entirely in order to rationally evaluate it, since this would require him to step back from his rationality itself, since it's part of his nature. Instead of leaving it an open question whether he should act as other members of his species do, his reason would vindicate whatever it is rational wolves rationally do. If hunting with the pack is something that rational wolves do, then a rational wolf's reason would require it. If it is something that only wolves without reason do, then of course the rational wolf may reject it, but this would not call into question the authority of his nature, since he is not a member of a non-rational wolf species. In this way, the strict normativity of natural norms is supposed to be vindicated — at least those pertaining to practical reason itself.

Having presented all this, there nevertheless seems to be a gap in the argument. Despite Lott, it seems to me that we've only really said enough to vindicate the broad normativity of rational norms. We've certainly shown that there can be natural norms by which rational activity is to be evaluated from the third-person point of view, but are these the same as those by which reason must evaluate itself?

To illustrate the problem, it is worth a brief digression to discuss Lott's treatment of Philippa Foot. There is some debate over whether this archetypal neo-Aristotelian provides a version of the practical reason response herself. On Lott's interpretation, Foot doesn't explicitly adopt this response, being more concerned with ensuring that morality meets what she calls "Hume's practicality requirement", which is the requirement that moral considerations be able to motivate us

to action.[32] She adapts an argument from Warren Quinn to argue that the moral standards of human goodness must be something practical rationality is concerned with in its own right. Now, this is indeed a necessary condition on morality being a source of rational obligations. But it is not a sufficient condition. The requirement that moral considerations actually be able to rationally obligate us is what Lott calls "Kant's practicality requirement", and the practical reason response is the only way to meet it. By giving the requirement this title, Lott references the Kantian view discussed above that reason must be practical "of itself" and not receive its ends from outside in order to be the source of moral obligation. Rational norms must not only be able to motivate us, but they must do so spontaneously and autonomously.

Lott thinks we need to go beyond Foot and not only establish that practical reason is concerned with morality but also emphasize that practical reason is a natural faculty and the norms of rationality are features of the human life form. Only this, he thinks, can give any fact about human life rational authority. However, now that Kant's practicality requirement is on the table, we should be able to see that there is still something missing from Lott's argument. In order to properly fulfill Kant's practicality requirement, I think neo-Aristotelianism must make a similar move to the constitutivist views we saw above and show how these norms are known and followed through the self-conscious activity of reason. Lott may have shown that we can evaluate rationality by these norms, but their *authority* for reason cannot be assured unless we can say they arise and can be vindicated from within the perspective of practical reason. We saw above that this was the only way to assert that the norms of rational agency were the self-given laws of an autonomous practical reason — the same holds if our subject matter is the norms of a specifically human rational agency.

---

[32] There is some disagreement here: Hacker-Wright (2009) interprets her as having something like the practical reason response in mind, while Jennifer Frey (2018b) interprets her as providing no adequate response of this kind.

We can supplement Lott's argument by borrowing from the Kantian approach. We need only assert that not only is practical reason a natural faculty of the human species, but it is a faculty that operates through an understanding of its own standards of proper functioning. Like the Kantian fact of reason, these standards are known to us not via empirical observation of humans but in virtue of reflection on our own reason's operation whenever we act. Other neo-Aristotelians have filled in the argument in a similar way, just without the reference to Kant. Most notably, Michael Thompson provides a neo-Aristotelian version of the fact of reason when, in "Apprehending Human Form" (2004), he makes room for a form of practical knowledge of the human life form, thus ensuring it can be known from "within" rather than from empirical observation. Here he is drawing on the Anscombean idea of a practical knowledge that is the cause of what it knows, distinct from theoretical knowledge that is caused *by* what it knows. Just as I can have practical knowledge of what I am doing as part of doing it, without needing to empirically observe it, I can have a knowledge of normative generalizations about my life form just by being a bearer of that life form. In the case of action, I can know that I am walking across the street since I possess the knowledge of what it is to walk across the street and the power to put this knowledge into practice. In the case of my life form, I can know that creatures like me engage in making and keeping promises because I have learned the norms of doing so and can put them into practice myself. This helps explain how natural norms might be grasped from within the practical standpoint, as Sharon Street would put it.

John Hacker-Wright and Jennifer Frey also appeal to the idea of practical knowledge on behalf of neo-Aristotelianism.[33] Frey (2018b) argues that norms must be the objects of practical self-

[33] Not every neo-Aristotelian concurs. Hlobil and Nieswandt (2019) suggest a different approach for neo-Aristotelianism that avoids the practical reason response. Instead of identifying certain norms of human nature as internal to practical reason or "a priori" in Thompson's sense, Hlobil and Nieswandt are content to accept these norms as what Korsgaard would call "external standards". Instead, they provide an infinite regress argument and other reasons why it is acceptable that the features of human rationality that ground practical norms may actually lie outside of what practical reason itself may know, arguing against McDowell. This

consciousness in order to be *practical norms*, that is, to have any motivational or normative significance for reason. She argues that we must go to Aquinas' account of how the human good may be known "as an end" for the full story on how this is possible for natural norms of human life. Hacker-Wright (2012) thinks that neo-Aristotelianism is in fact better situated than Kantian views to respond to Enoch's challenge, since the Thompsonian account of life-form knowledge ensures that Enoch's would-be shmagent has already grasped and endorsed the standards of human practical rationality as soon as he or she is motivated to act on any reason. So the idea of a rational detachment from such standards is a non-starter — not only is it impossible but it is not needed, since in every successful act of reason, we are aware of the validity of the relevant norms. That said, from what we've seen above, I think it is within the reach of the Kantian or Humean constitutivist to make a similar response. The only real difference is whether the norms that we bring ourselves under in our practical reasoning are norms that pertain to rational beings in general or to humans in particular.[34]

It should be clear that all these attempts to establish strict normativity on all sides point in the same general direction. The solution relies crucially on an account of practical reason as a self-conscious faculty. Thus, NN-Constitutivism should follow suit and adopt a "practical reason response" of its own. In the previous chapter, I identified the source of constitutive normativity as general facts about the powers, parts and operations of particular life forms. If we assert, with the neo-

---

approach is not compatible with NN-Constitutivism, however, as it requires a separation between the norms of reason and the features of reason as a natural faculty of human beings. Instead of making these things identical as NN-Constitutivism does, Hlobil and Nieswandt make the facts about human nature the potentially unknown metaethical explanation of the norms. However this explanation is supposed to work, it is not a relation of *constitution*.

[34] I think Kantian and neo-Aristotelian views are on reasonably equal footing when it comes to establishing strict normativity, and if there is to be a push toward Aristotelianism, it is likely to come either at the level of establishing the basis of broad normativity (as in my Chapter 4) or further down the line when we come to examine just what kind of rationality and moral motivation humans are endowed with.

Aristotelians, that reason is one of these powers and the principles of reason describe its operations, we establish the broad normativity of its principles. And if we emphasize the self-consciousness of its operation, that it grasps these natural-historical facts about itself through the process of reflective endorsement by which it operates, or "from within the practical standpoint", we have established its strict normativity.

5.3 Against absolutist constitutivism

In the previous section, I suggested that NN-Constitutivism can follow constitutivists and neo-Aristotelians in adopting a defense of strict normativity based on practical reason. However, though this is not usual for neo-Aristotelians, many other theorists add to the practical reason defense of strict normativity several further commitments. This is done often without argument, on the assumption that they are a necessary part of how constitutivism establishes strict normativity. These commitments are not compatible with the neo-Aristotelian aspects of NN-Constitutivism, so it will need to be shown that these commitments are unnecessary, if NN-Constitutivism is to adopt the practical reason response.

Let us review the commitments in question. In his paper "Forms of Rational Agency" (2017) Douglas Lavin articulates them as follows, as putative requirements on a successful practical constitutivism.

> 1. **Conceptualism**: It must show certain requirements on action to be *analytically contained* in the very concept of agency, a concept that applies to all subjects capable of acting for reasons, however different they may otherwise be.

> 2. **Formalism**: It must show how to *derive* the relevant substantive requirements from a conception of agency, in such a way that a subject who was skeptical of them, but who understood himself to be an agent in the relevant sense, and was capable of appreciating what this implies, would be compelled to accept the requirements on pain of inconsistency.
>
> 3. **Absolutism**: It must show that certain normative requirements hold for *all possible* agents-capable-of-acting-for-reasons. If a kind of rational agent which is not subject to these requirements were shown to be possible, the relevant constitutivist project would be shown to be a failure. (Lavin, 2017, 174)

It is conceptualism and absolutism that concern us most here, though Lavin holds that all three views are closely logically related. Most of the theorists in what I have called the constitutivist-constructivist school subscribe to at least two. Korsgaard holds that Kantian principles can be derived from the concept of a rational agent and thus must hold for all rational agents. Street holds that the Kantian categorical imperative cannot be derived in this way, but that Humean principles of instrumental rationality can. In accordance with conceptualism and absolutism, she concludes that only Humean principles possess rational authority.

Lavin, however, argues these commitments are not a necessary part of the constitutivist account of strict normativity. Against conceptualism, he argues that constitutive rational principles may be synthetic facts about a particular kind of rational being and thus not analytically derivable from the concept of rational agency. And against absolutism, he holds a principle's failure to characterize all possible rational agents does not undermine its normative necessity for the agents it does characterize.

Conceptualism and absolutism take hold, thinks Lavin, because people assume that if there are multiple logically possible forms of agency, then that means that a particular form of agency isn't truly necessary and inescapable. It is thought if we can't insist that the principles of rational agency are necessary for all rational beings, we won't be able to answer the would-be shmagent when they

ask why they should care about this or that rational principle. "Given any richer conception of agency, Enoch's question may seem to regain its grip: why should one strive to be 'an agent' in this richer sense, if it is possible to deliberate, choose, and act without being such a thing?" (Lavin 2017, 187)

Lavin argues to the contrary that Enoch's question can still be rebutted. The normative necessity of rational principles is not undermined even if reason might take several different forms. The worry that motivates absolutism confuses the metaphysical necessity that may prevent the existence of multiple types of reason with the normative necessity of a particular form of rational agency for individual agents who bear that form. Here is how Lavin puts the point:

> The constitutivists' claim need not be:
>
> > **The Form**: There is a form of agency such that, necessarily, any subject who acts for reasons will exhibit that form of agency.
>
> It might rather be:
>
> > **My Form**: For any subject who acts for reasons, there is some one single form of agency such that, necessarily, that subject will exhibit that form of agency. (Lavin 2017, 188)

On the latter reading, the starting point of the constitutivist argument is not the claim that a certain form of agency is inescapable full stop, but that for a given subject, a certain form of agency is inescapable. If this is true, then that agent is inescapable subject to whatever normative requirements are implied by that form of agency, whether or not it is possible for there to be other kinds of agents subject to other kinds of norms.

NN-Constitutivism must agree with Lavin on all of this. This is because it regards any practical reason as a natural faculty of a particular species, and its principles of operation as being authoritative only for members of that species. Thus, it cannot be part of the practical reason defense of strict normativity that the principles of reason be known as analytic truth about all rational beings. It must be sufficient for strict normativity that they be known from the practical standpoint as

synthetic truths about beings of my species, because that is all NN-Constitutivism provides. Lavin's argument serves to establish this: even if my form of agency is not the only form possible but only the human form, it may yet be necessary and inescapable for me, a human, to exhibit that form of agency. Now, I have largely avoided putting the constitutivist point in terms of inescapability since it is not inescapability itself that is the source of strict normativity. This source is rather the knowledge of norms from within the practical standpoint, and any inescapability is a consequence of that. For this reason, we might supplement Lavin's argument with the point that not only is analyticity unnecessary for strict normativity, it is not sufficient either. Even if some principle must by definition apply to all rational agents, it still possesses no strict normativity for those agents unless they themselves know it by applying it in their practical reasoning.

5.3.1 NN-Constitutivism and analytic principles of reason

In motivating his opening the door to a synthetic rather than analytic account of reason, Lavin presents arguments against Kantian and Humean attempts to establish absolutism. In their place, he advocates for a pluralism of possible types of reason. He also provides text to suggest that both Kant and Aristotle allow for multiple logically possible types of practical reason. However, the denial of conceptualism and absolutism does not entail that the principles of reason are indeed synthetic facts about particular species or that there are multiple logically possible forms of reason. This denial only means that these things are not ruled out. Thus, NN-Constitutivism is not actually committed to a pluralism about kinds of reason. To illustrate this, I would like to consider three objections that arise in connection with this pluralism and show how NN-Constitutivism can avoid them.

Various philosophers, even those sympathetic to neo-Aristotelianism have argued that the synthetic nature of reason on the natural normativity account creates challenges that must be

addressed if the approach is to succeed. Matthias Haase is concerned that pluralism will undermine the possibility of practical knowledge. And Korsgaard provides an argument that no kind of reason but the Kantian can account for the self-constitution of the rational agent and thus a Kantian account of reason is the only one possible. However, these are only problems if we mistake the indifference of NN-Constitutivism to pluralism about reason with endorsement for such pluralism. If such pluralism turns out to be problematic, NN-Constitutivism can abandon it. Now, if we grant that NN-Constitutivism does not rule out there being only a single account of reason, this raises another concern presented by Micah Lott.

5.3.1.1 First objection: Haase

In a discussion of neo-Aristotelian metaethics, Matthias Haase (2018) objects to a pluralism of forms of reason on account of the contingency that this introduces to reason's principles. Haase thinks that if these principles are contingent features of a particular life form they are not the proper object of practical knowledge. He thinks there can be no practical knowledge of a form of reason that is just one form amongst other possible forms that we might have happened to possess. He appeals to Marx for a conception of reason as both a faculty of a living thing and also something that is necessarily the same in all creatures that possess it. He also gives the example that for Aristotle, any form of perception must be accompanied by powers of desire and motion. This can be determined a priori by the analysis of these powers. Just as we can know things about perception in general, he thinks we can know things about reason in general, even if it is a natural faculty. He does not share the Kantian preference for considering rational beings in abstraction from organic life, but nevertheless wants to insist on the universality of rational principles across all life forms. We need not get into the details of his argument that such universality is required for self-conscious

knowledge of practical reason, because even if it turns out to be sound, the contingency of the form of reason we bear is not a crucial part of the practical reason response.

Lavin introduces the idea of contingency or pluralism, i.e. that there might be Humean, Kantian or Aristotelian forms of life, as a way of motivating giving up on absolutism. However, the only contingency that *follows* from NN-Constitutivism and the commitments it has taken on from Lavin is an epistemic contingency. The only knowledge of practical reason that we gain from exercising it is an awareness of the principles the sort of reason we have. This knowledge does not inform us as to whether other alien forms of practical reason are possible or not. For all we know, it may well be that there is only one kind of structure a faculty of practical reason could take, and for every species that has a faculty of reason, its principles, structure and operations are the same. Or it could be that humans possess an Aristotelian form of reason, Martians possess a Kantian form, and Venutians possess a Humean form. Whether any of this could be the case might yet be resolved by philosophical analysis of the concept of reason, but the answer is not provided to us by our practical knowledge of our own form. Thus, if NN-Constitutivism is committed to any kind of contingency of reason, it is only an epistemic contingency from the perspective of practical knowledge. If the Kantians, Humeans or Marxists have independently discovered that only one form of reason is possible, or only one form of reason is compatible with practical knowledge, then NN-Constitutivism can accept this result. NN-Constitutivism is compatible with the principles of reason being necessary, universal, and even analytic in same way that the relation between an animal's power of perception and power of motion might be analytic. What NN-Constitutivism takes on board from Lavin is that it does not depend on any of these being true. Strict normativity is derived not from the analyticity of the norms but from the combination of characterizing a life form and being self-consciously known.

5.3.1.2 Second objection: Korsgaard

Korsgaard (2019) makes one of the arguments that I have suggested might be made, arguing that only one form of reason is possible, namely, the Kantian form. According to Korsgaard, only a Kantian account of reason can capture the distinction between free, responsible, autonomous action in which one acts on principle, and unfree, heteronomous action where one is led around by one's desire for some object. Korsgaard argues, and the NN-Constitutivist should agree, that it is only the former kind of action that can be considered a form of self-constitution. And I have argued above that the NN-Constitutivist account of strict normativity relies on practical reason operating via an understanding of its own principles of operation — i.e. being self-constitutive. Now, according to Kant, autonomous action involves acting on the form rather than the content of a principle. Korsgaard argues that neo-Aristotelian virtue ethicists have no such account of autonomous action and thus have no account of how rational agents self-constitute. And she adds that while Aristotle himself has an account of self-constitution through our responsibility for forming our own moral character, this account is a flawed one.

I have two responses to make to these claims. Firstly, setting aside Aristotle's own account of how we choose our own character, it's not clear that there is no way for neo-Aristotelian virtue ethics to account for autonomous action. Korsgaard herself points out in other work (2008b) that Aristotle possesses a conception of acting on principle rather than merely from inclination. This is exhibited in his doctrine that the virtuous person acts for the sake of the noble and in accordance with reason. With this in mind, a neo-Aristotelian could easily say that when someone acts for the sake of the noble, their "rational choice" (*prohairesis)* is chosen for its accordance with the human life form, rather than for its object alone. It is chosen because the virtuous person knows such a choice to instantiate the human good in the way characteristic of the faculty of practical reason in humans. This seems to be at the very least a good start towards articulating in neo-Aristotelian terms the idea of

acting from the form of a principle as opposed to its content, and the distinction between action in which a rational agent constitutes themselves as such and degenerate forms of action.

Secondly, and more importantly, whether or not neo-Aristotelian virtue ethics can come up with a plausible account of self-constitution does not matter to the truth of NN-Constitutivism. As I argued in my response to Haase, it is compatible with NN-Constitutivism if it turns out there is only one coherent way practical reason could exist and it is, say, roughly Kantian. Despite its neo-Aristotelian underpinnings, NN-Constitutivism is not committed to moral norms taking the form of a virtue ethics, or any other form. The same goes if Frey (2018) is right that in fact it is an Aristotelian reason that is the only kind capable of a true practical self-consciousness. The only disagreement between NN-Constitutivism and the Kantian approach is about whether the necessary universality of the norms across all rational beings is required for their strict normativity. For NN-Constitutivism it is enough that the norms characterize a given agent's life form. This remains true even if there is only one form those facts can take.

5.3.1.3 Third objection: Lott

If we grant that NN-Constitutivism is in fact compatible with the principles of reason being analytic of every possible rational being, this may create other worries. Micah Lott raises the concern that neo-Aristotelianism has not done enough to rule out the possibility of a Kantian metaethics that applies to all rational beings, not humans. Thus we might also be concerned that if it is still possible for Kantians or Humeans to establish a universal conception of reason, this might undermine the neo-

Aristotelian underpinnings of NN-Constitutivism and ground normativity in the constitutive features of a rational being in general, rather than the features of the human life form.[35]

Having argued on behalf of the practical reason response for neo-Aristotelianism, Lott poses the question: "Why suppose that what we are judging is distinctly human soundness, rather than soundness for rational beings?"

> [O]ur fundamental practical judgments might just as well be tracking what is characteristic of some more abstract class of beings, such as rational agents—even if what is characteristic of rational agency is then applied, at a second stage, to our specifically human circumstances. Thus the Kantian might agree that we can only know what is true of creatures like ourselves through the application of our fundamental practical judgments but insist that the 'creatures' in question are, fundamentally, finite rational beings. (Lott 2014, 773)

He argues that the only way to combat this criticism is to show that our reason has distinctively human constitutive characteristics. That is, we must show that the reason we possess is not just conditioned and restricted by the features of the human life form, but is constituted by them.

I propose the following alternative answer: our practical reason must be judging distinctly human soundness because there is no other possible kind of normative "soundness" on constitutivist grounds. Strict normativity requires that broad normativity be established first, and on the basis of chapters 2 through 4, it is clear that this requires a concrete life form. Normativity does not arise from rationality alone in abstraction from its integration into the teleological system of a living thing. Thus a Kantian absolutism or conceptualism that relied on the analyticity or universality of its principles as a source of normativity could not be a successful constitutivism, on my account. Lott's

[35] In the final section of Lott (2014).

proposed project of exploring what aspects of reason might be uniquely human is certainly worthwhile, but not needed to make this point.

All this having been said, NN-Constitutivism does leave open the possibility of a certain version of Kantianism. It would have be argued for that human reason takes the Kantian form, but this would be compatible with NN-Constitutivism. This is a matter for debate between NN-Constitutivists, when they come to consider what possible normative ethical theories their metaethical approach allows or rules out, but it does not threaten the overall answer to the challenge of strict normativity.

5.4 A new biologism objection

In this section, I want to reconsider the biologism objection in the light of practical reason response. This objection was concerned that neo-Aristotelian ethical naturalism was committed to a crude 'biologism', whereby the theories of biology and anthropology about human behavior would automatically become normative for human beings, violating the autonomy of morality. I grouped this together with Enoch and McDowell's concerns about strict normativity and treated the practical reason response as an adequate answer to it, for both neo-Aristotelianism and for NN-Constitutivism, which is the same view at the metaethical level. However, the practical reason response opens up the possibility for a new kind of biologism objection which needs to be dealt with.

NN-Constitutivism is committed to the view that the facts about human reason are natural-historical facts about the human life form. In his response to the threat of biologism, Thompson (2004) argues that this does not give the empirical sciences any special authority to make pronouncements about what is rational or moral for humans. We need not defer to biology or anthropology to tell us about these facts since we already possess another form of knowledge of them

through the self-conscious operation of practical reason. We need not learn from science that humans make and keep promises, for example — we already know this norm applies to us from having learned to follow it in our practical reasoning. With this, it would seem that the threat of biologism is resolved. But there is more that needs to be done.

Let us give the name "radical biologism" to the claim that biology and the other sciences have complete authority to dictate the facts of reason and morality. Thompson has successfully shown that this view does not follow from neo-Aristotelianism or NN-Constitutivism. But what about "moderate biologism", the view that while we have an internal knowledge of human reason, we might just as well learn about it through empirical observation? It seems like the NN-Constitutivist might be committed to this, insofar as they think that the facts about human reason are natural facts.

However, moderate biologism creates just as much of a problem as radical biologism. This is because as long as we have two sources of knowledge of human reason, it is possible for these sources to come into conflict and for one to refute the other. Suppose an anthropologist comes to the conclusion that humans are aggressive creatures, that use violent sacrifice or combat in order to maintain social order.[36] Given that our practical knowledge is fallible, should we defer to the anthropological theory, if it is well evidenced, and put this norm into practice? This would still seem to be a violation of the autonomy of ethics: it seems that morality should not change based on current theories of human behavior, and that our internal grasp of human morality should be given more weight somehow.

One way to ensure that our practical knowledge is privileged over the empirical would be to argue that can *only* know about rationality and morality from the inside and they cannot be the subject of empirical investigation.[37] However, Frey (2018a, 64) and Moosavi (2018) argue that to do

[36] This is roughly Walter Burkert's thesis in *Homo Necans* (1972).

[37] Lott (2012b) attempts a response of this kind.

this would be give up on natural normativity: if the facts about human reason are known only by the virtuous through their self-knowledge and are not observable features of human life, it is unclear how they are natural-historical facts alongside the facts about human perception, nutrition and so forth. And Thompson himself argues that what we might know from our own practical activity (e.g. that humans engage in practices of making and keeping promises), a Martian anthropologist might know from observing us (Thompson 2004, 72).

So we need a way of maintaining the possibility of empirical knowledge of the human life form while still maintaining some sort of priority for practical knowledge. The best way to do this is by observing that rational activity is easy to know practically and difficult to know empirically. To begin with, practical knowledge of my life form may be compared with practical knowledge of my intentional actions — my knowledge of what I am doing. In practical knowledge of action, I have a non-empirical knowledge of an empirically observable event. Now, I can be wrong about what I am doing and a third-person observer can sometimes inform me of what I'm really doing. For instance, I might think I'm traveling north, but I got turned around and I'm actually traveling south. But usually I am the authority on what I'm doing. This is because I have more information than the observer: I am aware of my intentions, for instance, that may make a difference to what my action is. I know the full description of what I'm up to. An observer may be able to inform me that I'm actually traveling south when I think I'm traveling north, but it would be unlikely for them to be able to inform me that despite what I think, going to see my parents-in-law for Thanksgiving is not what I'm doing. In the same way, the principles that creatures of my type bring to bear when we come to act are going to be much more readily available to practical knowledge than to an outside observer.

Now let us consider what has to be done to make empirical observations of rational activity and the principles behind it. Any third-person account of a given action or ongoing practice as rational activity must impute a lot more than just certain visible motions to the agents involved. Indeed, most empirical observations of human behavior do not account for what happens *as rational activity*, since

it is so difficult. To make observations from which one can infer "This is how humans rationally act" with normative force, the observations would have include, for instance, a description of what the subjects thought they were doing, the imputation of further intentions, whether the subjects thought their actions were good or rational, whether other people normally think such behavior is good or rational and so on.

Thus, a human or Martian anthropologist who set out to study the nature of human rationality or morality on the basis of observed human behavior alone would have a very difficult interpretive task ahead of them. They would need to determine which actions were intentional and which were unintentional, which were done from a stable character and which from weakness of will, which character traits were the virtues and which the vices. Personal endorsement would be a sign of an action being intentional, collective endorsement would be a sign of an action being morally good, but both measures would be quite fallible. Hypotheses would have to be floated and revised. All this would be required just to empirically discover the facts about human reason as we commonly suppose it to be through practical knowledge. If our anthropologist wanted to make revisionary discoveries about morality, the situation would be even more dire. To do so, they would have to identify broad swathes of bad human behavior as actually intentional and known to be right, or broad swathes of good behavior as actually unintentional, based in vice, or known to be wrong. On the basis of observations of outward behavior alone, this would be a monumental task, made even more difficult due to countervailing claims from our own practical reason. Such a radical revision of morality through empirical observation would have to impute deep and sweeping dysfunction to all instances of opposing judgements rendered by practical reason — even calling into question the investigators' own moral knowledge if they themselves were human. This would require yet more evidence.

The conclusion to be drawn from all of this is that most biological and anthropological findings can only have a minor bearing on morality and could only hope to revise it in concert with a moral

philosophy based on our internal knowledge of the human life form. A veritable mountain of empirical evidence would be required to overturn common moral knowledge, unless there was a complete reorientation of our moral thought from within the practical standpoint. So even a moderate biologism cannot stand: practical and empirical knowledge of human practical reason are not on an even footing. The facts about reason are natural facts and can be known empirically, but we know them much better from constituting them ourselves.

5.5 Conclusion

To review the results of this chapter, we have seen that NN-Constitutivism can draw on the resources of the two schools it arises from in order to meet the strict normativity condition. The combined practical reason response allows it to avoid the pitfalls presented by McDowell and Enoch. And there is no threat of falling into "biologism", even in the more moderate form that Thompson's arguments seem to allow for. We saw that NN-Constitutivism's account of strict normativity does not ground it in the analyticity or universality of rational principles, but it is in fact compatible with these things. Chapter 4's largely Aristotelian account of broad normativity in terms of living things turns out to be suitable to ground various accounts of reason beyond just the Aristotelian. The "life form relativism" of natural normativity noted in that chapter is maintained, though its consequences are potentially limited. NN-Constitutivism could thus be the account of normativity supporting a neo-Aristotelian virtue ethics, a neo-Kantian deontology or a neo-Humean account of instrumental rationality, even though it may turn out that only one of these gives a coherent picture of practical reason. If only one logically coherent form of reason is possible, then while the norms will still be life form relative, they will be the same in content across all life forms — in much the same way that all

legal codes might contain the same laws on some topic. Of course, in the case of the latter two schools, the way of establishing the authority of rational principles will be different from how it usually is on those accounts and some proponents may find this objectionable — the point is that the substantive moral or rational principles of these approaches *can* be preserved on the basis of NN-Constitutivism. NN-Constitutivism provides a means of establishing the normativity of rational principles for any given constitutivist metaethics and provides a framework within which the debate can occur as to which is correct. Adopting one of these accounts of reason over the others will require arguing that through our self-conscious practical knowledge we know *this* to be the form that reason takes in human beings. It will be a separate matter whether it takes that form in all other rational beings.

6.0 Conclusion: Metaethical Constraints and Possibilities

We have now seen that NN-Constitutivism can meet all three of the conditions I laid out in Chapter 2: constitutivity, broad normativity and strict normativity. This makes it a theory that can live up to the promise of constitutivism to establish norms that are internal to a kind, normative while still being descriptive, and binding on rational agents. In this final section, I would like to look ahead to how this account would structure future debates in metaethics and the theory of normativity.

NN-Constitutivism does not determine what the first order principles of reason must be, or the overall structure and operation of reason. Instead it creates a framework for debate on this topic by identifying a key question and a methodology for addressing it. The key question is: what kind of reason do we possess as human beings? Since normativity is grounded in the parts and operations of living things, and strict normativity in the operation of the faculty of reason, to learn what its principles are we must examine how it operates in our species. This restricts certain possibilities for metaethics and opens up others. For comparison, Schafer (2018) raises the question of what characterization of my nature grounds my reasons, with the options including "rational agent" on the minimal Humean account of rationality, "rational agent" on the more substantive Kantian conception, or "human being" on the Aristotelian conception. Schafer thinks that the answer is to be found by investigating the nature of rationality, but NN-Constitutivism determines an answer to this question even before we consider rationality: our reasons stem from our nature as human beings. The kind "rational agent" cannot be a source of constitutive norms since it is not a life form. Constitutive norms arise in the context of the complete teleological system of a living being and indicating that a being has a faculty of reason does not yet tell us what role this reason plays in the life of that being and thus what its norms are.

Now, it might seem that the Kantian and Humeans cede a great deal to the Aristotelians if they adopt this theory of normativity but in Chapter 5 we saw that actually much less needs to be conceded than might be supposed. Despite focusing our attention on the nature of human reason specifically, NN-Constitutivism does not tell us what form human reason takes. Thus the debate about what characterization of my nature grounds my reasons becomes the question of what kind of reason we humans possess: the key question above. NN-Constitutivism leaves open the possibility that we are Humean life forms possessed of a purely instrumental reason, Kantian life forms possessed of a universalizing reason, or Aristotelian life forms possessed of a reason characterized by an integration of reason with dispositions of character, among others. NN-Constitutivism thus does not immediately entail an Aristotelian conception of reason, nor Aristotelian first-order ethical commitments.

NN-Constitutivism does provides us with a methodology for answering the key question about human reason. Constitutivists are to determine human reason primarily through introspection upon our practical knowledge of the norms that we use to guide our actions, supplemented in some small degree by observation. Though it may also be possible to make some progress through theoretical considerations of what any possible faculty of reason may be like, this may not narrow down the possibilities enough and any discoveries must in any case be verified through practical knowledge to be authoritative.

Furthermore, it will not be enough to characterize reason or rationality in a vacuum without identifying it as part of a life form and determining its role in a teleological system. A part of teleological system is defined in terms of how it interacts with other parts of that system. Thus, in order to understand reason, we must answer two questions. Firstly, what moves reason? For instance, is it non-rational appetites? Reason's own representation of a law? A representation of an action as choiceworthy for its own sake? The answer to this question will mark a distinction between Humean, Kantian and Aristotelian constitutivisms. The second question is, how does reason move us? For instance, what kinds of desires or affects do the representations of reason produce? Does it

operate via producing a feeling of respect for the moral law, as in Kant? Or does it produce various feelings appropriate to the domain of life and the relevant virtue, as in Aristotle? In this way, the moral psychology surrounding reason will be crucial for determining its role within the human organism.

A final note on work that remains to be done: we have seen plenty of applications of the approach to biological, rational and moral norms. We were not able to show within the scope of this project how it might be applied in every domain. For instance, we did not extend the account to norms of social practices and artifacts. But we have laid the foundation for this extension. To give a rough sketch of how it might be done, what must come next is an account of practical reason — the norms of rational living activity. This must be a social form of reason, belonging to a social life form. We may then introduce as one of the joint operations of such a life form the idea of a *social practice* — a particular way of acting rationally as a group. These practices will have norms just as any other operation of a living thing has norms. They will be binding just so long as the practice is in accord with the proper operation of reason in these social creatures, just as the norms of any vital operation derive normativity from what is needed for the characteristic self-maintenance of the life form. Once practices are on the table, we might introduce the idea of a practice of *production*, where the practice shapes and uses certain things in the environment. The norms for these things—artifacts—will derive from the role they place in the practice of production and use. Of course, there are many details to be worked out here, but this is how I think the extension of NN-Constitutivism would go.

Bibliography

Artiga, M. & Martínez, M. 2016. "The Organizational Account of Function is an Etiological Account of Function." *Acta Biotheoretica* 64 (2): 105-117.

Boyle, Matthew. 2012. "Essentially Rational Animals." In *Rethinking Epistemology, Volume 2,* edited by Günter Abel and James Conant. Berlin, Boston: De Gruyter.

Burkert, Walter. 1983. *Homo Necans*: *The Anthropology of Ancient Greek Sacrificial Ritual and Myth*. Translated by Peter Bing. Berkeley; Los Angeles: University of California Press.

Copp, David and David Sobel. 2004. "Morality and Virtue: An Assessment of Recent Work in Virtue Ethics." *Ethics* 114.

Davidson, Donald. 1987. "Knowing One's Own Mind." *Proceedings and Addresses of the American Philosophical Association* 60 (3): 441-458.

Drewery, Alice. 2005. "The Logical Form of Universal Generalizations." *Australasian Journal of Philosophy* 83 (3): 373-93.

Enoch, David. 2006. "Agency, Shmagency: Why Normativity Won't Come from What Is Constitutive of Action." *The Philosophical Review* 115 (2):169-198.

———. 2011. "Shmagency Revisited." In *New Waves in Metaethics*, edited by Michael Brady. London: Palgrave Macmillan.

Ferrero, Luca. 2009. "Inescapability and the Constitution of Agency". *Oxford Studies in Metaethics* 4: 303-333.

Fitzpatrick, William. 2000. *Teleology and the Norms of Nature*. New York: Garland Publishing.

Fix, Jeremy. Forthcoming. "Two Sorts of Constitutivism." *Analytic Philosophy*.

Foot, Philippa. 2001. *Natural Goodness.* Oxford: Clarendon Press.

Frey, Jennifer A. 2018a. "Against Autonomy: Why Practical Reason Cannot Be Pure." *Manuscrito* 41 (4): 159-193.

———. 2018b. "How to Be an Ethical Naturalist." In *Philippa Foot on Goodness and Virtue*, edited by John Hacker-Wright. Cham: Palgrave Macmillan.

Geach, Peter. 1977. *The Virtues.* Cambridge: Cambridge University Press.

Haase, Matthias. 2009. "The Laws of the Thought and the Power of Thinking." *Canadian Journal of Philosophy* 39 (S1): 249-297.

———. 2018. “Practically Self-Conscious Life.” In *Philippa Foot on Goodness and Virtue*, edited by John Hacker-Wright. Cham: Palgrave Macmillan.

Hacker-Wright, John. 2009. “What is Natural about Foot’s Ethical Naturalism.” *Ratio* 42: 302-321.

———. 2012. “Ethical Naturalism and the Constitution of Agency.” *Journal of Value Inquiry,* 46 (1): 13-23.

Hlobil, Ulf and Katharina Nieswandt. 2019. “Foot without Achilles’ Heel.” *Philosophia*. https://doi.org/10.1007/s11406-019-00062-y

Huddleston, Andrew. 2016. “Normativity and the Will to Power: Challenges for a Nietzschean Constitutivism.” *The Journal of Nietzsche Studies* 47: 435-56.

Hursthouse, Rosalind. 1999. *On Virtue Ethics.* Oxford: Oxford University Press.

Kant, Immanuel. (1788) 1996. *Critique of Practical Reason.* In *Practical Philosophy*, edited by Mary Gregor. Cambridge: Cambridge University Press.

———. (1790) 1951, *Critique of Judgement*, translated by J. H. Bernard. New York: Hafner Press, 1951.

———. (1793) 1999. *Religion within the Boundaries of Mere Reason*, translated by Allen Wood and George di Giovanni. Cambridge: Cambridge University Press.

Katsafanas, Paul. 2013. *Agency and the Foundations of Ethics: Nietzschean Constitutivism*. Oxford: Oxford University Press.

Korsgaard, Christine M. 1996. *The Sources of Normativity*. Cambridge: Cambridge University Press.

———. 2008a. *The Constitution of Agency: Essays on Practical Reason and Moral Psychology*. Oxford: Oxford University Press.

———. 2008b. “From Duty and For the Sake of the Noble: Kant and Aristotle on Morally Good Action.” In *The Constitution of Agency: Essays on Practical Reason and Moral Psychology*. Oxford; New York: Oxford University Press.

———. 2009. *Self-Constitution: Agency, Identity, and Integrity*. Oxford: Oxford University Press.

———. 2019. “Constitutivism and the Virtues.” *Philosophical Explorations* 22 (2): 98-116.

Lavin, Douglas. 2004. “Practical Reason and the Possibility of Error.” *Ethics* 114 (3): 424-457.

———. 2017. “Forms of Rational Agency.” *Royal Institute of Philosophy Supplement* 80: 171-193.

Leslie, Sarah-Jane. 2007. “Generics and the Structure of the Mind.” *Philosophical Perspectives* 21 (1): 375–403.

———. 2008. “Generics: Cognition and Acquisition.” *Philosophical Review* 117 (1): 1-47.

Liebesman, David. 2011. “Simple Generics.” *Noûs* 45 (3): 409–442.

Lindeman, Kathryn. 2014. “Grounding Constitutivism.” PhD dissertation, University of Pittsburgh.

———. 2017. “Constitutivism Without Normative Thresholds.” *Journal of Ethics and Social Philosophy* 12 (3): 231-258.

Lott. Micah, 2012a. “Have Elephant Seals Refuted Aristotle? Nature, Function, and Moral Goodness.” *Journal of Moral Philosophy* 9 (3): 353-375.

———. 2012b. “Moral Virtue as Knowledge of Human Form.” *Social Theory and Practice* 38 (3): 407–431.

———. 2014. “Why Be a Good Human Being? Natural Goodness, Reason, and the Authority of Human Nature.” *Philosophia* 42: 761-777.

MacFarlane, John. 2014. *Assessment Sensitivity.* Oxford: Clarendon Press.

Maitra, Ishani. 2011. “Assertion, Norms and Games.” In *Assertion*: *New Philosophical Essays*, edited by Jessica Brown and Herman Cappelen. Oxford: Oxford University Press.

McDowell, John. 1996. “Two Sorts of Naturalism.” In *Virtues and Reasons*, edited by Rosalind Hursthouse, Gavin Lawrence and Warren Quinn, 149-79. Oxford: Clarendon Press.

Millikan, Ruth. 1989. “In Defense of Proper Functions.” *Philosophy of Science* 56 (2): 288-302.

Moosavi, Parisa. 2018. “Neo-Aristotelian Naturalism and the Evolutionary Objection: Rethinking the Relevance of Empirical Science.” In *Philippa Foot on Goodness and Virtue*, edited by John Hacker-Wright.

Moravcsik, Julius. 1994. “Essences, Powers and Generic Propositions.” In *Unity, Identity, and Explanation in Aristotle’s Metaphysics,* edited by T. Scaltsas, D. Charles, and M. L. Gill. Oxford: Clarendon Press.

Mossio, Matteo, Cristian Saborido and Alvaro Moreno. 2009. “An Organizational Account of Biological Functions.” *The British Journal for the Philosophy of Science* 60 (4): 813-841.

Nickel, Bernhard. 2016. *Between Logic and the World: An Integrated Theory of Generics.* Oxford: Oxford University Press.

Nussbaum, Martha. 1995. “Aristotle on Human Nature and the Foundations of Ethics.” In *World, Mind, and Ethics: Essays on the Ethical Philosophy of Bernard Williams*, edited by J. E. J. Altham and Ross Harrison, 86–131. Cambridge: Cambridge University Press.

Railton, Peter. 1997. **“**On the Hypothetical and Non-Hypothetical in Reasoning about Belief and Action**”** in *Ethics and Practical Reason*, edited by Garrett Cullity and Berys Gaut, 53-80. Oxford: Oxford University Press

Rödl, Sebastian. 2012. *Categories of the Temporal*, translated by Sibylle Salewski. Harvard University Press.

Scanlon, T.M. “The Unity of the Normative.” *Philosophical Studies* 154: 443-450.

Schafer, Karl. 2015. "Realism and Constructivism in Kantian Metaethics 1: Realism and Constructivism in a Kantian Context." *Philosophy Compass* 10 (10): 690-701.

———. 2018. "Constitutivism about Reasons: Autonomy and Understanding." In *The Many Moral Rationalisms*. Oxford; New York: Oxford University Press.

Searle, John R. 1970. *Speech Acts: An Essay in the Philosophy of Language*. Cambridge: Cambridge University Press.

Setiya, Kieran. 2012. *Knowing Right from Wrong.* Oxford: Oxford University Press.

Smith, Michael. 2013. "A Constitutivist Theory of Reasons: Its Promise and Parts." *Law, Ethics and Philosophy* 1: 9-30.

———. 2015. "The Magic of Constitutivism." *American Philosophical Quarterly* 52 (2): 187-200.

Street, Sharon. 2010. "What is Constructivism in Ethics?" *Philosophy Compass* 5 (5): 363-384.

Teichman, Matthew. 2015. "Characterizing Kinds." PhD dissertation, University of Chicago.

———. (2016). "The Sophisticated Kind Theory." *Inquiry*. DOI: 10.1080/0020174X.2016.1267407

Thompson, Michael. 1995, "The Representation of Life." In *Virtues and Reasons: Philippa Foot and Moral Theory*, edited by Rosalind Hursthouse, Gavin Lawrence and Warren Quinn, 247–96. Oxford: Clarendon Press.

———, 2003. "Three Degrees of Natural Goodness." Originally in Italian translation as "Tre gradi di bontà naturale", *Iride* 38: 191–197; author's preprint in English available online.

———. 2004. "Apprehending Human Form." In *Modern Moral Philosophy* (*Royal Institute of Philosophy Supplement 54*), edited by Anthony O'Hear. Cambridge: Cambridge University Press.

———. 2008. *Life and Action: Elementary Structures of Practice and Practical Thought*. Cambridge, MA: Harvard University Press, 2008.

Thomson, Judith Jarvis. 2008. *Normativity*. Chicago: Open Court.

Velleman, J. David. 1996. "The Possibility of Practical Reason." *Ethics* 106, 694-726.

———. 2000. "On the Aim of Belief." In *The Possibility of Practical Reason*, 244–281. Oxford: Oxford University Press.

———. 2009. *How We Get Along*. Cambridge: Cambridge University Press.