

In Defense of an End-Relational Account of Goodness

By

Brian Preston Coffey
B.A. (University of California, San Diego) 2004
M.A. (University of Florida) 2008

DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Philosophy

in the

OFFICE OF GRADUATE STUDIES

of the

UNIVERSITY OF CALIFORNIA, DAVIS

David Copp, Chair

Marina Oshana

Adam Sennet

Cody Gilmore

Stephen Finlay

Committee in Charge

2014

Dedication

To my loving and supportive parents
Doug and Karen Coffey

Acknowledgments

Many people have helped me complete this project, so I worry that the following list of acknowledgments—although lengthy—will nonetheless be incomplete.

I dedicate the project to my parents, who in addition to their support, ignited my interest in ethics by being excellent moral role models—although each in their own way. My mother taught me that there are rules that must not be broken, even when doing so would be easy, undetected, and beneficial. My father taught me to treat everyone with kindness and respect. Both have shaped my moral view of the world profoundly.

I have been very lucky to be a part of two very supportive and lively philosophy departments (University of Florida, and University of California, Davis). I would like to acknowledge and thank those that made those departments into such positive communities: Bert Baumgaertner, John Biro, Chris Buckels, Leslie Byrnes, Ron Claypool, Robert D’Amico, Greg Damico, Virginia Dampier, Scotty Dixon, Andreas Falke, Hana Filip, Tyrus Fisher, Kyle Fritz, Catherine Galko, Moti Gorin, Joy Hayes, John Hendrix, Michael Hunter, Marija Jankovic, Erik Johnson, Hanna Jones, Nicole Kramer, Christopher Lubbers, Kirk Ludwig, G.J. Matthey, Shawn Miller, Roberta Millstein, Bernard Molyneux, Evan Moore, Edward Perez, Ted Randles, Elizabeth Rard, Greg Ray, Sarah Roe, Kevin Savage, Ted Shear, Edie Stasulat, Holly Stillman, David Tostenson, Jon Tresan, Heidi Williams, Gene Witmer, Casey Woodling, Casey Woolwine, and Eugenio Zaldivar.

Special thanks go to the members of DaGERS (the Davis Group in Ethics and Related Subjects) for their patience reading (and re-reading!) early drafts of much of this material, and their helpful comments on how to improve it. I would especially like to thank Kyle Adams, Russell Disilvestro, Chris Healow, Tim Houk, Noel Joshi-Richard, Charles Lang, Jared Poon, Johnnie Pederson, and Kyle Swan.

Writing a dissertation can be quite a slog— and many of the people listed above helped me find ways to balance my professional and personal lives. I cannot stress enough how much their friendship has meant to me. I would also like to thank my loving and patient girlfriend, Laura Kopp, the hardworking volunteers at the UC Davis Craft Center, and the fine men and women of Sophia’s for making my life outside of my work so happy and fulfilling.

Finally, I wish to thank my committee: David Copp, Stephen Finlay, Cody Gilmore, Marina Oshana, and Adam Sennet. Their insights and comments have vastly improved the project.

David and Marina deserve special mention for being two of the best mentors a student could ask for. I owe them a great debt for all the support, attention, and encouragement they have given me these many years.

In Defense of an End-Relational Account of Goodness

Abstract

What is it exactly that we are attributing to a thing when we judge it to be good? According to the orthodox answer, at least in some cases when we judge that something is good we are attributing to it a monadic property. That is, good things are “just plain good.” I reject the orthodox view. In arguing against it, I begin with the idea that a plausible account of goodness must take seriously the intuitive claim that there is something in common between moral and non-moral goodness—something in virtue of which it makes sense to call ‘good’ both the things that are morally good and the things that are non-morally good. However, it is implausible that all judgments about the goodness of things attribute a monadic property to those things, as this does not capture what we mean when we judge something to be non-morally good. Instead, I propose and defend my own relational theory of goodness according to which goodness is a relation that holds between things and ends (or goals). That is, goodness is a relational property such that to be good is to be related in the relevant way to some goal or another. This is true in both the mundane everyday cases and in loftier cases concerned with moral goodness.

Table of Contents

Dedication	ii
Acknowledgments	iii
Abstract	iv
Introduction	1
Chapter One – The Debate over the Nature of Goodness	9
Chapter Two – My Argument that Goodness is Relational	31
Chapter Three – Why We Do Not Need Intrinsic Goodness	80
Chapter Four – Other Relational Accounts of Goodness	124
Chapter Five – My End-Relational Account of Goodness	164
Bibliography	203

Before turning to those moral and mental aspects of the matter which present the greatest difficulties, let the inquirer begin by mastering more elementary problems.

Sherlock Holmes, Study in Scarlet, Part 1, Ch. 2

Introduction

I. Preliminaries

What is it exactly that we are attributing to a thing when we judge it to be good?

This is the principal question I will be investigating in this dissertation. Since Plato, the orthodox answer to this question has been, at least in some cases, that when we judge that something is good we are attributing to it the monadic, non-relational property of goodness. That is, some good things are “just plain good.” I reject the orthodox view.¹ I argue instead that goodness is a relation; a thing is good just in case, and because, it is related in the relevant way to a goal or end. Goodness is end-relational.² This is true of both the mundane everyday cases, and of the cases concerned with moral goodness.

To clarify how my view differs from the orthodox view, let me explain some terminology by way of an illustration. Consider the coffee in the mug that is sitting on my desk. That coffee has many qualities or attributes: for instance, it is brown in color, it is

¹ There are a growing number of philosophers who have rejected the orthodox view. For instance: R.B. Perry (1926), Peter Geach (1956), Paul Ziff (1960), Georg Henrik von Wright (1963), J.L. Mackie (1977), Phillipa Foot (1985, 2001), Judith Jarvis Thomson (1992, 1993, 1997, 2008), David Copp (1995), Christine Korsgaard (2013), and Stephen Finlay (2001, 2004, 2005, 2006, 2008, 2014).

² Stephen Finlay is the first to elaborate and defend an end-relational account of goodness. See (2001, 2004, 2006, 2008, 2014). While I had developed much of my own end-relational account prior to reading the fullest statement of his view—developed in his 2014 book—I found that book to be a helpful foil for my arguments.

bitter, hot, it was made a few minutes ago, and it is currently sitting in a ceramic mug on the desk in front of me. All of these features are either *properties* or *relations* that the coffee has, or *instantiates*. So what is the difference between a property and a relation? Properties are attributes that are monadic—which simply means that there need only exist one thing to instantiate that quality. The color brown, then, is an example of a monadic property. This coffee is brown. Coffee could instantiate the property of being brown by itself, without requiring the existence of another object. Relations, on the other hand, typically are instantiated by more than one thing—a relation holds between two or more entities. Coffee cannot occupy a mug unless there is a mug to occupy, so occupying is a relation that holds between the coffee and the mug. Now I said that relations ‘typically’ relate two or more objects because there are some relations that a thing can stand in to itself. These relations are *reflexive*. An example of a reflexive relation is the relation of identity. An entity can only stand in the identity relation to itself.³ As I stated above, the account of goodness that I will be arguing for says that goodness is a relation—but on my account, it is not a reflexive relation. So for our purposes, we can mostly ignore reflexive relations, and continue speaking as if relations hold between two or more entities.⁴ I bring up reflexive relations merely to note that some relations could be instantiated by a single entity.

At first glance, there is a strong presumptive case that goodness is (at least in some instances) a monadic property. Sentences like: “Coffee is good,” appear to ascribe

³ Other examples of relations that are reflexive: Being the same shape as, being the same size as, etc. While it may sound strange to say that the coffee mug is the same shape as itself, it is nonetheless, true of the coffee cup that it is the same shape as itself.

⁴ The issue of reflexive goodness will be discussed in Chapter Five when I discuss Stephen Finlay’s end-relational account of goodness. On his view, a state of affairs can be good relative to itself. I argue otherwise.

goodness *to coffee*. This sentence appears to only be about coffee; it does not appear to be about anything else. That is, the sentence does not appear to be relating coffee to some further entity. Thus, it certainly looks like this sentence is ascribing to coffee the monadic property of goodness.

The previous argument trades on the fact that the *predicates*, or adjectives and adjectival phrases that we use to express properties and relations, can also be monadic or polyadic. For instance, the predicate ‘is brown’ requires only one noun—or noun phrase—to form a complete sentence. It has one empty *place*; it is thus a monadic predicate. On the other hand, the predicate ‘is in front of’ requires at least two nouns (or noun phrases) to form a complete sentence. It has two empty places, and is thus dyadic. It is tempting to think that monadic predicates always express monadic properties, and that polyadic predicates always express relations, but this is not quite right. Compare, for instance, the polyadic predicate ‘is in front of,’ with the monadic predicate, ‘is in front of me’. The latter predicate is monadic because it only has one empty place; but being in front of me is a relational property. The attribute of being in front of me is the kind of attribute that only needs one entity to instantiate it, and is thus, a monadic property, but it is nonetheless *relational*; it is constituted by a relation to me. Being in front of me is thus a monadic relational property.

There are two lessons to draw from this last example. First, the example shows that there can be monadic *relational* properties. For the sake of simplicity, I will continue to use the phrase ‘monadic property’ to refer only to *non-relational* monadic properties, and will add the qualifier ‘relational’ whenever I wish to speak about monadic relational properties. So when I say that the orthodox view in philosophy is that goodness is a

monadic property, I mean that it is a monadic, *non-relational* property. Second, since the predicate used to express the monadic relational property of being in front of me was a monadic predicate, it follows that in some cases, monadic predicates can express relations or relational properties. The presumptive argument for the claim that goodness is a monadic property is thus, inconclusive. Although ‘good’ is a monadic predicate, it might nonetheless express a relation or a relational property. And that is precisely what my thesis is: that ‘good’ always expresses a relation that holds between things and goals.

Having clarified my thesis, let me now state briefly how I go about supporting it. I began this introduction with a quote from Sherlock Holmes that quite succinctly and beautifully describes the method of doing philosophy that I pursue in this dissertation. Questions about the nature of goodness are difficult, philosophical questions with broad implications. The way we answer them will have implications for both meta-ethics and normative ethics. And although those are the areas where philosophers are most eager to get answers about the nature of goodness, I have opted to begin elsewhere; I have opted to begin my inquiry with the ‘more elementary’ questions. My thought was that I should start in familiar territory, and only venture further out into the unknown after having a good understanding of the lay of the land close to home. To be less metaphorical: I thought that I should begin by investigating the everyday, common occurrences of the word ‘good,’ and see what they reveal about the nature of non-moral goodness.⁵

After all, everyday we judge countless things as good in ways that are not (typically) regarded as moral. We judge the goodness of newspapers, websites, parking

⁵ G.E. Moore was perhaps the first philosopher to pursue such a strategy. Since Moore, most philosophers concerned with describing the nature of goodness have pursued this strategy. Notably: R.M. Hare (1952), Paul Ziff (1960), Georg Henrik von Wright (1963), J.L. Mackie (1977), Paul Ziff (1960), Peter Geach (1956), Judith Thomson (1992, 1993, 1997, 2001, 2003, 2008), and Stephen Finlay (2001, 2004, 2006, 2014).

spaces, restaurants, conversation topics, clothing, hairstyles, routes to work, student papers, etc, etc. Due to the pervasiveness of our non-moral goodness judgments, one would expect that we would be well situated to answer questions about the nature of non-moral goodness. So why not begin our investigation into the nature of goodness with an investigation into non-moral goodness? Only after having developed an adequate theory of non-moral goodness do I then move onto investigating the philosophically vexing questions about moral goodness. I call this the ‘Bottom-Up Strategy.’”

I believe that much more can be said about the virtues of the Bottom-Up Strategy, or about the perils of pursuing the opposite strategy (the ‘Top-Down Strategy’). I will save that discussion for another time.⁶ Ultimately, the true test of a methodology is whether or not it gets the right results. I believe that my pursuit of the Bottom-Up Strategy delivers a plausible, elegant theory of all goodness—both moral and non-moral. It is hard to argue with those results.

II. Chapter Summaries

In Chapter One, I review the contemporary discussion of the nature of goodness. I frame the discussion by focusing on the views of G.E. Moore, who defended the view that goodness is a simple, indefinable non-natural property, and Peter Geach, who argued that goodness is a relation between things and kinds. I begin by evaluating Moore’s arguments. I argue that he was unsuccessful in showing that ‘good’ is undefinable. I then present Geach’s argument, and consider some objections to it offered by R.M. Hare and Michael J. Zimmerman. Ultimately, I end up agreeing with the objectors: Geach’s

⁶ I discuss the issue very briefly in Chapter Two.

argument simply does not fully support the conclusion that all goodness is relational. I take the shortcomings of Geach's argument as instructive, however, of what more is needed from an argument to fully support the claim that all goodness is relational.

In light of the lessons learned from Geach's failed argument, in Chapter Two, I construct my own argument that goodness is relational. Briefly put, I argue that the best account of the nature of non-moral goodness is that it is a relation that holds between the things that are good and ends. I then argue that the word 'good' is univocal—and thus we should conclude that our moral uses of the word ascribe the same attribute to things that our non-moral use of the word does. Thus, all goodness is end-relational. There are three possible ways of objecting to this argument: (a) deny that non-moral goodness is end-relational, (b) argue that 'good' is ambiguous, or (c) point to something other than their end-relational nature that non-moral goodness and moral goodness have in common, in virtue of which it remains appropriate to insist that 'good' is univocal. I argue that each of these options conflict with too many of our pre-philosophical judgments about goodness. We are thus correct to conclude that all goodness is end-relational.

In Chapter Three, I look at what I believe is the most glaring and troubling (apparent) shortcoming of an end-relational account of goodness: according to such a view, there is no such thing as intrinsic goodness. After all, on an end-relational account of goodness, things are good because of how they relate to goals. They are good for the sake of goals and thus cannot be good for their own sakes. The fact that my view requires denying the existence of intrinsic goodness might make my view seem too revisionary to be true. However, I argue that the roles that intrinsic goodness has been taken to play in

our practical lives and practical reasoning can be filled even if we only posit end-relational goodness. It is thus unnecessary to posit intrinsic goodness.

In Chapter Four, I further defend my premise that non-moral goodness is end-relational. A handful of philosophers (Peter Geach, Judith Thomson, Philippa Foot, Richard Kraut, to name a few) have defended relational views of goodness where goodness is not relative to *goals*, but to other things instead. Notably, they have defended the views that goodness is relative to kinds, or relative to interests. I describe these different views and evaluate the motivations for them, and the theoretical virtues they enjoy. I argue that these theories either conflict with our pre-philosophical intuitions, or lack important theoretical virtues that an end-relational theory of goodness enjoys. I conclude that a theory that views goodness as relative to ends is thus the most plausible relational account of goodness.

Finally, having settled that an end-relational theory of goodness is the most plausible relational theory, the last task remaining is to develop and defend my preferred end-relational account of goodness in greater detail. I do this in Chapter Five. In doing so, I distinguish my view from that of Stephen Finlay. I focus on his view because it is the most developed end-relational theory that can be found in the literature. As I read him, Finlay disagrees with me about three key components of an end-relational theory of goodness: (1) What type of things can be good, (2) How we should characterize ends, and (3) What the goodness relation is like. I provide reasons to think that my view is more plausible than Finlay's in regards to all three of these points, thus showing my view to be the more plausible end-relational theory of goodness overall.

Finally, in the last section of the dissertation, I discuss the possibility of extending my end-relational account of goodness to account for moral goodness. I sketch a couple of ways in which someone taken by my account of goodness might develop a moral theory. The goal of that section is not to argue for these moral theories, but merely to show that sensible and plausible moral theories can be developed, even if we only posit end-relational goodness. Further developing and defending those theories will have to be left for future work. This last section, I hope, vindicates my pursuit of the Bottom-Up Strategy. If that section is successful, it shows that one can construct a philosophically plausible moral theory by starting from the ‘elementary’ questions and working up.

Chapter One – The Debate Over the Nature of Goodness

I. Introduction

In the last few decades, a handful of philosophers have argued against the orthodox view that goodness is a property, insisting instead that it is a relation (although, they tend to disagree about the details). These *relationalists*,⁷ for lack of a better term, noticed that some monadic predicates, like ‘in front of me,’ express relational properties. The surface structure of these predicates was not a good indicator of what kind of attribute they express. Since the surface grammar of sentences can be misleading, the relationalists insist that we need to inquire into the meaning of sentences to understand the nature of goodness. We should not conclude that goodness is a monadic property simply because ‘good’ is a monadic predicate in sentences like: ‘Coffee is good.’

The first such relationalist was Peter Geach, who in his article, “Good and Evil,” made a very interesting observation about the behavior of certain adjectives.⁸ Geach observed that sentences containing the adjective ‘good’ had very different logical properties than the sentences containing adjectives that clearly express monadic properties. Thus, he concluded, ‘good’ does not express a monadic property. For better or for worse this argument has been used as a starting point for many of the relationalists who have followed Geach.

⁷ For instance: R.B. Perry (1926), Peter Geach (1956), Paul Ziff (1960), J.L. Mackie (1977), David Copp (1995), Philippa Foot (2001), Judith Jarvis Thomson (2008), Richard Kraut (2007, 2011), Christine Korsgaard (2013), and Stephen Finlay (2014). In Chapter Four I will discuss to what extent it is proper to call some of these philosophers ‘relationalists’.

⁸ Geach (1956)

In this chapter, I will evaluate Geach's argument that all goodness is relational. Before looking at the details of his argument, however, I will first discuss the views of G.E. Moore—a defender of the orthodox view of goodness, and one of Geach's primary targets in "Good and Evil." In fact, in that article Geach was less concerned with developing a positive account of goodness as he was with arguing against the view of goodness offered by G.E. Moore.⁹ This is understandable; Moore's arguments revitalized an interest in the question of the nature of goodness. Moore's conclusion, however, was that not much could be said about the nature of goodness—goodness is unanalyzable. If Moore is correct, then the project of this dissertation is doomed from the start. After all, I wish to analyze goodness as a relation that holds between things and ends. If goodness is unanalyzable, then my theory of goodness is mistaken. Thus, my first task in this dissertation is to argue that Moore's argument that goodness is unanalyzable is unsound.

Once I have dealt with Moore's argument, I will turn my attention to Geach's. After explaining the argument in more detail, I will consider some of the criticisms that have been offered by R.M. Hare and Michael J. Zimmerman.¹⁰ They argue that Geach's observation about the logical properties of sentences containing 'good' is limited to a much smaller class of sentences than Geach imagined. His argument thus only supports the claim that goodness is relational—in some cases. Ultimately, I will side with Geach's critics. Geach's argument does not adequately support his thesis that all goodness is relational. My hope is that by identifying what goes wrong with Geach's argument, I

⁹ Geach also spent much of that article arguing against the theory of goodness offered by R.M. Hare. I will discuss Hare and his view of goodness in the next chapter, when I discuss non-cognitivist views of goodness.

¹⁰ Hare (1957), Zimmerman (2001)

stand a better chance of constructing my own argument that is not vulnerable to similar objections.

II. Moore's Account of Goodness

G.E. Moore's work, *Principia Ethica*, launched the contemporary meta-ethical discussion of goodness. Contemporary debate over the nature of goodness has largely been a response to G.E. Moore—with one side defending, extending, and developing the views he offered in *Principia Ethica*, and the other side criticizing those views and offering alternatives. While I disagree with much of what Moore and his followers have written, I will concede that there are a few points on which we agree. For instance, Moore and I agree that the question about the nature of goodness is of central importance for moral theorizing. He writes:

Ethics is undoubtedly concerned with the question what good conduct is; but, being concerned with this, it obviously does not start at the beginning, unless it is prepared to tell us what is good as well as what is conduct. For 'good conduct' is a complex notion: all conduct is not good; for some is certainly bad and some may be indifferent. And on the other hand, other things, beside conduct, may be good; and if they are so, then, 'good' denotes some property, that is common to them and conduct; and if we examine good conduct alone of all good things, then we shall be in danger of mistaking for this property, some property which is not shared by those other things: and thus we shall have made a mistake about Ethics even in this limited sense; for we shall not know what good conduct really is...
(G.E. Moore, *Principia Ethica*, 1903, p.54)

In this passage, Moore expresses two other thoughts with which I agree. First, he claims that there must be something in common between those things that are morally good and those that are non-morally good, and that the word 'good' denotes this shared property. My only wish is that he did not use the word 'property' when making this point. After all,

I claim that it is a relation—not a property—that all good things share in common. However, were we to use a neutral term like ‘feature’ or ‘attribute’ that could equally apply to properties as well as to relations, then Moore and I would agree that ‘good’ always ascribes the same attribute to things, whether we are talking about non-moral goodness or moral goodness. In the next chapter, I call this the Thesis of Commonality, and it is a crucial premise in my main argument. Second, Moore and I agree that there is ‘danger’ in beginning an inquiry into moral goodness by focusing on moral goodness alone (what I called a ‘Top-Down Strategy’ in the Introduction). Doing so would likely result in a theory of goodness that is not unified—that is, a theory that denies the Thesis of Commonality. Such a view would claim that moral goodness is one type of property or relation, while non-moral goodness is another (or others). I suspect that Moore would agree with me that there is a presumptive case in favor of the Thesis of Commonality, and thus, in favor of a unified theory of goodness. It is those who wish to argue that ‘good’ sometimes ascribes one type of attribute to things, and other times ascribes another, that owe us reasons for abandoning the presumption of the Thesis of Commonality. Despite these similarities between our views, Moore goes on to make many claims I disagree with, and consequently, offers a very different theory of goodness than the one I defend.

For example, one of Moore’s central claims in *Principia Ethica* is that ‘good’ is simple and indefinable. This appears to be a claim about the word ‘good,’ and not about goodness itself, but Moore is clear that his conclusion is not about words or lexical definitions. He writes:

The most important sense of ‘definition’ is that in which a definition states what are the parts which invariably compose a certain whole; and in this

sense ‘good’ has no definition because it is simple and has no parts. It is one of those innumerable objects of thought which are themselves incapable of definition, because they are the ultimate terms by reference to which whatever *is* capable of definition must be defined (G.E. Moore, *Principia Ethica*, 1903, p.61).

This quotation clarifies Moore’s thesis in two ways. First, it shows that Moore was not looking for ‘definitions’ of words, but of the properties and objects themselves that our words pick out. His thesis is about goodness as an *object of thought*—not as a lexical item. It is thus the *property* of goodness that he thinks is simple. Second, the quote clarifies what Moore means by ‘definition’. A definition, in Moore’s sense, involves describing the nature of an entity by describing its component parts. In summary, we should really understand Moore’s thesis as claiming that the *property* of goodness is simple, and thus, not *analyzable* (at least not in the sense of ‘analysis’ that involves describing an entity by describing its constitutive components).

According to this picture, some properties are *complex*, that is, built up out of more basic component parts. We analyze these properties by breaking them into their component parts. So for instance, being a horse can be defined, according to Moore, as being: “A hoofed quadruped of the genus *Equus*.”¹¹ Sometimes while analyzing an entity, we will discover that some of its component parts are themselves complex. However, Moore thinks we can simply analyze those components into their components. For instance, in the above analysis of what it is to be a horse, we can further analyze what it is to be a quadruped as “a being with four feet”. Eventually, Moore thinks any successful analysis will reveal the set of simple properties or entities (that cannot be broken down further into more basic entities) that constitute the complex entity that is being analyzed. Moore believed that goodness was one of these simple, unanalyzable

¹¹ Moore (1903) p.60

properties. If Moore is right about this, then my entire project is doomed from the start. After all, I believe that we can give an end-relational definition of goodness. If goodness is simple, then it certainly cannot be given an end-relational definition. I need to show that there is something wrong with Moore's arguments for this claim.

Moore's primary (and most famous) argument for the claim that goodness is simple and unanalyzable is what has come to be called the "Open Question Argument." In brief, Moore proposes a test we can use to determine if we have reached the correct analysis of a complex entity. Suppose we are attempting to analyze a kind of entity, E, and we propose the analysis A. If A were the correct analysis of E, says Moore, then the question: "I recognize that X is A, but is X really E?" when asked of a particular entity, X, will sound closed.¹² When a question is *closed* anyone who understands the question cannot fail to thereby know the answer. So for instance, "I recognize that Brett is an unmarried man, but is Brett really a bachelor?" should sound closed to us. We know that by definition bachelors are unmarried men. Since we understand this fact, there is really no way for us to wonder about whether or not Brett, the bachelor, is an unmarried man. The question is settled—and thus closed. A question is *open*, on the other hand, if a person can understand what the question is asking and can sincerely entertain more than one possible answer to it. Moore's thought was that if A were the correct analysis of E, then A would mean the same thing as the word we use to refer to E. Thus, the test

¹² This is not exactly the form of the question that Moore proposed. His Open Question Test was of the form: "Is E really A?" I believe that nothing crucial turns on my decision to formulate the Open Question Test in the manner that I have. Notice that if Moore's test question were open, it would explain why my test question is too: "Is a horse really a hoofed quadruped of the genus *Equus*?" sounds closed, thus, so would, "I recognize that Mr. Ed is a hoofed quadruped... but is Mr. Ed really a horse?"

question: “I know that X is A, but is X really E?” should sound closed. Moore concludes that if the test question sounds open however, then A is not the correct analysis of E.

Moore claimed that any possible analysis we could give for goodness would fail the Open Question Test. He writes: “whatever definition be offered, it may be always asked, with significance, of the complex so defined, whether it is itself good.”¹³ So for instance, my account of goodness claims that when we judge that X is good, we are judging that X promotes a goal G. We can apply Moore’s test to this analysis by asking: “I recognize that X promotes G, but is X really good?” This question sounds open, so Moore would conclude that my analysis of goodness is mistaken. Now of course, Moore does not go through each and every possible analysis of goodness, showing them all to fail his Open Question Test. Instead, Moore challenges us to think of an analysis that would make the test question sound closed. He holds that there is no such analysis; therefore, there must not be a *correct* analysis of the property of goodness. That is, goodness is simple.

In response, I must concede to Moore that the question: “I recognize that X promotes a goal G, but is X really good?” does indeed sound open. In order to defend my proposed analysis¹⁴ of goodness and simultaneously show that Moore is wrong to conclude goodness is unanalyzable, I need to explain how (a) we have the intuition that the test question is open (what I will be referring to as the ‘Openness Intuition’), even though (b) my analysis is still the correct one.

There is a vast literature on the Open Question Argument; it would be another book length project to address the plentitude of alleged problems (and solutions) that

¹³ *Ibid*, p.67

¹⁴ In the sense of analysis that Moore has been calling a ‘definition’

philosophers have raised with respect to this test. Instead, I will focus on one response that is open to relationalists—one that has been given by J.L. Mackie and Stephen Finlay.¹⁵ Their response involves identifying another factor that could (at least in the case of goodness) generate the Openness Intuition when we are confronted with the test question. They argue that certain questions with the same grammatical structure as the test question can sound open when they are asked of words that are context-sensitive.¹⁶ Since contextual clues are required to understand the meaning of the questions, our having the Openness Intuition can be explained, at least in these cases, pragmatically. In short, they claim that the Open Question Test will occasionally give us false negatives. This shows that Moore was wrong to conclude that ‘good’ was unanalyzable simply because it fails the Open Question Test. The Openness Intuition we have when we run the test question on a proposed analysis of goodness *could be* an indication that our proposed analysis is incorrect, or it could be that ‘good’ is context-sensitive.

Consider the relational property of tallness. Now, whatever tallness is, it is not merely a quantitative measure of height. We can measure height with a ruler—and give more or less precise quantitative judgments about a thing’s height. However, given any particular measurement of height, that height can be rightfully called both tall, and not tall (relative to different standards). This is because tallness is always relative to some class of entities, whereas height is not relative. A height of five feet is pretty tall for a kindergartner; it is not, however, very tall for a basketball player. It is extremely short for a redwood tree, and extremely tall for a blade of grass. A height of five feet is thus

¹⁵ Finlay (2014) p.21, Mackie (1977) p.60

¹⁶ Finlay calls these ‘incomplete predicates.’ His examples are ‘tall’, ‘old’, ‘fast’, ‘cold,’ and ‘eager’ (2014) pp.20-21. Geach uses the example of ‘big,’ (1956) p.33, Ziff uses the example of ‘heavy’ (1960) p.203, and Thomson uses the example of ‘famous’ (2008) p.14. In what follows, I stick to the example ‘tall’.

not just plain tall. The same could be said of the height of 500 feet, 5,000 feet, or even five million feet for that matter. A height of five million feet would be extremely tall for a building on Earth. However, we can imagine a race of giant aliens for whom a building of five million feet is quite short. Tallness, I conclude, must always be understood as relative to some comparison class.

Often, we will use sentences to express the thought that something is tall, without explicitly stating which group relative to which we are judging that thing's tallness. In those cases, our audience will have to determine, by contextual clues, which comparison class we are using. Standing in front of a rollercoaster that my friends are trying to coax me to ride, I might say: "I don't know; it looks pretty tall." It would be inappropriate for my friends to respond: "You think *that* is tall? Have you ever seen Mt. Everest? Now *that's* tall!" This response is inappropriate because my friend's claim invokes the standards of tallness for mountains (or perhaps of terrestrial objects), and it is uncharitable for my friend to think I was judging the tallness of the rollercoaster relative to those standards (surely I know that even the tallest rollercoaster is short compared to mountains!). This interpretation is uncharitable since it requires interpreting me as either an incompetent judge of tallness, or as blatantly flouting one of Grice's Conversational Maxims.¹⁷ One Gricean maxim—the Maxim of Relation—states that one's contribution to the conversation should always be relevant. Given that we were talking about rollercoasters, and more specifically, about riding this particular rollercoaster, it would have been irrelevant to our conversation for me to suddenly start comparing rollercoasters to mountains without signaling the change of topic in any way. More charitably then, one should interpret my comment about the ride's tallness to reflect a judgment about the

¹⁷ Grice (1989)

rollercoaster's tallness relative to the group: rollercoasters (or perhaps, the group: structures I would feel safe sitting atop of with little more than a seatbelt securing me). These interpretations of my comment do not require either interpreting me as an incompetent judge of tallness, or as flouting any Gricean maxims.

I hope that my rollercoaster example illustrates why I am calling our judgments about tallness context-sensitive. Since I did not explicitly state it, my audience needed to appeal to contextual clues to determine which comparison class it was that sets the standards for tallness that I was using when judging the rollercoaster as tall. In this case, there was no real mystery; we were talking about rollercoasters after all. In fact, it is often unnecessary to explicitly state which group it is relative to which we are judging an object to be tall. I do not need to say, "It's tall for a rollercoaster," I can just say, "It's tall." However, we should still understand the latter sentence to include an implicit reference to the standards of tallness for rollercoasters. No explicit reference to standards is necessary if the context of the utterance provides sufficient clues to determine what standards are being utilized.

Now consider a child, Kevin, who is 5 feet tall. That is quite tall for a kindergartener. However, the following question still sounds open: "I know that Kevin is tall for a kindergartener, but is he *tall*?" My claim is that 'tall' introduces a context-sensitive relation. Notice, that for the first occurrence of 'tall' in that question we are alerted to which group of entities sets the criteria for tallness—kindergarteners. However, the second occurrence of 'tall' lacks an explicit reference to which group it is that sets the standards for tallness. The hypothesis that Finlay and Mackie propose is that the reason we hear this question as open is due to pragmatics. We hear the question as if the second

occurrence of ‘tall’ introduces a different standard of tallness from the first occurrence. Why do we hear the sentence this way? Were we simply to interpret *both* occurrences of ‘tall’ as introducing the standard of tallness for kindergartners, then the question would clearly be closed. “I know that Kevin is tall for a kindergartener, but is he tall for a kindergartner?” is a closed question. Since this question is closed, we know that the asker already knows the answer to the question (unless she does not understand the words she is using). Why would she ask this question if she knew the answer to it already? Interpreting the sentence as closed thus requires interpreting the asker as either incompetent with respect to the words she uses in the question, or as flouting the Gricean Maxim of Relation that requires that our contributions to a conversation be relevant. To avoid interpreting our interlocutor as intentionally flouting Grice’s maxims, we consequently avoid interpreting both occurrences of ‘tall’ as introducing the standards of tallness for kindergartners. Instead, we interpret the second occurrence as introducing some further standard. The Openness Intuition we have in response to the question about Kevin is due to this imagined switch in standards.¹⁸ After all, the question, “I realize that this thing is tall relative to this set of standards, but is it tall on this other set of standards?” will always sound open. Thus, there is at least one possible confounding factor in the Open Question Test—if the term we are judging in the Open Question Test is one that is potentially sensitive to the same kind of contextual standard-shifting as exhibited by ‘tall,’ this standard shifting will make the question sound open.

It is worth noting one shortcoming of my description of the argument so far. The question about Kevin’s tallness is not one where the first part of the question contains a

¹⁸ Stephen Finlay (2014) pp.21-2 and Mackie (1977) p.60 both explain the Openness Intuition we have in response to the Open Question Test as being a result of perceived standard-switching.

putative analysis of ‘tallness’. My example is thus not perfectly analogous with the Open Question Test. My opponents might conclude that no lesson about the Open Question Argument can thus be extracted from consideration of my example.

Insofar as the Open Question Argument relies on our intuitions about which questions sound open, however, I see no problem in drawing conclusions about the Open Question Argument by considering different ways in which questions might sound open. If a presumed shift in standards of evaluation can generate the Openness Intuition in cases that do not involve analysis, why couldn’t a presumed shift in standards also generate the Openness Intuition in cases of analysis? I agree there is a difference between my example question and the Open Question Test—but it is a difference that makes no difference.

In summary, because there is no such thing as just plain tallness and that tallness must always be understood as relative to a standard, the term ‘tall’ in the second clause of the test question can be understood in many different ways (by simply switching what standard relative to which we understand it). When confronted with the question: “I know he is tall, but is he tall?” we could of course assume there is no switch in standard, and thus that the question is closed. This requires us interpreting our interlocutor in an uncharitable way. Alternatively, we could interpret our interlocutor as intending to switch standards—thus explaining both the fact that the question sounds open, and why anyone would ask such a question in the first place. We expect that our interlocutors are making a genuine effort to communicate with us in an efficient and informative manner—and thus expect that their contributions to our conversation will be made in accordance with

Grice's Conversational Maxims. Thus if we can interpret their question as involving a switch in standard—instead of being a trivial, closed question—we will do so.

My analysis of goodness fails the Open Question Test (which, again, I am granting to Moore) since the question, “I know that X promotes G, but is X good?” sounds open. I have been arguing that this does not mean that my analysis is incorrect. The Open Question Argument returns false negatives in the case of putative analyses of certain relations—namely, those relations that are ascribed by context-sensitive terms. Moore's argument is thus inconclusive. If goodness were a relation that introduced different standards of evaluation in different contexts, then it might be the case that goodness is not simple. In the next chapter, I will provide reasons for thinking that our judgments about goodness are *in fact* context-sensitive in the same way that our judgments about tallness are.

III. Geach's Linguistic Argument

In the last few decades, a handful of philosophers have challenged the traditional view that goodness is a monadic property.¹⁹ Most of them, at one point or another, rely on an observation about the “the peculiar logic of the term ‘good’”²⁰ described by Peter Geach in his now seminal article, “Good and Evil.” There, Geach distinguished between two kinds of adjectives: those that were logically *predicative*, and those that are logically *attributive*.²¹

¹⁹ See footnote 1

²⁰ Geach (1956) p. 40

²¹ For simplicity's sake, I will henceforth simply use ‘predicative’ and ‘attributive’ to mean ‘logically predicative’ and ‘logically attributive’ respectively.

Predicative adjectives, Geach tells us, exhibit two noteworthy logical features. First, consider sentences of the form: “That is a ADJ NOUN” where ‘ADJ’ is a *predicative* adjective. If such a sentence is true we can infer the truth of the sentences “That is ADJ” and “That is a NOUN”. So for instance:

- (1a) That is a gray car.
- (1b) That is a car.
- (1c) That is gray.

Because (1a) is true, it follows that (1b) and (1c) are also true.

Second, Geach observed that if there are two or more (common) noun terms that correctly describe the same thing, then we can interchange these noun terms in sentences of the form ‘That is a ADJ NOUN’ without loss of truth, if the adjective is predicative. So for instance, my car is both a car and a Toyota Corolla. Thus, if “That is a gray car” is true of my car, so is:

- (1d) That is a gray Toyota Corolla.

The other type of adjective, what Geach calls an *attributive* adjective, does not exhibit these same logical features. Consider the following example:

- (2a) George is a big flea.
- (2b) George is a flea.
- ** (2c) George is big.
- ** (2d) George is a big animal.

Because ‘big’ is an attributive adjective, we cannot infer the truth of (2c) from (2a). And even though all fleas are animals, and it is true that George is a big flea, we cannot conclude that (2d) George is a big animal.

Geach noticed that the predicate ‘good’ does not exhibit the logical features of a predicative adjective, and thus concluded that it is an attributive adjective. Consider:

- (3a) That is a good chess move.

(3b) That is a chess move.

** (3c) That is good.

And although a chess move is a human act, even though it is true that some particular move is a good chess move, it does not follow that:

** (3d) That is a good human act.

Why do these two types of adjective exhibit different logical features? Geach suggests an answer. Attributive adjectives like 'big' do not pick out a property. There is no such property as that of being *just plain big*. Instead, any judgment of bigness must be understood as relative to some criteria, and the criteria for bigness judgments are determined by what kind of thing we are talking about. A big atom might only be a few picometers across (that is, a few trillionths of a millimeter), whereas a big galaxy will be measured in light years (that is, in trillions of kilometers). So for any particular measurement of length, a thing of that length will be big for some things, but small for others. There is thus no particular size that is *just big*. Similarly, concludes Geach, there is no such thing as *just plain goodness*. The criteria for goodness are determined by what kind of thing it is we are talking about. So, sentences of the form "That is good" or "Pleasure is good" that do not specify what kind of thing we are talking about, are either meaningless, or are rendered meaningful by some contextually indicated kind term. As Geach puts it:

Even when 'good' or 'bad' stands by itself as a predicate, and is thus grammatically predicative, some substantive has to be understood; there is no such thing as being just good or bad, there is only being a good or bad so-and-so (Peter Geach, "Good and Evil," 1956, p.34).

Earlier, I noted that sentences like 'Coffee is good' on their surface, appear to be ascribing a monadic property. What Geach says is that this surface grammar is deceiving.

Although sentences like ‘Coffee is good’ appear to ascribe monadic properties—they have implicit references to kinds that suggest that ‘good’ actually ascribes a relation.

Geach claims that ‘good’ is an attributive adjective. He concludes from this fact that all goodness, including moral goodness, is relational. How does Geach move from observation about the logical properties of sentencings containing certain kinds of adjectives to the metaphysical claim that all goodness is relational? Unfortunately, the article is very brief, and much of his argument is not explicitly stated. In the next section, I will do my best to extract an argument from what Geach does say explicitly.

a) A Geach-Inspired Metaphysical Argument

Geach might be taken to argue as follows: For every monadic property, P, and every entity, E, either E has P, or it does not.²² Any sentence, S, that ascribes P to E would be true if and only if E has P. It should not matter for the truth of S how E is described (so long as the description is accurate), nor should it matter what kind term it is that E falls under; all that matters for the truth of S is (a) that S ascribes P to E, and (b) that E instantiates P. So, if the truth of S *does depend* on how we describe E, or on what kind term E falls under, then it follows that S does not ascribe a monadic property to E. The truth of sentences predicating an attributive adjective of a noun depends on what noun or noun phrase we use to refer to E (This premise is the one Geach spends much of his time arguing for). Thus, sentences containing attributive adjectives do not attribute a monadic property to E. Thus, attributive adjectives do not express monadic properties. ‘Good’ is an attributive adjective (Geach argues for this premise as well). Therefore, ‘good’ does not express a monadic property.

²² I should probably add, “at time T_1 , to allow for cases where E has P at one time, then loses it later (or lacks it at an earlier time, and gains it later). To keep things simple in the main text, I am omitting these references to time.

Central to this argument is the claim that the truth of sentences that ascribe monadic properties to things does not depend on how we describe those things. What can be said to motivate this premise? The thought is this. Suppose I have a term ‘X’ that successfully refers to an entity E, and an adjective ‘A’ that expresses a monadic property P. Suppose that it is true that E instantiates P. Any sentence that expresses the proposition that E instantiates P will be true, and any sentence that expresses the proposition that E does not instantiate P will be false, since E really does instantiate P. When I predicate my referring term ‘X’ with my adjective ‘A’ to form a sentence, that sentence expresses the proposition that E instantiates P. That sentence is thus true. Notice, however, that I explained this in a way that was completely content-neutral. This argument will work with whatever referring term we plug in for ‘X’. Thus, the truth of sentences that ascribe monadic properties to things does not depend on how we describe those things.

b) Objections to Geach’s Argument

Geach’s argument also crucially depends on the claim that ‘good’ is an attributive adjective. This premise appears explicitly in the article, and is one that Geach spent a fair amount of time defending (That is, it is not a premise that I cooked up to make sense of Geach’s implicit argument). R.M. Hare was the first to take issue with Geach’s support for this premise. In “Geach on Good and Evil,”²³ Hare argues that the linguistic data that Geach relies on do not adequately support the claim that ‘good’ *always* operates as an attributive adjective. Recall, Geach supports this claim by way of examples—only a very small handful of examples at that. At best, then, Geach has shown us that *some* uses of ‘good’ operate attributively.²⁴ Perhaps the goodness of *artifacts* like cars must always be

²³ Hare (1957)

²⁴ Hare, (1957) pp.107-8

relative to a kind term. However, unless he can show that there is no predicative use of ‘good,’ it remains possible that ‘good’ sometimes behaves predicatively, and thus, possible that ‘good’ sometimes expresses a monadic property.

Michael Zimmerman expands on this objection. He argues that not only is it fallacious for Geach to conclude from the fact that ‘good’ sometimes operates as an attributive adjective that it never operates as a predicative adjective, but he also claims that in fact, ‘good’ *does* sometimes operate as a predicative adjective.²⁵ Consider, for example:

(4a) Giving to charity is a good human action.

If (4a) is true, it seems like so is (4b), (4c), and (4d):

(4b) Giving to charity is a human action.

(4c) Giving to charity is good.

(4d) Giving to charity is a good way to spend one’s time and energy.²⁶

What is most worrying about this objection, for Geach’s supporters, is that the kind of goodness that Zimmerman appeals to in (4a)-(4d) might be *moral* goodness. As I mentioned before, many of the people who believe that goodness is monadic are only really interested in discussing moral goodness—they thus often ignore non-moral goodness. If Zimmerman’s counterexample is successful, it remains open for the supporters of the orthodoxy to say that non-moral goodness is relational, while moral goodness is not. Geach has not shown otherwise.

Geach believes, however, that he has already explained why Zimmerman’s alleged counterexample falsely appears to be a predicative use of ‘good’. He says that

²⁵ Zimmerman (2001) p.18

²⁶ This is my own example. Zimmerman uses the sentence “X is an intrinsically good state of mind.” (2001) p.21. I avoided using his example due to his focus on ‘intrinsic’ goodness (which he later goes on to equate with moral goodness).

when ‘good’ is used without an explicit reference to a kind term, we either contextually determine what kind term is implied, or we read a kind term into the sentence. Giving to charity belongs to the kind human action, so we naturally read into (4c) “Giving to charity is good (qua human action)”. Geach can accept that the truth of (4c) is guaranteed by the truth of (4a) when interpreted in this way (since they would mean the same thing). What Geach cannot accept, however, is that (4c) expresses that charity is just plain good—or in Zimmerman’s words ‘generically good.’ He must insist that despite appearances, (4c) really states that charity is good relative to a kind.

Zimmerman has two replies to the previous defense of Geach’s position. First, suppose we agree that Geach is correct and that there is no such thing as generic (or just plain) goodness. Who cares? That is not the notion of goodness that is important for moral philosophy. We could be rid of it, and could still make all the goodness judgments we believe to be relevant to ethics. Those judgments, after all, are about *moral* goodness—not generic goodness.²⁷ Second, since our interest is in *moral* goodness, the real question should be whether sentences of the form “X is morally good” exhibit the logical features that we expect from sentences containing predicative adjectives or not. We can thus set aside the issue of the relation of moral goodness to generic goodness. But, says Zimmerman, once we run Geach’s test on an example discussing moral goodness, we no longer see the entailments that Geach predicts for ‘good’:

- (5a) Giving to charity is a morally good human action.
- (5b) Giving to charity is a human action.
- (5c) Giving to charity is morally good.
- (5d) Giving to charity is a morally good way to spend one’s time and energy.

²⁷ Zimmerman in fact believes there is such a thing as generic goodness, but notes that this is “strictly by the by.” *Ibid*, p.21

(5a) appears to have all the logical features we expect from a sentence with a *predicative* adjective. Thus, ‘morally good’ is a predicative adjective—and it thus ascribes a monadic property.

I am inclined to agree with Zimmerman that his counterexamples successfully show that ‘morally good’ operates like a predicative adjective. Yet we have to be careful about what conclusion to draw from this fact. One of Geach’s premises was that if a predicate ascribes a monadic property, it will operate as a predicative adjective. This is a necessary condition for a predicate’s ascribing a monadic property, not a sufficient one. That means that we are not licensed to conclude that *any* term that operates as a predicative adjective is thus one that ascribes a monadic property. Further, even if ‘morally good’ introduces a monadic property, this property might still be relational. Recall that monadic properties are properties that have only one place. A relation that has two places, such as the relation expressed by ‘tall,’ can become monadic if we fill one of those places. It would then be a relational property. So for instance, ‘tall for a kindergartner’ expresses a monadic relational property. Given that ‘tall for a kindergartner’ expresses a monadic property (albeit, a relational one), we would expect for ‘tall for a kindergartner to behave like a predicative adjective.

Perhaps, then, a similar thing is happening in Zimmerman’s counterexamples. That is, perhaps goodness is a two-place relation, and in sentences about moral goodness, one of those places has been filled—making it a monadic, relational property. If this were the case, then even though ‘morally good’ expresses a relational property, it would still behave as a predicative adjective. To test this hypothesis we need to run Geach’s test on a sentence containing a predicate that clearly expresses a monadic, relational property

and see if it behaves as an attributive adjective. One notion of goodness that is clearly relational is the notion of well-being, or goodness for. This is the notion of goodness used in sentences like: “Eating a little bit of chocolate everyday is good for you.” We know this type of goodness is relational, since what is good for me is not always good for you, what is good for human beings is not always good for dogs, etc.²⁸ There is no generic goodness for, it is always goodness for *somebody*. Let’s try the following test sentence:

(6a) Broccoli is a vegetable that is good for you.

If (6a) were true, then it seems, the following sentences would be true as well:

(6b) Broccoli is a vegetable.

(6c) Broccoli is good for you.

(6d) Broccoli is a snack that is good for you.

We thus have a predicate, ‘good for you,’ that expresses a relational type of goodness, and yet behaves like a predicative adjective.

The above example demonstrates that Zimmerman’s own counterexamples should not be taken to show that moral goodness is non-relational. After all, moral goodness could be a monadic, relational property. Zimmerman’s counterexamples merely show that Geach’s test is inconclusive. Adjectives that express monadic properties, whether they are relational or non-relational, behave like predicative adjectives. Thus our discovery that “morally good” behaves like a predicative adjective does not give us any insight into whether or not moral goodness is relational.

²⁸ Here I follow Richard Kraut in his claim that “it is a conceptual truth about a type of thing that is good for someone that possibly that same type of thing is not good for someone else” (2011) p.70. To support this claim, Kraut uses a cross-species example of horses and humans. I prefer the cross-species example of dogs and chocolate, given by Finlay (2014), p.25.

IV. Conclusion

I have here considered two different kinds of linguistic tests that were meant to support metaphysical conclusions about the nature of goodness. Moore supported his claim that goodness is simple with his Open Question Argument. He argued that incorrect definitions would fail his Open Question Test. But Moore needed to make the stronger claim that *all* and only correct definitions would pass his Open Question Test—and this claim is not supported. As I argued, there can be other confounding factors that might make a question sound open: notably, if an adjective ascribes a relation where one of the places of that relation is filled by some contextually indicated entity.

Geach's linguistic tests did not fare any better. He proposed that there are two kinds of adjectives—predicative and attributive adjectives—and that sentences ascribing these different kinds of adjectives have different logical properties. 'Good' does not behave like the adjectives that clearly express monadic properties, thus goodness is not a monadic property. I agreed with Hare and Zimmerman, however, that Geach's argument does not support his conclusion that all goodness is relative to kinds. For all that Geach has said it remains open that moral goodness might be non-relational.

Chapter Two – My argument that goodness is relational

I. Introduction

In the last chapter, I argued that Geach's famous argument in favor of a relationalist account of goodness was insufficient. The problem was that while 'good' clearly operates as an attributive adjective in most cases, it certainly seems to operate as a predicative adjective in others. Now, for many of the cases where 'good' operates as a predicative adjective, Geach wants to claim it nonetheless introduced a kind-relational attribute. For those cases that cannot be easily explained in this manner, Geach simply dismisses these predicative uses of 'good' as deviant. Unfortunately, this move is unwarranted and ad hoc.²⁹

In this chapter, I propose my own argument in support of the relationalist view of goodness. I focus, not on the so-called 'logical properties' of judgments that ascribe goodness to things, but on a definition of goodness. Unlike Moore, I am optimistic that there is a common definition that applies to all good things. My argument supports the conclusion that when we judge something to be good we ascribe to that thing end-relational goodness. That is, we judge that it stands in the relevant relation to a goal. After providing my argument I consider a series of potential objections. I consider what theoretical costs there are in denying one or more of the premises of my argument. Ultimately I conclude that denying any of my premises requires accepting a view of

²⁹ As I will discuss in Chapter Four, Geach thinks that 'good thing' or 'good event' are meaningless expressions. He concludes that they involve an "illegitimate" use of 'good' (1956) p.34.

goodness that is less plausible than my end-relational account of goodness. There is thus good reason to affirm the premises of my argument, and consequently, to affirm an end-relational account of all goodness.

II. My Argument

I begin this section with a brief statement of my argument that all goodness is (end)³⁰ relational. I will then proceed to explain and justify each premise in what follows. My hope is that by having this sketch of the argument come early in the chapter, it will make the next sections easier to follow. I do not want us to miss the forest for the trees, so to speak!

- (1) When we judge a thing to be non-morally good, we ascribe end-relational goodness to that thing.
- (2) Unless the word ‘good’ is ambiguous, it would ascribe the same attribute in moral judgments as well as in non-moral moral ones.
- (3) ‘Good’ is not ambiguous.
- (C) Therefore, when we judge something to be morally good, we ascribe end-relational goodness to that thing.

I should note that I have intentionally decided to gloss over a few complexities and finer points to keep this representation of my argument pithy. The messy details will all come out in the discussion that follows.

a) Premise (2) – The ‘Thesis of Commonality’

I will begin the defense of my argument with a discussion of the second premise. This might initially seem like a counterintuitive place to begin, but discussion of Premise

³⁰ While I speak of end-relational goodness in the premises, the argument of this chapter alone is, strictly speaking, not sufficient to show that all goodness is *end*-relational. The argument of this chapter shows that non-moral goodness is *relational*—but is inconclusive on what it is, exactly, that good things are good in relation to. In Chapter Four, I argue that it is ends—or goals—that good things are good in relation to.

(2) will be revealing of some of my philosophical and methodological starting points and background assumptions. I am of the belief that it is best to get these out into the open as soon as possible. Most importantly, my defense of premise two will reveal something about how I view the connection between language and metaphysics. My hope is that my claims about the connection between language and metaphysics do not require taking on too many substantive positions on controversial questions. I hope that the story that I give is one with which nearly everyone can agree. Unfortunately, I suspect that I am being naive in this wish.

Premise (2) states that unless the word ‘good’ is ambiguous, it will ascribe the same attribute³¹ in moral judgments as well as in non-moral moral ones. At first glance, this premise might strike one as clearly false. My imagined objector might say that upon very little reflection, it becomes clear that our judgments about goodness ascribe *very different* attributes to the things we are judging. There is no common feature among good things that makes it appropriate to call them all ‘good.’ For instance, when we judge that a piece of sandpaper is good we are judging that it is gritty and scratchy to the touch, and could be used to scrape away the roughness of a piece of lumber.³² On the other hand, when we judge that a piece of silk is good, we are judging that it is smooth and soft. We would not want to run a piece of good silk over a rough piece of lumber (and were we to do so, surely we wouldn’t worry about the outcome the contact would have on the wood)!

³¹ As in Chapter One, I am using the word ‘attribute’ instead of the word ‘property’ at least partly because, in what follows, I will be arguing that goodness is not a property, but a relation. I thus think it is false to say that our judgments about goodness attribute the same *property* to the things we are judging. However, it would also be premature to use the word ‘relation’. Instead, I opt for a neutral term, ‘attribute,’ that I have been using to stand for something that is either a property or a relation.

³² Ziff uses the examples of sweetness (good strawberries are sweet) and sourness (good lemons are sour) to make this same point (1960) pp.202-3. Finlay uses the examples of sedatives/stimulants, and buoys/anchors (2014) p.20.

Thus, concludes my objector, sometimes our judgments about goodness ascribe softness to the thing we are judging, while other times they ascribe roughness. Further, given the variety of things that we can judge to be good, and the variety of attributes that make those things good, there will be many attributes, F, such that there will some things that are good in virtue of having F and some things are good in virtue of having non-F.³³

Therefore, concludes my objector, our judgments of goodness do not ascribe the same attribute—in fact, they not only ascribe a multitude of attributes, they even often ascribe *opposing* attributes!

This objection seems to be a specific instance of a more general argument given against conceptual analysis by Ludwig Wittgenstein in *Philosophical Investigations*. He wrote:

Consider for example the proceedings that we call ‘games’. I mean board-games, card-games, ball-games, Olympic games, and so on. What is common to them all? -- Don’t say: ‘There *must* be something common, or they would not be called “games”’ – but *look and see* whether there is anything common to all... And the result of this examination is: we see a complicated network of similarities overlapping and criss-crossing: sometimes overall similarities, sometimes similarities of detail. I can think of no better expression to characterize these similarities than ‘family resemblances’... (Wittgenstein, *Philosophical Investigations*, 1953, Aphorisms 66 & 67)

As it is traditionally understood, conceptual analysis involves finding necessary and sufficient conditions for the analysandum. That an activity involves a board on which to place pieces, or the use of a ball, or teams competing, or a winner and loser, etc. are not necessary conditions for being a game; some games do not have those features. Further,

³³ My objector might be attracted to the stronger claim: that for literally *every* attribute F, there will be some thing X, such that our judgment that X is good ascribes F to X. I am doubtful this will be true. However, all my objector needs to make her point is that for many, if not most, attributes this will be the case. We can thus expect a very wide variety of attributes that we ascribe to things when we judge them to be good.

none of those conditions are sufficient for an activity to count as a game. A person might use pieces on a board to demonstrate the relative locations of buildings downtown, a ball can be used when doing certain yoga positions, teams of researchers might compete to create and release the next great technical device, wars can have winners and losers, etc. So, concludes Wittgenstein, no set of necessary and sufficient conditions can be given even for simple, everyday notions like that of a game. If he is right, our chances of giving an analysis of a more complicated, philosophical notion—like goodness—will be bleak. According to Wittgenstein, just because we use the same word to refer to a group of entities, this does not mean they share some set of common features. More importantly for my project: If Wittgenstein’s argument is successful and generalizes, then ‘good’ does not ascribe the same attribute across both moral and non-moral uses.

I believe that Wittgenstein is mistaken about this. In some cases, we can identify the common feature(s) among things that are appropriately grouped under a common term. Because I see my imagined objector as offering a specific instance of Wittgenstein’s objection (though specifically directed towards proposed analyses of goodness), a good response to Wittgenstein would be, at the same time, a response to my objector. As such, I will now rehearse a few compelling refutations of Wittgenstein’s objection.

First, it is difficult to see exactly how Wittgenstein’s objection is meant to lead to skepticism about conceptual analysis. He has not shown it to be *impossible* to give necessary and sufficient conditions for the concept of a game. He has merely shown that giving necessary and sufficient conditions can be a very difficult thing to do—a claim that I suspect few philosophers would deny! It is a big jump, however, from showing that

a task is difficult to accomplish to showing that it is impossible. One might further support Wittgenstein's conclusion, however, by noting that if a good number of smart philosophers work quite long and hard at finding an analysis of a concept, and yet fail, this is good evidence that the task is impossible. This would be a strange defense of Wittgenstein, however, given his particular example: games. After all, part of the set-up of that alleged counterexample was that the concept of a game was an ordinary concept—one that has not received much philosophical investigation. Thus, there is no long history of philosophers trying and failing to analyze the notion of game.

In fact, quite to the contrary, in his book *The Grasshopper: Games, Life and Utopia* Bernard Suits appears to have succeeded at analyzing the notion of a game.³⁴

Suits claims that:

To play a game is to engage in activity directed towards bringing about a specific state of affairs, using only means permitted by rules, where the rules prohibit more efficient in favour of less efficient means, and where such rules are accepted just because they make possible such activity (Bernard Suits, *The Grasshopper*, 1978, pp.48-9).

Whether this analysis is correct or not is not for us to decide here. It is at least quite close to being successful at a task that Wittgenstein thought was impossible.³⁵ One point worth noting, however, is that this analysis is *general* and attempts to identify *essential* features

³⁴ Suits' book was published after Wittgenstein's death, so we cannot fault Wittgenstein for not discussing Suits' analysis.

³⁵ Consider two of Suits' examples: in the game of golf the goal is to get the golf ball into the cup. The easiest way to do so is probably just to walk over and drop the ball in. That method is forbidden, and only a less efficient manner of accomplishing this task (hitting the ball with clubs from a distance) is permitted. People accept this limitation in order to enjoy the game of golf (1978, p.38). Similarly, the goal of chess is to immobilize your opponent's king. The easiest ways of accomplishing that might be to physically overpower your opponent or to glue her king to the board. These methods are not permitted by the rules of chess, and people who do not accept these limits would no longer be playing *chess* (*Ibid*, pp.67-9).

of games.³⁶ As Thomas Hurka points out in the Introduction to *The Grasshopper*, the structural features that Suits points to could easily be true of many various types of game (including board games, card games, etc).³⁷ On the other hand, Wittgenstein’s attempt at analyzing the concept of a game was superficial, resting “only [on] surface differences between games... without even wondering whether they may not be consistent with a deeper commonality”³⁸ It is silly to think that an analysis of the concept of a game will be given by looking only at the very specific and superficial attributes of particular games. Similarly, it is silly to think that an analysis of goodness will be given by looking only at the very specific and superficial attributes of the things we judge to be good. In both cases, we would want an analysis that is quite *general* and that identifies some *essential* attribute(s) of the analysandum. If there is anything in common among all games, it will clearly not be the number of players, equipment used, scoring method, etc. Similarly, if there is anything in common among all good things, it will not be specific properties such as the roughness or softness of those things.

The word ‘game’ refers to a general type of entity—one that admits of several different sub-types (e.g. ballgame, board game, card game, etc). The definition given by Suits suggests a general definition of games that serves as a way of distinguishing games from other entities. Once we have identified the more general definition of game, we can then give specific definitions for what it is to be a ballgame, board game, etc by tacking on additional necessary conditions. A board game, for instance, is a game (however that

³⁶ Technically, it is an analysis of playing a game. I suspect one could easily construct an analysis of games derivatively once she has an analysis of game-playing in her pocket.

³⁷ *Ibid*, p.12

³⁸ *Ibid*, p.12

is analyzed) that is played with a board.³⁹ The mistake that Wittgenstein made was confusing the additional necessary conditions, which demarcate types of games from one another, with necessary conditions that demarcate games from other entities. Now we could have tried to arrive at a general analysis of game by first giving an account of board games and then trying to subtract all and only those conditions that are required to distinguish board games from other types of games. I suspect that this would be tricky. In doing so, it would be easier to make Wittgenstein's mistake—thinking that features that are necessary for distinguishing board games from other types of games are the features that might distinguish games from non-games. Suits' approach was more promising. He started by considering many types of game and looking for very general features they shared in common. He took seriously Wittgenstein's advice to "*look and see* whether there is anything common to all". It is very plausible that he was successful—there really is something that all games share in common, and in virtue of that fact, it is appropriate to call them all 'games'.

My argument for Premise Two (unless the word 'good' is ambiguous, it would ascribe the same attribute in moral judgments as well as in non-moral moral ones) is that we can give a very similar story about the nature of goodness as the one we just gave about the nature of games. Just as in the case of 'game', the word 'good' expresses a general type of attribute—one that admits of several different sub-types (or to use

³⁹ Two caveats are needed. First, surfing is 'played with a board,' but is not a board game. For this reason, if I were really interested in giving an accurate account of board games, I would have to more clearly define what sense of 'board' I am talking about. Second, there are some who would want to categorize role-playing games like Dungeons and Dragons as a 'board game' since it is typically played sitting around a table, and rolling dice. There is no board in D&D, though, so on my definition, it would not be a board game. Whether or not that is correct would be an important issue to resolve were I sincerely trying to give an account of board games. Since I am merely using this as an example, I will not be taking stances on these issues.

Thomson's language, *varieties of goodness*).⁴⁰ There is moral goodness, prudential goodness, aesthetic goodness, etc. Most philosophers who discuss goodness are ultimately interested in saying something about *moral* goodness. They will, for that reason, focus their inquiry on just that variety of goodness. It is rare that these philosophers ever think about the other varieties of goodness. I call this the 'Top-Down Strategy'—and it is one that can lead to the very same mistake that Wittgenstein made. That is, it can lead us into thinking that features that distinguish moral goodness from other kinds of goodness are also features that distinguish goodness from all the things that are not goodness. For example, it is widely agreed upon by moral philosophers that moral facts (like those about moral goodness) are normative. This feature of moral facts has led some philosophers (Derek Parfit and David Enoch, for instance) to conclude that moral goodness must be a non-natural property. Roughly, their argument is this: No facts that ascribe purely natural properties to things are normative; facts that ascribe moral goodness to things are normative; thus facts that ascribe moral goodness do not ascribe a purely natural property.⁴¹ That is, moral goodness is a non-natural property. Now to be clear, neither Parfit or Enoch were interested in saying anything about the other types of goodness other than moral goodness. Notice that if we were to think that normativity is an essential feature of *all* varieties of goodness, however, then Parfit and Enoch's argument would push us into accepting non-naturalism about non-moral goodness. If on the other hand, we were to think that normativity is only a feature of moral goodness, their argument sheds no light on what distinguishes good things from those things that are

⁴⁰ Thomson (1992, 1993, 1997, 2001, 2003, 2008)

⁴¹ Derek Parfit gives a variant of this argument in *On What Matters Volume Two*, pp.324-7; David Enoch gives it in *Taking Morality Seriously*, pp.80-1.

not good.⁴² I do not raise this example to discuss the merits of Parfit and Enoch’s argument—I raise it to discuss the pitfalls of the Top-Down Strategy. I believe that we cannot know if normativity is a general feature that all goods share if we do not spend any time thinking about non-moral goods.

For this reason, I suggest a ‘Bottom-Up Strategy’ similar to the one Suits uses in his discussion of games: start by considering many types of goodness and looking for very general features they share in common. We should not focus on moral goodness, thereby overlooking the fact that everyday we judge countless things as good in ways that are not (typically) regarded as moral. We judge the goodness of newspapers, websites, parking spaces, restaurants, conversation topics, clothing, hairstyles, places to eat lunch, student papers, etc, etc.⁴³ Due to the pervasiveness of our non-moral judgments, one would expect that we would be well situated to answer questions about the nature of non-moral goodness. So why not begin our investigation into the nature of goodness with an investigation into non-moral goodness? If we can identify very general features that many kinds of goodness share in common, then this is likely to be what it is in virtue of which it is appropriate to call them all ‘good.’ I call this the *Thesis of Commonality*—that there is something in common among all good things such that it is appropriate to call all of them ‘good.’⁴⁴ It is my main justification for Premise Two. My thought is that we should

⁴² This is not a problem for Parfit and Enoch’s projects—since they are not inquiring into the nature of goodness. It would be a problem for *our* project, however.

⁴³ Finlay provides a similar list (2014), p.19

⁴⁴ G.E. Moore explicitly endorses the Thesis of Commonality. He writes: “And on the other hand, other things, beside conduct, may be good; and if they are so, then, ‘good’ denotes some property, that is common to them and conduct” (1903), p.54. Other places we see philosophers asserting the Thesis of Commonality: Ziff, (1960), p.203, Mackie, (1977), pp.51-2, Thomson, (2008), p.10, and Finlay, (2014), p.19.

expect that the Thesis of Commonality is true so long as the word ‘good’ is not ambiguous. In the next section I will turn to my argument that ‘good’ is indeed univocal.

First, however, I would like to finish my response to my imagined objector. As you will recall, she was skeptical that there were any features in common among good things—since good silk must be soft, and good sandpaper must be rough. I think we are now in a position to give her the start of an answer to the question: what does good silk have in common with good sand paper? A first stab at an answer is that to say that a piece of silk is good is to say that it *meets some standard of what is expected from silk*.

Similarly, a good piece of sandpaper *meets some standard of what is expected from sandpaper*. We might then generalize and say that X is good if and only if X meets some standard of what is expected from things of its kind. While I don’t think this analysis is quite right, it is at the very least, the kind of analysis we are looking for.⁴⁵ For one, it is the kind of analysis that is (a) general enough that it might apply to all things that we judge to be good, and further, (b) it explains why it is the case that the same feature can make one thing good, while at the same time making some other thing bad. The standards that we judge silk by will surely mention that we expect silk to be soft. For a piece of silk to meet this standard, and thus be deemed ‘good,’ it must be soft. Thus, softness is a good making property *for silk*. The standard we evaluate sandpaper by, on the other hand, will *not* include a condition that sandpaper be soft. In fact, it will require that sandpaper be rough. So softness is not a good making property *for sandpaper*.

⁴⁵ As stated, this proposal is not fully developed. What sets the standards? If it is the *kind* that sets the standards, the view will be kind-relational, and thus, quite like the view of Geach (1956), Thomson (2009), and Foot (2001). If it is *interests* that people take in those kinds that sets the standards, the view will be interest-relational and then, more akin to the views advanced by Ziff (1960), and Mackie (1977). Finally, if it were *ends* that set the standards, the resulting view would be similar to the accounts offered by Finlay (2014) and myself.

What I have shown is that there is at least one plausible contender for a general feature that all good things have in common, in virtue of which it is appropriate to call all those things ‘good.’ This is what we should expect if the word ‘good’ is not ambiguous.

b) Premise (3) – That ‘good’ is univocal

A word is lexically ambiguous when that word has two (or more) distinct meanings.⁴⁶ The word ‘bat,’ for instance, in some instances refers to a flying mammal, while in other instances it refers to a piece of sporting equipment.⁴⁷ ‘Bat’ is thus ambiguous. Premise three, on the other hand, states that the word ‘good’ has only one meaning. Let me say a few things about how I understand this claim.

Since ‘bat’ has at least two meanings, a sentence containing the word ‘bat’ can have multiple meanings as well. Consider the sentence: ‘That is a bat.’ Because ‘bat’ is ambiguous, this sentence could express either of two judgments: it could express the judgment that the thing that is being pointed at is a flying mammal, or it could express the judgment that it is a piece of sporting equipment. Not so for ‘good,’ says the third premise. The sentence: ‘This is good soup’ just expresses the judgment that ascribes to the soup the attribute of goodness (whatever that attribute turns out to be). There is no other judgment that might be expressed by the sentence ‘This is good soup.’

What evidence is there that ‘good’ is univocal? The first—and least technical—bit of evidence that ‘good’ is univocal is that on reflection the word does not seem to have two or more distinct meanings. It might be strange of me to say this, especially given how

⁴⁶ Most of my discussion of ambiguity is informed by Adam Sennet’s Stanford Encyclopedia of Philosophy entry: “Ambiguity” (2011).

⁴⁷ To keep things simple, I am ignoring the various other meanings of the word ‘bat’. For instance, I am bracketing the uses of this word that describe the action of swinging a bat.

frequently I am drawing the distinction between moral and non-moral goodness. Doesn't this distinction suggest that 'good' does have (at least) two meanings?

While it is tempting to think of non-moral and moral goodness as two distinct things, we should not confuse the *generality* of the word 'good' for ambiguity. Consider obviously ambiguous words like 'bat,' 'crane,' or 'bank.' For these terms, it seems to be a contingent and coincidental fact about English that we use them to talk about two (or more) different types of entity.⁴⁸ There is no connection between the flying mammals and the baseball equipment that would be lost if we started calling one of them by a different name. Thus, there is good reason to think that these words are ambiguous. Now consider, instead, a general term like 'ship'. I say that 'ship' is a general term because the analysis of the concept that it expresses applies to a variety of different types of entities. We call many vessels that travel by water 'ships'—as well as some vessels that travel through the air or through space. We apply the same term, 'ship' to these distinct types of entities because they share in common certain features such that they all meet the necessary and sufficient conditions for being a ship. Spaceships and seafaring ships are thus two distinct *categories* of ship. If we stopped calling one of them 'ship', it would not change the fact that (a) the thing we are talking about is a ship, and that (b) it shares this feature in common with some other types of entity. Because we do call both seafaring and space-faring vessels 'ships,' our language expresses that we are aware of the facts (a) and (b). Were we to only call one type of ship a 'ship', our language would no longer express that

⁴⁸ Perhaps 'crane' is a bad example. One might think that the piece of construction equipment called a crane was given that name due to its resemblance to the bird. If we stopped calling either a crane, we would lose this connection. However, we could have equally named the construction equipment a 'giraffe,' if all we wanted to express was that the machine resembles an animal with an elongated neck.

spaceships and seafaring ships are all related to one another insofar as they are both categories of the more general type: ship. Our language would be impoverished.

My intuition is that moral and non-moral goodness are two *categories* of goodness, and not two distinct notions contingently picked out by the same English word. There is something in common between things that are non-morally good and those things that are morally good, in virtue of which they are all properly categorized as *good*. To test my intuition, we could ask if anything would be lost if we started using two very different terms for non-moral and moral goodness. It seems to me that we would lose something. Consider, for instance, the debate between philosophers like L.W. Sumner who believe that moral goodness should be identified with *goodness for*, or well-being,⁴⁹ and philosophers like Michael Zimmerman who believe that moral goodness should be identified with *intrinsic goodness*. If we did not use the same word ‘goodness’ for these three notions—goodness for, intrinsic goodness, and moral goodness—this debate might sound quite bizarre. The possible version of English that has three very different terms for these three notions of value is less expressive, and thus impoverished when compared to the actual version of English where those three notions share a term.

In addition to imagining possible versions of English and asking whether or not those variations would continue to express certain relations between things, we can also look to actual non-English languages and see if they mirror our uses of words. If in other languages the speakers use a single word to talk about moral and non-moral goodness, this is evidence that the notion of goodness is in fact a general notion (and not ambiguous). If, on the other hand, it is merely a coincidence that we use the word ‘good’ to refer to two distinct notions—then other languages will likely use two distinct words to

⁴⁹ Sumner (1996)

refer to those notions.⁵⁰ The cross-linguistic data, however, show that speakers of other languages use one word to talk about both notions.⁵¹

There are also a number of linguistic tests we can run on a word to see if it is ambiguous. These tests, unfortunately, are not all that informative. They sometimes confuse generality and/or context sensitivity with ambiguity. Since I believe that ‘good’ is both general *and* context sensitive—it comes as no surprise to me that it sometimes fails these tests.

A term is *context sensitive* if the content that the term contributes is variable “due purely to changes in the context of utterance without a change in the convention of word usage.”⁵² As we saw in the last chapter, the word ‘tall’ is context sensitive. When we are talking about kindergartners, a child that stands 5 feet would be tall. In the context of professional basketball players, a person that stands 7 feet would be tall. We should not conclude that the word ‘tall’ sometimes means standing 5 feet, sometimes standing 7 feet, and sometimes standing several hundred feet tall (when we are discussing redwood trees, for instance). If that were the case, ‘tall’ would be massively multiply ambiguous. Instead, we should understand ‘tall’ as requiring some information from the context—in this case, what group it is relative to which the thing in question counts as tall—to determine its content.

As noted above, we might think that to say of something that it is good is to say that the thing in question meets some standard. Which standard we are judging the thing

⁵⁰ We would, of course, need to look at languages whose word for good is not either derived from the English word ‘good,’ or from the same word that the English word is derived from. It would not be so surprising if words in other languages that share an etymology with ‘good,’ also share some features of how that word is used.

⁵¹ Robert Shanklin discusses the data in Chapter One of his dissertation, *On Good and ‘Good’* (2011)

⁵² Sennet, (2011).

relative to might be contributed by context. Is it the silk standard that requires the thing be soft? Or is it the sandpaper standard that requires roughness? Again, we *could* think that ‘good’ sometimes means soft, other times rough, and still other times any number of other attributes. In that case, ‘good’ would be multiply ambiguous. I think this is a mistaken way to think about the meaning of the word ‘good’. ‘Good,’ like ‘tall,’ seems to have only one meaning, but that meaning is partially determined by context.

Again, I think the common tests that linguists use to determine if a word is ambiguous or not occasionally give us false positives for ambiguity when the term we test is either general or context sensitive in the way that I believe ‘good’ is. But let’s look at a couple of these tests anyway.

The Contradiction Test: A sentence that looks as if it would be contradictory (if there is only one meaning of the tested term), will not sound contradictory if that term is ambiguous. So, for instance, ‘That is a giraffe, but it isn’t a giraffe’ should sound like a contradiction to us because ‘giraffe’ is univocal. On the other hand, we should be able to imagine contexts where ‘That is a bat, but it isn’t a bat’ would sound perfectly fine and non-contradictory since the word ‘bat’ is ambiguous.⁵³ How does ‘good’ fair in the contradiction test? Consider the sentence: ‘The soup is good, but it is not good’. This sounds contradictory to me, though, I can also hear an interpretation where it does not sound contradictory. The non-contradictory interpretation would have us switching which standard it is relative to which we are judging the soup to be good from the first half to the second half of the sentence. Suppose, for instance, we are dining at a restaurant that

⁵³ You will likely notice that this test is very close in structure to Moore’s Open Question Argument. It should not surprise you, then, that my hypothesis about what occurs when you run this test on ‘good’ will sound very similar to the response to Open Question Test that I discussed above.

you believe to be subpar—you find most of the food on the menu to be terrible. However, you find this particular soup to be less bad than everything else on the menu. So, you might say: “This soup is good (relative to the rest of the slop on the menu here), but it is not good (relative what I could get at another restaurant)”. A similar thing would occur if we said of the 5 foot tall kindergartner, ‘He is tall, but he is not tall.’ In that example, we might be saying that he is tall (for a kindergartner) but not tall (for an adult male). That switch in standard makes the sentence not sound contradictory. But we should not conclude that ‘tall’ is thus ambiguous. It is just context sensitive, and the Contradiction Test can give us false positives with context sensitive terms. For these reasons the Contradiction Test is inconclusive on ‘good.’

Conjunction Reduction: In this test, we start with two sentences using allegedly different meanings of the term, and combine them via conjunction. If the resulting conjunction sounds absurd, then that is evidence that the term is ambiguous. So for instance: “The feathers are light” and “The colors are light” combine to give us “The feathers and the colors are light”. That sentence sounds odd—so we have evidence that ‘light’ is ambiguous.⁵⁴ How does ‘good’ fair on this test? Let us consider the sentences: ‘The soup is good’ and ‘Giving to charity is good’. The combined sentence ‘The soup and giving to charity are good’ does indeed sound strange. Should we conclude that ‘good’ is thus ambiguous? Again, I think here we are confusing generality with ambiguity. What makes the sentence sound odd is the switch from two different categories of goodness. In the first conjunct, we are ascribing non-moral goodness to the soup. In the second, we are ascribing moral goodness to charity. As I argued above, I do not think we should conclude that moral and non-moral goodness are two distinct types

⁵⁴ This is Sennet’s example (2011).

of thing, but merely two categories of goodness. The Conjunction Test does not seem to distinguish generality from ambiguity, and does not answer whether ‘good’ is ambiguous.

In summary, I think the standard linguistic tests are inconclusive when it comes to determining whether ‘good’ is ambiguous or not. The best evidence we have that ‘good’ is univocal is that our language would be less expressive if we used completely different terms for all the various types of goodness. ‘Good’ is general and context sensitive—but not ambiguous.

c. Premise (1) – That non-moral goodness is end-relational

So what then, we might ask, do good kites, silk, sandpaper, murder weapons, hammers, and parking spaces have in common? Earlier, I suggested that a good first guess might be that to say that X is good is to say that X meets some standard of what is expected from things of its kind. There would be different standards for what it is to be a good kite, or a good murder weapon etc. However, we might wish for a more illuminating answer. In particular, we might still ask: *what is it that determines these standards?*

To begin answering that question, it is worth noticing that depending on the set of standards we use, the very same object can be good relative to one set, but bad relative to another.⁵⁵ For instance, on one set of standards, a particular hammer will be good, but that same hammer will be bad relative to another set of standards.⁵⁶ The hammer that is good for nailing boards together to make a birdhouse would not be good for driving

⁵⁵ Ziff notes that the same object (a dead body) can be good when described as a cadaver, but that it would be strange to say of it that it is good when we describe it as a corpse (1960) p.211. He does not explicitly note that a thing could thus be good relative to one set of standards and bad relative to another—but the seeds of this thought are there.

⁵⁶ Georg Henrik von Wright uses the example of a hammer (1963) p.8.

railroad ties into the ground. A sledgehammer would be good for driving in railroad ties, but bad for putting together a birdhouse. In these cases, the *standards are set by what goal one is trying to accomplish*. To be fair, there are two other contending types of answer to what determines the standard. First, there is the view that it is the *kind* that a thing belongs to that determines the standard (so in this case, we have two kinds: sledgehammer and hammer, and they each have different standards). Second, there is the view that it is human interests that determine the standard. I reject both these kinds of views in Chapter Four. For now, I hope that it is enough that at least one plausible answer is that standards are set by goals. If one wants to build a birdhouse, one needs to precisely drive small nails through thin boards of wood. A good hammer—relative to the goal of building a birdhouse—then, is one that would adequately accomplish this task.

Hammers (but not sledgehammers) are typically used for driving nails through wood. Consequently, it is tempting to conclude from the above example that a good X is an X that adequately achieves some goal that X's are typically used for. The standards for judging the goodness of an X would be set by what it takes to achieve that goal. This analysis is unsatisfactory for a few reasons. First, we can evaluate things relative to goals that they are not typically used for. Suppose that an artist is constructing a sculpture that will resemble the skeleton of a Tyrannosaurus Rex, though it will be made entirely of modern tools painted white. She needs a small, 'T' shaped tool to stand in as one of the bones. A hammer might adequately achieve this goal: representing a certain Tyrannosaurus Rex bone in the artist's sculpture. That hammer is good relative to this goal, even though it is a strange goal that people do not usually use hammers for. It would be good for this goal, even if certain flaws in the handle, for instance, make it a

bad hammer relative to the goal of hammering in nails. Noting this does not force us to reject an end-relative account of goodness. It merely suggests that we should drop the talk of goals that X's are typically used for. Instead, we should note that because hammers are typically used for hammering, it is in general safe to assume that a person that says, 'This is a good hammer' is evaluating that hammer relative to the goal of hammering.⁵⁷ Were the artist to say 'This is a good hammer,' she might have to give her audience some contextual clues to eliminate the assumption she is evaluating that hammer relative to the goal of hammering. She should let us know that she is judging its goodness relative to the goal of representing a dinosaur bone in a sculpture, or we will likely misunderstand the claim she is making.

A deeper objection to a goal-based account of goodness is that for many things we judge to be good it seems incorrect to think of our judgments as being about those things achieving goals. What goals could I be evaluating a tiger by when I say 'That is a good tiger?'⁵⁸ Or, more troubling: What would the end-relational account say of the judgment that 'Sarah is a good person'?

First, it is worth stressing that we do not need to claim that tigers and people have *essential* functions that they serve, much as a hammer has the function of driving nails. Many criticized Aristotle for making such a claim, and I think their criticisms were apt. Above, I argued that for some objects, it is safe to assume that a judgment that that thing is good is to be understood as a judgment about that thing's achieving a goal for which it is typically used. Hammers are typically used for driving nails, so we can generally

⁵⁷ That is, it is a conventional implicature that a person is evaluating the hammer as a device for driving nails, unless this implicature is defeated by context.

⁵⁸ This is Philippa Foot's example from *Natural Goodness*, (2001), p.49. Paul Ziff expresses doubts about the grammaticality of such sentences in *Semantic Analysis* (1960). He thinks that the sentence, 'That is a good Gila Monster' sounds ungrammatical *Ibid*, p.212.

conclude that ‘good hammer’ refers to a hammer that adequately achieves that goal. However, I also noted that we could produce different judgments about the goodness of the hammer by referencing other goals. So the end-relative account of goodness I prefer is compatible with things lacking any essential function. All it says is goodness is relative to *some* goal or another.

So what goal is it that we are evaluating a tiger relative to when we say ‘That is a good tiger’? One suggestion might be the goal of being a representative member of its species.⁵⁹ Consider a zookeeper who is looking to add a tiger to his zoo. She might want to acquire an albino tiger as such a tiger is rare and will draw crowds. Or, she might want a tiger that looks and acts like most tigers in the wild; she might want a representative tiger. That is one possible interpretation of ‘That is a good tiger’. We might, alternatively, think that we are evaluating how well that tiger meets the goal of thriving in the wild. Philippa Foot seems to have this in mind. On her view, a good tiger is one that is well-adapted to survive in its environment, and perhaps to contribute to the propagation of its species.⁶⁰ Often, in the natural world, these two interpretations will return the same results. A representative tiger is also a fit-to-survive tiger. Some animals are less lucky. The Giant Panda, for instance, is generally a finicky eater and mater. In times of plenty, pandas flourish, and their species thrives. But, mostly due to human invasion of their habitat, representative pandas currently are having trouble finding food they like, and will

⁵⁹ Ziff suggests a similar answer in *Semantic Analysis*. Having already claimed that ‘That is a good Gila Monster’ sounds ungrammatical (*Ibid*, p.212), he goes on to say that ‘That would be a good Gila Monster to make a pet of...’ sounds fine p.236. Mackie explicitly states that there are uses where ‘good’ seems to mean ‘paradigmatic’ (1977) p.58. Thomson’s account of goodness properties begins with the property of “being a model, exemplar, paradigm, or good specimen of a K” (2008) p.19.

⁶⁰ She writes that the natural goodness of “plants and non-human animals... [has] to do, directly or indirectly, with self-maintenance, as by defence and the obtaining of nourishment, or with the reproduction of the individual, as by the building of nests” (2001) p.31.

rarely mate successfully. In this case, representative pandas are not well-adapted pandas. I will discuss this example in more detail in Chapter Four, when I turn to evaluating the plausibility of other relational accounts of goodness. For now, it is enough to show that what we mean when we judge that an animal (such as a tiger or a panda) is good is going to vary from case to case. I do not think I have here identified every goal relative to which we might judge a tiger good (I suspect, for instance, when an animal trainer picks a tiger to train to perform in circuses, she is not going to be interested in a tiger that is well-adapted to survival in the wild). I have instead, pointed to a couple of plausible interpretations of what someone might mean when they say of a tiger that it is good.

I have a similar answer to what goal it is relative to which we judge people to be good. Consider the sentence, spoken of Bert, that he is a good man. This sentence might express the judgment that Bert is a morally upstanding person. That is, he adequately achieves the goal (whatever it is) of being moral. Or, suppose that I am discussing Bert's merit as a possible romantic partner. In that context, 'Bert is a good man' might mean that he adequately achieves the goal (whatever it is) of being a fitting romantic partner. Again, there are many different goals relative to which we might judge someone to be a good man, woman, or person. All the end-relational account of goodness must say is that for any claim of the form, 'X is a good person,' there is some goal that sets the standards relative to which we make that claim.

In summary, when we judge that something is good, we seem to be judging it relative to some standard, and that standard seems to be set by the goal we are trying to

accomplish with the thing we are judging to be good.⁶¹ To judge that X is good, then, is to ascribe to X the attribute of adequately accomplishing a goal, G.⁶²

III. Objections

According to my view, goodness is a relation that things can stand in to ends. This is true both of non-moral and of moral goodness. In contrast, according to many other accounts, moral goodness is a different type of attribute than non-moral goodness. At first glance, there is something intuitive about the position that moral goodness is a different type of attribute from non-moral goodness. After all, why would we distinguish moral goodness from other kinds of goodness if they weren't *different*? As I hope to show in the subsequent chapters, my view can accommodate our intuition that moral goodness is different from non-moral goodness (in some respects), without needing to claim that they are two different attributes.

In this chapter, I argued that the best analysis of non-moral goodness, roughly, is that goodness is a relation that holds between things and goals.⁶³ I also argued that unless the word 'good' is ambiguous, then we ought to believe that all good things share some common feature(s) in virtue of which it is correct to call them 'good'. I suggested that this common feature of all good things is that they stand in the relevant relations to goals. Thus, the end-relational analysis I provided for non-moral goodness extends to *all* types of goodness—including moral goodness. All goodness is relational.

⁶¹ Some philosophers believe that it is kinds or interests that set the standards. I discuss the virtues and vices of these theories of goodness in Chapter Four.

⁶² In Chapter Five, I will spell this analysis out in more detail.

⁶³ Again, I will give this analysis further precision in Chapter Five. For now, nothing rides on the details of the analysis.

Many philosophers reject my conclusion. To do so, they must reject at least one of my premises. That is, in order for a philosopher to claim that not all goodness is relational, she must either (a) claim that non-moral goodness is not relational, (b) deny the Thesis of Commonality, or (c) find some other feature that is in common amongst non-moral and moral goodness such that the Thesis of Commonality is satisfied. I have already given a few reasons to think that the first two of these options are unattractive, though I will flesh out my objections further in this chapter. One possibility that I have as of yet not discussed is that our goodness judgments do not ascribe any property at all, but merely express our attitudes. This is a view that has had many supporters, especially when it comes to moral goodness. I will argue, however, that it does not give us a very plausible account of non-moral goodness.

I will then focus my attention on one plausible attempt to pursue the third strategy, one offered by Michael J. Zimmerman. I will argue that it is wrongheaded. Of course, this does not show that *any* such attempt at pursuing the (c) strategy for defending the claim that moral goodness is not relational will be equally wrongheaded. At the very least, I see my argument as offering a challenge to my opponents; if we are to believe that moral goodness is not relational, you, my opponent, owe us a plausible account of goodness that is also compatible with the Thesis of Commonality. I will give some reasons to doubt any such account is forthcoming.

a) Objection: Non-moral goodness is not end-relational

Suppose that someone wanted to support the claim that moral goodness is not relational by way of (a). There are two ways of doing so. First, she could claim that *all* goodness is monadic. According to this view, even non-moral goodness is monadic. So,

when a hammer is good, it is good because it has the monadic property of goodness.

Second, she could argue that what is *really* essential to goodness is its relation to the acts of recommending, commending, or prescribing. We do not ascribe goodness to X unless we express, in the relevant way, some pro-attitudes towards X. I think all of these answers, however, give us quite implausible accounts of non-moral goodness.

i) Non-moral goodness is monadic

Let me start with the first suggestion: that all goodness is monadic. I can see two ways of developing this view. The first is to claim that a thing's goodness supervenes on, but is distinct from its other attributes. The second is to claim that a thing's goodness does not supervene on its other attributes—it is just one additional property that good things have regardless of what other attributes they have. I believe that the latter view is, quite simply, untenable. The most plausible version of the former view appears to only really differ with my end-relational account in one respect: it needs to posit an additional property that I do not need to posit. We thus ought favor my view to it.

Suppose one were to insist that a thing's goodness does not supervene on its other properties. The resulting view would not be elastic enough to accommodate the context sensitivity exhibited by our goodness judgments. Recall the hammer that was good as a bone in a sculpture, but bad at driving nails. Does this hammer have the property of goodness? It seems that it both does and it does not—which is clearly a contradiction. We thus have a *reductio* of the view that goodness is a non-supervening monadic property.

Perhaps this view could be salvaged if we were to claim that there are a variety of types of goodness properties, and the hammer has one of them, and lacks another one.

Thus there is no contradiction. This ‘solution,’ however, creates more problems than it solves. Suppose we have a hammer that is both good for driving nails and as a component in a sculpture. We would thus have a hammer with the following attributes: goodness₁, the ability to adequately drive nails, goodness₂, being well suited to serve as a component in a sculpture, etc. What is bizarre about this scenario is that the hammer’s having goodness₁ is not related in any relevant way to its ability to drive nails (or to its serving as a part of a sculpture for that matter). The same goes for the hammer’s having goodness₂. These are four unrelated properties that the hammer just so happens to instantiate. We might wonder, then: How many goodness properties does this particular hammer have? What constraints are there on how many goodness properties the hammers can have? Why do we call all these separate monadic properties ‘goodness properties’? Perhaps the strangest consequence of this view is how the goodness of this hammer relates to how I ought to respond to it. Suppose I want to make a birdhouse, and I learn that this hammer adequately drives nails. Do I also need to know that it is *good* before I buy it? It seems as if the goodness of the hammer is irrelevant to whether or not I should buy it, use it, favor it, etc. It would be mere coincidence if this hammer was both good and could adequately drive nails. Intuitively, however, I should buy, use, or favor good hammers (well, at least when I am interested in driving nails). I conclude that the view that non-moral goodness is a non-supervening monadic property is wildly implausible.

Several of the problems with the above view are fixed if we simply add the claim that goodness does supervene on the other features of good things. For instance, it might no longer be the case that there is no connection between a hammer’s ability to drive nails and its being good. If the hammer’s goodness is related in the appropriate way to its

ability to drive nails, then there would be an easy story to tell about why I should buy, use, favor, etc the hammer (given that I am looking to drive nails). However, I worry that this supervenience view is not really very different from an end-relational account of goodness, since the goodness of a thing will always supervene on the features that make it the case that the thing promotes some goal.

As we noted above, for almost any feature, F, some things will be good because they have F, others things would cease to be good if they were to have F. Thus, either goodness does not supervene on F alone (and thus, F needs to be combined with some other feature(s) for goodness to supervene on it), or goodness does supervene on F alone, but the supervenience can be prevented by pairing F with some other feature(s) of X. Consider, for example, the property of softness. Soft sandpaper is bad; soft silk is good. Thus, either goodness does not supervene on softness alone, or it does supervene on softness alone, but that supervenience can be prevented due to the presence of other attributes in the thing in question. Either way, what other attributes the thing has other than softness will bear on whether or not it is good.

Which other features of X bear on whether or not it is good? One possible answer is that the relevant feature is which kind X belongs to. Perhaps it is the case that if X has the attributes of being soft and being silk, then X is good. If it has the attributes of being soft and being sandpaper, then it is not good. This answer is mostly right, but again, it does not allow that things might be good even if they are not good members of their kind. A piece of sandpaper that is well-used, and thus soft, might no longer be good as *sandpaper*, but if I were interested in a surface on which to jot down a quick note, it would be good. This suggests that the further feature other than softness, on which X's

goodness supervenes, is whether or not X promotes a goal, G. I conclude that on a supervenience view, the relation of promoting a goal will always be present in the set of subvening attributes of an object.

Since the above argument is neutral with respect to what feature ‘F’ stands for, it seems to suggest that for *any* attribute, X’s having that attribute is not sufficient for X’s being good. This would thus cut against my thesis that goodness is the relation of promoting an end. The argument above suggests that the goodness of X would depend not just on its promoting a particular goal, but also on the other attributes it has. This suggests that goodness cannot be identified with the promotion of goals.

Yet the attribute of promoting goals is not an attribute like every other—if X promotes a goal, then there must be some other attributes of X in virtue of which it promotes that goal. For instance, if a piece of sandpaper would promote the goal of smoothing a piece of rough lumber—that means the sandpaper is rough, durable, etc. If a piece of silk had those features (being rough, durable, etc), it would also promote the goal of smoothing lumber. We can thus distinguish the relational property of promoting goal G, from all the further attributes that X has such that it has the relational property of promoting G. These further features make it the case that X promotes G. X’s goodness supervenes on its promoting G—and its promoting G supervenes on what other attributes the thing has. Softness promotes some goals but not others, so softness alone does not ground X’s goodness. It is the promoting of G that grounds X’s goodness.⁶⁴

I conclude that the most plausible supervenience view must agree with my end-relational account of goodness on the following claims: (i) there is a relation of

⁶⁴ Finlay also notes that X’s being able to promote a goal depends on what other features X has. He calls these additional features of the ‘good-making properties’ (2014) p.47.

promoting a goal, (ii) all good things stand in this relation to some goal, and (iii) the goodness of those things depends on their standing in that relation to a goal, as well as the presence of other features. Our accounts only differ in how we understand the relations among the goodness of X, the relation of promoting some goal, and the further features of X that contribute to X's goodness. My view says that if X has features such that it stands in the relation of promoting a goal, then X is good relative to that goal. Standing in that relation just is being good relative to that goal. The supervenience view says that if X has features such that it stands in the relation of promoting a goal, then X also has some further monadic property, goodness, that supervenes on those features and the relation of promoting a goal. My view does not require positing a further, monadic property of goodness. It thus posits fewer entities and should be favored to the supervenience view.

In this section, I have been evaluating the plausibility of views that claim that non-moral goodness is a monadic property. On the first such view, the property of goodness is not supervenient, and thus, is not related in any relevant ways to the other properties an object has. This means that it is possible for a hammer to have the ability to drive nails, and yet, still not be good. My interest in hammers, then, should not be directed by which ones are good, but by which ones adequately achieve the goals that I am interested in using hammers to accomplish. That is a counterintuitive result. The other view I considered says that non-moral goodness is a monadic property that supervenes on the other attributes that a thing has. I argued that the most plausible version of this view agrees with my view on many points, but that it needs to posit an additional property whose existence I deny. My view is thus simpler, and should be favored. Of course, if it turns out that we have other reasons for positing monadic goodness, then the fact that my

view is simpler than the supervenience view might no longer count against it. Yet so far we have not seen any reasons for positing monadic goodness.

ii) Non-moral goodness is not essentially end-relational

Let us move now to the other type of strategy someone might take if they wish to reject my end-relational analysis of non-moral goodness. Pursuing this strategy involves claiming that what is *really* essential to goodness is its relation to the acts of recommending, commending, or prescribing. There are two ways one might pursue this style of theory. First, one might simply deny that our uses of the word ‘good’ ever introduce a property at all. ‘Good’ does not ascribe a relation, nor does it ascribe a property. Instead, this objector might insist, our uses of ‘good’ express the speaker’s mental states. I am speaking, of course, of a view that is parallel to the family of views in meta-ethics referred to as ‘Non-Cognitivism.’ This view has struck many as a plausible account of our moral uses of ‘good,’ but I argue that it is an implausible view when applied to our non-moral uses of ‘good.’⁶⁵ Second, my opponent might allow that in non-moral uses ‘good’ does have some descriptive content (and that content might be about ends), but the essential feature of our use of the word ‘good’ is that we use it to recommend, commend, or prescribe. In *The Language of Morals*, R.M. Hare advances a view of this sort.⁶⁶ Against such views, I will argue that any connection between our judgments of goodness and prescriptions is not a necessary one. It is possible to judge something as good without commending it, thus we do not need to reference commendation in our account of the meaning of ‘good.’

⁶⁵ I agree with Finlay when he states that a non-cognitivist view of non-moral goodness would be ‘implausible’ (2014) p.19. In what follows, I hope to provide additional arguments in support of this claim.

⁶⁶ Hare (1952).

To start, let us look at what might be called a naïve ‘expressivist’ account of non-moral goodness. This view can be quickly rejected. According to a naïve expressivist view, when we say that something is ‘good’ all we are doing is expressing a pro-attitude for that thing. We are not asserting anything that can be evaluated as true or false—we are not, for instance, asserting that the thing we are calling ‘good’ has a property, relation, or set of properties and/or relations. Our judgments of the form, ‘X is good,’ are not descriptive of the world; they are purely *evaluative*. They serve the same function as pure expressions of approval such as, “Hooray X!”⁶⁷

This purely evaluative account of the meaning of ‘good’ is something that sounds somewhat plausible when discussing moral goodness. Speaking for myself, it is certainly true that when I say that something is morally good, that I have positive evaluative attitudes (so-called, ‘pro-attitudes’) towards the thing in question.⁶⁸ In telling you that charity is morally good, for instance, I express that I have the relevant pro-attitudes towards charity. Is it plausible that that is *all* I am doing when I judge that charity is morally good? Is it plausible that I do not intend to assert something true about the world? I have my doubts, but several philosophers have found this plausible. Let us set that issue aside for now. What I want to stress is that when an expressivist account of goodness is applied to our non-moral uses of ‘good,’ however, this view is quite simply implausible. Consider what the expressivist would have to say about the following story:

⁶⁷ See, for instance, A.J. Ayer (1936). Also, see Mark Schroeder’s recent book (2010) for a thorough and illuminating discussion of the various types of non-cognitivism.

⁶⁸ Is it possible to imagine someone who can make sincere moral judgments without the relevant pro-attitudes? That is, can there be so-called ‘amoralists’? Brink (1986). This questions have been hotly contested for years. I am not here taking a side on this issue—I am merely expressing that it seems true of me that when I make sincere moral judgments that I have the relevant attitudes.

Hardware Store. You have decided to build a birdhouse for your backyard. You have drawn up a plan, and amassed the appropriate materials (wood, nails, etc). However, it occurs to you that you cannot begin the project since you lack a crucial tool: a hammer. To remedy the situation, you head to the local hardware store: Jolene's Hardware. Jolene has been running the business for years, is extremely knowledgeable, and adores do-it-yourself projects. Upon entering the store you find yourself face to face, however, not with Jolene, but with her teenage son Junior. Lucky for you, though, Junior has worked in the store with his mother for years, and has learned the trade well. He is now truly an expert on do-it-yourself projects. Upon hearing about the nature of your project, Junior takes you to the same section of the store that Jolene would have, and hands you the same hammer Jolene would have, and declares (as Jolene would have): 'This is a good hammer.'

However, Junior does not have the same passion for hardware as his mother. In fact, he *hates* hardware, and sees do-it-yourself jobs as a waste of time. He would rather spend his time and energy studying philosophy. Junior is a naturally cheerful fellow, and nothing about the manner in which he utters the sentence indicates to you his utter disdain for everything having to do with hardware. I stipulate then that when Junior utters the sentence 'This is a good hammer' he does *not* express his approval of that hammer, and/or of hammering, and/or of do-it-yourself projects. He has no pro-attitudes for that hammer, and thus, does not express any pro-attitudes. Yet, by my lights, he judges it to be good. The expressivist would have to say that Junior's judgment was insincere—but this just does not seem right. Even if we accept that in calling a hammer 'good' Junior would typically be expressing a pro-attitude about that hammer (more on this claim later),

it is hard to imagine that that is *all* he is doing. It is hard to imagine that the further descriptive claim (that hammer will adequately drive nails in the construction of your birdhouse) is not an essential part of the judgment that the hammer is good.

R.M. Hare was well aware that our judgments of non-moral goodness contain descriptive content—but he denied that that was all there was to them. He argued that we could not define ‘good’ by way of descriptive content alone. He pointed out that the primary function of our judgments of goodness is to *commend*. He thus rejected any account of goodness that did not make reference to the role it plays in commending. Junior commends the hammer that he hands you in *Hardware Store*. He tells you that it is choice worthy—he suggests that you should choose it. A similar thing, says Hare, happens whenever we judge that something is good. He writes:

The meaning of ‘good motor-car’ is something that might be known by someone who did not know the criteria of its application; he would know, if someone said that a motor-car was a good one, that he was commending it; *and to know that, would be to know the meaning of the expression.*
(R.M. Hare, *The Language of Morals*, 1952, p.117. Emphasis my own).

So, according to Hare, we can imagine that you, as the uninformed customer, might not know the ‘criteria of application,’ for the term ‘good’ when applied to hammers. That is, you do not know which features are the goodness-making features of a hammer. However, when Junior hands you a hammer and says, ‘This hammer is good,’ you understand what Junior is saying. You understand since the primary meaning of ‘good’ is evaluative. To say of X that it is good is to commend X.

I believe that Hare is mistaken about the primary meaning of ‘good’ having to do with commendation, though it is not hard to see why one would be tempted to claim that judgments of goodness necessarily involve some kind of commendation, approval, or

other sorts of pro-attitudes. Suppose that Tim says of a certain kite that it is a good kite. Now Tim would have to be a very strange person indeed if he was offering up this judgment, unprompted, and yet, lacked pro-attitudes towards kites and/or kite flying. First, to judge a good kite from a bad one, Tim would need to have some knowledge about kite flying. It would take time and energy to learn what it is that makes the difference between good and bad kites. Consequently, people rarely come by such knowledge unless they have an interest in flying kites. When a person has knowledge in a subject area, this is often a sign that the person has (or had) interest in that subject area. Second, for Tim to offer his judgment without prompting suggests, at the very least, that he is interested in *talking* about kites. It would be very strange for a person that is truly disinterested in kites and kite flying to be nonetheless interested in talking about them. Finally, it would certainly be strange if Tim told us that although he judges this kite as good, he is in no way recommending to us that we should use that kite, or go buy a kite that is similar to it. Why would he even talk about the kite being good unless he wanted us to respond in certain ways to kites like it? It is for these reasons that we generally assume that a person who makes a judgment about the goodness of X has some pro-attitudes towards X. And generally speaking, this assumption is safe—it would be a rare occurrence for one to judge that X is good and to talk about X's goodness, if she did not have pro-attitudes towards, and want to commend X. Regardless, I contend that it is possible to judge the goodness of X without having pro-attitudes and without commending X.

For example, suppose my friend Jared is writing a murder mystery novel. For his novel to be a success, he needs the murderer to kill his victims in an unusual but effective

manner. Jared is suffering from a bit of writer's block and having trouble thinking of any creative methods of murder, so he asks his friends to help him brainstorm about good murder weapons and good methods for killing. I might, for his sake, start thinking about these topics, and come to make some judgments about what would be a good murder weapon. I might come to think: "Suppose the murderer was a scientist working on creating and programming airborne nano-bots. He could program an swarm of nano-bots to invade the victims' bodies through their mouth or nose, find their way into their victims' bloodstreams, scrape the cholesterol from their victims' arteries, and collect it into a clot in the victims' brains causing them to suffer strokes. That would be a good way to kill someone." I can form this judgment, I believe, without thinking that murder is a good thing, or without commending murder by nano-bot swarms. In fact, I could even deplore murder, and believe that no one should ever murder! I need not even have pro-attitudes towards murder mystery novels or the sorts of events depicted in such novels. As described, I have pro-attitudes towards helping my friend Jared write his novel, but those attitudes only serve to explain why I was engaged in thinking about murder weapons to begin with. Granted, I am commending the scenario as a good thing to write in his book. However, it seems to me that the reason that I commend the scenario is that I honestly judge that murder by nano-bot swarms would be a good way to kill someone. I see no reason why we should interpret me as either having pro-attitudes, or as commending, murder by nano-bots.

Hare would likely respond to this example by suggesting that if I do not have the relevant pro-attitudes or intend to commend murder by nano-bots, then my use of 'good' is not sincere. He would say that my use of 'good' is an example of what he called, the

‘inverted commas’ sense of ‘good.’⁶⁹ When we use the inverted commas sense of good, we signal that our use of ‘good’ is not entirely sincere. As Hare explains: “We are, in this use, not making a value-judgement ourselves, but alluding to the value-judgments of other people.”⁷⁰ We intend to express that X has the attributes that *others* say meet the standards for goodness for things like X, but in so doing, we do not intend to express our agreement with those standards. The inverted commas use of ‘good’ is often an ironic or sarcastic use of ‘good’ that includes a negative evaluative judgment. Since I do not seem to share the positive evaluative judgment when I claim that it would be good to murder someone with nano-bots, Hare might claim that this is not a true judgment of goodness, but merely an inverted comma use of the word ‘good.’

To use ‘good’ in the inverted commas sense, I need to be aware of the conventional standards of goodness by which others would judge things like X. How might I come to have knowledge of the conventional standards by which people would judge the goodness of murder weapons, for instance? I see two avenues. First, there is a history of murder full of people who were interested in murdering, and I am familiar with this history. I have seen what features people have looked for in murder weapons, and thus, knows what people who are interested in murdering will judge to be good. But what if there were no such history? Could one dispassionately judge something to be good without there being a history to draw on? Consider, for instance, what we would say about the following case: A group of friends are hanging out at the bar, and are involved in a heated discussion about what strategy is best for surviving a zombie

⁶⁹ *Ibid.*, p.124

⁷⁰ *Ibid.*

outbreak.⁷¹ Laura, a common friend of the discussants, shows up to the bar and wishes to join the conversation. Laura is not interested in zombies, movies about zombies, etc and would normally find such a conversation boring. However, she enjoys puzzles, and sees the set-up of the conversation as a kind of puzzle: How would one plan to survive in a world populated with zombies? She thus joins in on the conversation and declares: “A samurai sword would be a good weapon for that scenario.”

There is no history of fighting hordes of zombies, as zombies are fictional creatures. There is thus, no convention about what would be good to fight zombies with that Laura appeals to in making her judgment. Laura’s judgment that a samurai sword would be good for fighting zombies does not rely on or presuppose her knowing any conventional standard of goodness for zombie fighting. One might insist that there is a history of zombie fighting, though it is a *fictional* history portrayed in movies and books—and that Laura could be drawing on this history in making her judgment. Let me stipulate then that she is not. Imagine she has never seen a zombie movie. She simply hears her friends describe how zombies attack, and what their weaknesses are, and she concludes that a samurai sword would be a good weapon for fighting them. If you doubt this is possible, think back to the very first movie or book featuring zombies. Before the author(s) of that fiction created zombies, there was no history to draw on. The author(s) had to imagine ways the heroes of the stories would escape and or fight off this new kind of creature. They made judgments about what would weapons would be good and bad for fighting zombies. They made these judgments, not by appealing to a history of what

⁷¹ This is a rare use of ‘zombie’ in philosophical writing. My interest here is reanimated, or walking dead, humans—and not, so-called ‘philosophical zombies’.

people chose—but by imagining an end (defeating zombies in a fight) and then figuring out what weapons would advance that end.

This has been a long-winded way of suggesting that even when one judges an X to be good by appealing to a history full of people choosing X, the explanation for those peoples' choices is ultimately grounded in the fact that X advanced some relevant goal. So we really do not need to reference the people or their choices in explaining why X is good, or why I currently judge X to be good. We can simply appeal to the goal and my judgment that X advances that goal. That is what Laura is doing in the zombie-killing conversation, that is what I am doing in the nano-bot murder conversation.

There is one last move for the proponent of a commending account of goodness to make. She might insist that since I do not endorse the relevant standards, when I judge that a swarm of nano-bots would be a good way to murder someone, I am not sincerely and literally using 'good'. The nano-bot example is my own invention, so I could of course just stipulate that I was in fact being sincere in the example. I doubt my opponent would allow me this stipulation; she would insist that it is impossible for me to judge that X is good without the relevant evaluative attitudes. As a reply, let me try an example that comes from my own, actual experience. That way, I can speak to how sincerely the judgment was uttered.

To begin, I need to admit to watching a lot of bad television shows and movies. One of my favorite bad movies, *The Room*, is truly a paradigm example of how not to make a movie. The writing, acting, directing, set-design, casting, music, story, and pacing are all just terrible.⁷² It is a really, really bad movie. I have, nonetheless, uttered the

⁷² I suspect that a lot of this can be explained by the fact that one man, Tommy Wiseau, wrote, directed, produced, and starred in the film.

following judgment quite sincerely: “I know that *The Room* is not a good movie, but you should really watch it.” If this were an inverted commas use of ‘good,’ the first clause would allude to standards of goodness that I do not endorse. But I can assure you that I do in fact endorse the standards of goodness for movies by which *The Room* is truly judged a very bad movie. Now presumably, if a sincere use of ‘good’ expresses pro-attitudes, then the first clause of this sentence (when spoken sincerely) expresses a lack of pro-attitudes. It should consequently sound contradictory (or at the very least, incongruous) for me to go on in the next clause and commend the film to you (since, on the view we are considering, sincere commendation requires the relevant pro-attitudes). This sentence is thus very puzzling on a commending account of goodness. However, intuitively, it is not a puzzling sentence at all.

The sentence is not puzzling if we think that what Hare calls “the evaluative meaning of ‘good’” is not essential.⁷³ For instance, if we believe that ‘good’ ascribes end-relational goodness, we can make sense of this sentence quite easily. The first claim—that *The Room* is a bad movie—is sincerely uttered because I believe that there are a variety of goals that set the standards for what a good movie should be like, and *The Room* meets few, *if any*, of those goals. I can endorse those standards, without having a corresponding negative attitude about *The Room*. I can like things that I know to be bad. Thus, it is mistaken to think that the correct definition of ‘good,’ in non-moral uses, is essentially about expressing pro-attitudes.

b) Objection: Thesis of Commonality is false

Suppose instead, that my opponent wished to defend the thesis that moral goodness is monadic by way of (b), that is, by rejecting the Thesis of Commonality.

⁷³ Hare (1952) p.121

Without the Thesis of Commonality, there is no problem with saying that non-moral goodness (like the goodness of hammers) is relational, while moral goodness is monadic. At first glance, this sounds like an open and attractive way to argue; yet rejecting the Thesis of Commonality comes at a great cost.

To see this cost, consider a silly example. Suppose I were to pick an arbitrary monadic property (say, the property of redness for instance) and an arbitrary relational property (say, being the brother of Brianna). I could decide to use the same adjective, 'red' to refer to both of these things. Now for the most part, my use of 'red' would coincide with the regular use of that term in English. I would agree that apples, fire trucks, and roses are red. However, while other English speakers would claim that 'red' picks out a monadic property, I would disagree. When asked to defend my view, I would point out that 'red' sometimes refers to the relational property of being the brother of Brianna. My detractors could argue: "Look, Brian. You clearly have the concept of redness and the concept of being the brother of Brianna. For some reason, you have picked out these two distinct attributes with the same predicate, but you could just as well have used two predicates (and things would have been much less confusing)". "Yes, yes," I might reply, "but there is also, it strikes me, a *further* property that is the disjunction of either the property of redness or the relational property of being the brother of Brianna. I have been calling that disjunctive property 'red', but I could very well just call it 'Blarb' to avoid confusion." "Do what you like. You can arbitrarily gerrymander disjunctive properties however you see fit, but just know that the rest of us aren't interested in the property you are calling 'Blarb', we are interested in redness."⁷⁴

⁷⁴ Frank Jackson (2000) proposes that one could simply identify a moral property (like goodness) with the massively disjunctive property of being either G1 or G2 or G3... where those variables

The accusation that the disjunctive property picked out by the predicate ‘Blarb’ is the result of arbitrary gerrymandering is apt, if not only because I set up the example that way (remember, I arbitrarily selected a property and relation for the example). As I understand it, the philosopher that pursues (b) is also introducing a disjunctive property. She is calling ‘good’ anything that has either the monadic property of moral goodness or (at least) one of the relational properties of non-moral goodness. We should ask, then, whether this disjunctive property is also the result of an arbitrary gerrymandering. If it is, I see no reason to continue using a single predicate, ‘good’ to pick out the various properties and relations that have been gerrymandered together. We should be just as comfortable picking out new predicates for non-moral goodness, moral goodness, and for the disjunction of moral and non-moral goodness, as we are for redness, being the brother of Smith, and the disjunction of redness and being the brother of Brianna.

If it strikes you, as it does me, that these two cases are distinct—that it makes sense to use a variety of predicates in the redness/brother of Brianna case, but not in the goodness case—then you agree with me that the predicate ‘good’ does not pick out an arbitrarily gerrymandered disjunctive property. And if ‘good’ does not pick out an arbitrarily gerrymandered disjunctive property, there must be something in common among the things we call ‘good.’ That is, the Thesis of Commonality holds for goodness.

c) Objection: The common feature of good things is not that they are relational

Thus (c) is the only remaining option for the philosopher who wishes to deny that goodness is relational. That is, such a philosopher must claim that though non-moral goodness is relational and moral goodness is monadic, there is still something that these

refer to the members of the set of all properties we would intuitively view as ‘goodness making’. This view is not illuminating; it tells us nothing of the nature of goodness.

attributes have in common such that it makes sense to call them both types of goodness. They thus agree with me that the Thesis of Commonality is true—they simply disagree with me about which feature of good things it is that they all share in common. One such proposal, offered by Michael J. Zimmerman, is that what unites all good things is the *appropriateness of favoring them*.⁷⁵

I am confident in describing Zimmerman as a pursuer of strategy (c). He argues that moral goodness is monadic. But he also says, after considering the claim that moral goodness is just one of many ways of being good: “But what, then, might be said to unite the ways of being good, making them ways of *being good*...? My answer is: the idea of *valuableness* (that is, the worthiness of being valued).”⁷⁶ I take it that the question—what unites the ways of being good?—is asking for some way to satisfy the Thesis of Commonality. If ‘good’ just means ‘worthy of being valued,’ then it is certainly possible that ‘good’ sometimes picks out a monadic property, and other times a relational property.

According to Zimmerman, to say that it is appropriate to favor something is to say that certain favoring attitudes are fitting, or even required.⁷⁷ That is, it is appropriate or fitting to have certain favoring attitudes if there are reasons that you ought to have those attitudes. In the roughest terms, then, Zimmerman explains goodness in terms of oughtness. Good things are those things that you ought to favor.

⁷⁵ Wherever possible, I will do my best to present Zimmerman’s analysis as accurately as I can without getting too bogged down in the somewhat excruciating details. Zimmerman takes almost an entire book to clarify what his analysis amounts to.

⁷⁶ Zimmerman, p.27

⁷⁷ According to Zimmerman’s final analysis of moral goodness, we are morally required to favor it to a certain degree. *Ibid*, p.122

This view is attractive for several reasons. First, it seems correct to say of most good things that we ought to favor them. On Zimmerman's account of goodness, this connection is no coincidence. What it means for something to be good is for it to be the case that we ought to favor it. Second, thinking of goodness in this way gives us a very natural way to distinguish between moral and non-moral goodness. Suppose a hammer is a good hammer. According to Zimmerman, this means I have reasons to favor the hammer. Of course, if I am not interested in hammering, then these reasons do not have to factor into my practical deliberation. The reasons that count towards favoring non-moral goods will be *hypothetical*—their normative force depends, in some important way, on the agent's desires. But, says Zimmerman, we have moral reasons to favor things that are morally good. Presumably, part of what makes a reason a *moral* reason is that its normative force is independent from the agent's desires. That is, we have *categorical* reasons to favor things that are morally good. Again, by distinguishing these two types of reason, we can distinguish moral goodness from non-moral goodness while still insisting that what makes them both types of goodness is that it we ought to favor them.

Despite these virtues of the view, Zimmerman's account of goodness strikes me as seriously mistaken. While his account of goodness explains the tight connection between the goodness of things and the appropriateness of valuing those things, it gets the explanation backwards. In the style of the Euthyphro Dilemma, we can ask: Ought we favor good things because they are good, or are they good because we ought favor them? It is hard to do more here than appeal to intuitions. If I tell my friend he ought to see it, and he asks me why, a sensible and natural response is 'You ought to see it because it is good'. If, on the other hand, I tell my friend that a movie is good, and he

asks me why, it is not a very natural response to say ‘It is good because you ought to see it.’ Intuitively then, it is a thing’s goodness that explains why we should favor it (or perhaps, the fact that something is good merely signals that there are other features of the thing that explain why we ought to favor it).⁷⁸ In either case, that a thing is good *justifies* our favoring it.

Zimmerman predicts and replies to this objection, by way of clever observation about how the term ‘because’ operates. As I am interpreting things, the fact that something is good *explains* or *justifies* its being the case that I ought to favor it. I ought to favor X *because* it is good. On this view, the goodness of X is prior to what I ought to do, so we cannot analyze goodness in terms of what we ought to favor. However, says Zimmerman, the explanatory / justificatory use of ‘because’ is not the only such use of the term. Consider the following example. We know that every point of a circle is equidistant from a fixed, central point. If we were asked why this is the case, we would reply: “Every point on a circle is equidistant from a fixed point *because* that is what it is to be a circle.” Let us call this the *constitutive* use of ‘because’. Zimmerman asks: “Might it not be that it is *this* sense of ‘because’ that makes ‘It is appropriate to favor S because it is good’ sound ‘natural’? If so, the suggested analysis is not threatened after all”.⁷⁹

Now of course, it *might* be the case that we use ‘because’ in its *constitutive* use when we say that a person ought to favor good things because they are good. Zimmerman himself goes on to express doubt that this is in fact how we use ‘because’ in these cases. I know of no sure way to settle the question; the best I offer is brute appeal to my sense of

⁷⁸ I am here describing T.M. Scanlon’s ‘buck-passing’ account of value (1998) pp.96-8.

⁷⁹ Zimmerman, (2001) p.115

how I use the term. When I say that we ought to favor good things because they are good, it seems to me that I am using ‘because’ in the explanatory / justificatory way. Similarly, when a friend asks why I think she ought to go see a movie, and I reply: “Because it is good” I seem to be justifying my answer (albeit, in a not very descriptive manner) rather than repeating my recommendation in a different way.

I reject Zimmerman’s account of goodness because it gets backwards the explanation of how goodness relates to what we ought to do. This is, unfortunately, far from showing why any attempt at strategy (c) is bound to fail—although it is suggestive. Recall that to satisfy the Thesis of Commonality we need to point to some feature that all good things have in common in virtue of which it is appropriate to call them ‘good.’ Any such feature will have to be quite general and abstract, and at least be part of an analysis of what it is to be good. Zimmerman’s account is paradigmatic of a wide range of theories that suggest that goodness be analyzed in terms of reasons, fitting reactions, or what we ought to do. Such theories would succumb to the same objection—that the goodness of things *explains* why it is appropriate to respond to them in the specified ways. So while any such theory is compatible with moral goodness being monadic, and also with there being something in common between moral goodness and non-moral goodness, such theories are objectionable.

For what it is worth, my end-relational account of goodness also accounts for the intuition that goodness is closely related to what we ought to do. After all, on my view, if X is good, then X adequately promotes a goal G. Now if a person were to want to achieve G, it seems to follow that she would have reasons to favor X. So if I tell you that X is good, I tell you that *if you were interested in achieving G, you ought to favor X*. Notice

that this would be a form of commendation—but it is merely hypothetical commendation. I could always believe that no one is interested in achieving G—or stronger still, that no one *should* be interested in achieving G. I can, nonetheless, judge which things would be good relative to G. Suppose, for instance, that you want to build a solid piece of furniture. You ought to build it out of hardwoods (not softwoods), because hardwoods are good to use in building sturdy pieces of furniture. I tell you that maple is a good wood for building furniture. It seems that I have signaled to you that there is reason for you to use maple in your project, and I have thus commended maple. So on my view, goodness is still closely related to both commendation and to what we ought to do. Better yet, however, my view also explains why the goodness of things *explains* why we ought to favor them.

IV. A noteworthy virtue of my end-relational analysis

If the arguments of this chapter are sound, then goodness is a relation that holds between things and ends. Though I do not wish to make much of the point, it is worth noting that one of the theoretical virtues of an account of goodness wherein all goodness falls under one unified analysis is that the resulting account of goodness is highly parsimonious.

I have two reasons for not making much of the advantage that my account enjoys due to parsimony. First, I wish to prevent a certain misunderstanding of my argument. The motivation behind my appeal to the Thesis of Commonality is *not* parsimony; the motivation behind my appealing to the Thesis of Commonality is the accommodation of a common sense intuition. This is as it should be, since common sense intuitions are a good place to begin philosophical theorizing while parsimony is not. If parsimony were a

starting point for philosophical theorizing, things would be too easy; parsimony comes very cheaply. I could construct a theory, for instance, wherein I claim that all our evaluative terms, ‘ought,’ ‘good,’ ‘beautiful,’ ‘just,’ ‘right,’ etc. pick out one and the same property. This would certainly be a parsimonious theory; but no such theory is acceptable to our common sense intuitions. Intuitively, there is a difference between calling something ‘just’ and calling it ‘beautiful.’ We can reject such a simplistic theory forthrightly. Perhaps some argument could be given that would convince us that such a theory is correct despite conflicting with our common sense intuitions, but a mere appeal to parsimony will not do.

This leads me to my second reason for not putting much emphasis on the parsimony of my account: as the above discussion suggests, parsimony is a fairly weak virtue. Parsimony can tip the scales when we are deciding between two or more competing theories that strike us as equally plausible.⁸⁰ What parsimony does not do, however, is make a weaker, less plausible theory trump a stronger, more plausible theory.⁸¹ This discussion of the virtue of parsimony is my way of acknowledging the work I have ahead of me. I have shown that my account of goodness is fairly parsimonious, and accommodates some of our commonsense intuitions. At best, I have shown that my account of goodness cannot be dismissed outright. To show that we ought to adopt my account of goodness, however, I need to show that it is more plausible—or at

⁸⁰ As it did above when I discussed the view that non-moral goodness is a monadic property that supervenes on the other attributes of good things.

⁸¹ This might be putting the point too strongly. Suppose we are evaluating theories T1 and T2. Perhaps T1 is only slightly less plausible, but much more parsimonious, than T2. In this situation, I am willing to concede that it is at least possible that we ought to favor T1 to T2. The point I am trying to make, however, is that the two (or more) competing theories need to be very nearly equally plausible before parsimony matters.

least nearly as plausible—as less parsimonious theories. Arguing for this point is the primary task of the remaining chapters.

V. Conclusion

I have here argued that all goodness is relational, and we thus ought to reject the orthodox view that moral goodness is a monadic property. Were we to adhere to the orthodox view, we would be faced with several unattractive alternatives. Two of those alternatives: (a) rejecting the end-relational analysis of *non-moral* goodness, and (b) giving up the Thesis of Commonality, are implausible. It is tempting to try the third alternative, (c) giving an account of goodness that accommodates the Thesis of Commonality, while still insisting that *at least* moral goodness is monadic. My suspicion is that the most plausible accounts of goodness that can be utilized for strategy (c) will analyze goodness in terms of what we ought to do, what we ought to favor, or what would be appropriate to respond or the like. But such approaches require abandoning the claim that the goodness of things *justifies* or *explains* why we ought to do things. Thus, I conclude, there is no plausible account on which moral goodness is a monadic property.

Upon reflection, it actually seems a bit odd that philosophers would take sides on whether goodness is monadic or not unless something else rides on how we answer this question. Geach and some of his followers believe that certain meta-ethical and normative positions can be rejected if goodness is not a monadic property; I agree (I will discuss some of these theoretical consequences of my view in a Chapter Five). Consequently, one motivation for taking a side on the nature of goodness is an interest in arguing against, or defending, some of these meta-ethical and normative positions. Most philosophers, I suspect, defend the claim that moral goodness is monadic due to their

belief in how that claim connects with the claim that some things are intrinsically good. I believe it is really this second claim, about intrinsic goodness, that is of most interest to moral philosophers. I turn now to discussing why that claim is false.

Chapter Three – Why we do not need intrinsic goodness

I. Introduction

In the last chapter I argued that goodness is end-relational. That is, a thing is good just in case and because it stands in the relevant relation to a goal. Briefly, my argument began with the premise that the best account of non-moral goodness reveals that it is end-relational. The second premise was what I called the ‘Thesis of Commonality’—that is, that moral goodness and non-moral goodness share something in common in virtue of which it is true that they are both types of goodness. My suggestion was that because our best analysis of non-moral goodness reveals that it is end-relational, then the Thesis of Commonality recommends that we should likewise conclude that *moral* goodness is end-relational.

I believe that absent good reasons to reject it, we ought to accept the Thesis of Commonality. In this chapter, I will entertain one line of reasoning that, if successful, would justify abandoning the Thesis of Commonality. More specifically, I concede that if affirming the Thesis of Commonality leads us to a theory of goodness that is too revisionary—one that conflicts with too many of our commonsense intuitions and expectations about what a theory of goodness does or should do—then perhaps we ought to reject the Thesis of Commonality. For instance, if the extension of an end-relational account of goodness into the moral realm has too many counterintuitive consequences, perhaps we should conclude that although non-moral goodness is end-relational, moral goodness is not. This is to say that for my master argument from the last chapter to

succeed, I need to show that adopting an end-relational theory of moral goodness would not do too much damage to our ordinary understanding of morality.

At this point, I must confess that the end-relational theory of goodness that I prefer does have one consequence that will strike many as quite revisionary: according to my view, nothing is intrinsically good. For many, this consequence itself constitutes a *reductio ad absurdum* of my end-relational theory of goodness. In this chapter, however, I will show that even though my view entails that nothing is intrinsically good, the view is not too revisionary. My argument involves examining what roles intrinsic goodness plays in our theoretical and practical lives, and then showing that an end-relational theory can fill those roles without positing intrinsic goodness; we do not *need* intrinsic goodness.

II. Preliminaries

Before looking at any arguments, it is best to be absolutely clear about (a) how I understand the notion of intrinsic goodness, and (b) why my view entails that nothing is intrinsically good. Let me start with (a). I intend to be discussing the same notion of intrinsic goodness that has been discussed for centuries by moral philosophers. Traditionally, moral philosophers have used the label ‘intrinsic goodness’ when talking about the value a thing has when it is ‘good in itself,’ ‘good for its own sake,’ or ‘just plain good.’ For instance, John Stuart Mill says that happiness is the only thing that is valuable for its own sake.⁸² Immanuel Kant says that the only thing that is good in itself is a good will.⁸³ Aristotle says that eudemonia is the “chief good” that “we choose always

⁸² Mill, John Stuart. *Utilitarianism*, From *On Liberty and Other Essays*, Edited by John Gray, (1991) Oxford University Press, p.137

⁸³ Kant, Immanuel. *Groundwork for the Metaphysics of Morals*, 4:393

for itself and never for the sake of something else”.⁸⁴ I take all of them to be making substantive claims about what things are intrinsically good.

Christine Korsgaard notes that labeling this kind of goodness ‘intrinsic goodness’ is misleading.⁸⁵ The label suggests that things are good for their own sake due only to their intrinsic properties. However, some philosophers have argued—quite compellingly—that there are things that are worthy of being valued for their own sakes at least in part due to their non-intrinsic properties.⁸⁶ For instance, Picasso’s *Guernica* is good for its own sake—but at least part of the reason this painting is valuable is due to who painted it, its cultural and historical significance, and the event it depicts.⁸⁷ Those are relational properties of the painting. If this is correct, then the intrinsic value of the painting is at least partially determined by its non-intrinsic properties. For this reason, Korsgaard suggests that we refer to goodness for its own sake with a new label: ‘final goodness.’⁸⁸ I am friendly to the suggestion, but am opting to stick with the longstanding tradition of referring to things that are good for their own sake as ‘intrinsic goods’.⁸⁹

Turning now to issue (b): Why does my view entail that nothing is intrinsically good? As I understand it, for X to be intrinsically good is for X to be good for its own sake. But on my view, nothing is good for its own sake; things are good only because

⁸⁴ Aristotle. *Nicomachean Ethics*, EN I.2.1097a27-1097b

⁸⁵ Korsgaard (1983) p.170

⁸⁶ See Shelley Kagan (1998) and W. Rabinowicz & T. Rønnow-Rasmussen (2000, 2003).

⁸⁷ Picasso’s *Guernica* depicts a truly despicable event of the Spanish Civil War. It depicts when, on April 26th 1937, German and Italian warplanes bombed the civilian population of the Spanish town of Guernica (at a busy outdoor market) at the request of then aspiring Spanish dictator Francisco Franco.

⁸⁸ Korsgaard (1983) p.170

⁸⁹ I am not alone in this decision. A recently published volume of essays titled *Recent Work on Intrinsic Value*, Edited by T. Rønnow-Rasmussen and Michael J. Zimmerman (2005), features 32 essays, many of which discuss the notion of value that Korsgaard calls ‘final goodness,’ under the label ‘intrinsic value’.

they stand in the relevant relation to a goal. They are good for the sake of promoting those goals. Thus, on my view nothing is intrinsically good.

Finally, let me introduce a bit of terminology for the ease of discussion. For the rest of this chapter, I will be using the term ‘nihilism’ to refer to the view that nothing is intrinsically good, and ‘realism’ to refer to the view that there is at least one thing that is intrinsically good. I will use the terms ‘nihilist’ and ‘realist’ to refer to the proponents of these views. I wish to stress, however, that these terms should only be understood quite narrowly. The nihilist whom I am considering in this chapter does not deny the existence of other kinds of goodness or value.⁹⁰ She simply believes that there is nothing that is intrinsically good. Further, the realist that I am concerned need not be a ‘realist’ in the sense of affirming the existence of Platonic universals. She can accept any metaphysical view about the nature of properties that allows for it to be the case that at least one thing is intrinsically good.

The realist view is the orthodox one in philosophy. As such, my end-relational view is revisionary. The question of this chapter, however, is whether my view is *too* revisionary. Does my view conflict with too many of our commonsense intuitions about moral goodness? To answer this question, I will start with an account of what roles or functions intrinsic goodness is believed to play in our practical lives. I will focus on two: (i) It explains why some of our judgments about the goodness of things are true, and (ii) it grounds basic practical reasons. Clearly, these are important aspects of our practical lives. Were it the case that we needed to posit intrinsic goodness in order to accomplish either (i) or (ii), I would concede to my opponent that this would constitute a fatal shortcoming

⁹⁰ Unless, as I go on to argue in Section III, her nihilism about intrinsic goodness commits her to nihilism about some other kinds of goodness.

for nihilism. I will argue, however, that we do not need to posit intrinsic goodness to accomplish (i) or (ii). My end-relational theory is not so revisionary that it requires giving up on basic practical reasons or the truth of (most) of our judgments about goodness.

III. Explaining the Truth of Goodness Judgments

The first role that intrinsic goodness plays in our practical lives is explaining why some of our judgments about the goodness of things (henceforth, ‘goodness judgments’) are true.⁹¹ Let us begin by dividing our goodness judgments into two categories. Some of our goodness judgments purportedly ascribe *intrinsic* goodness to the things we are judging; other goodness judgments ascribe *non-intrinsic* goodness. The truth of our *intrinsic goodness* judgments requires realism about intrinsic goodness, so the nihilist must deny that these judgments are true. This might seem like a big bullet for the nihilist to bite. After all, as I just mentioned, the history of moral philosophy has seen many very smart people arguing over which intrinsic goodness judgments are true. I will argue that a nihilist can give a plausible explanation for why so many people mistakenly believed that there are intrinsic goodness judgments that are true. This explanation will also show that although nihilism requires denying that any intrinsic goodness judgments are true, nihilism is nonetheless, not implausibly revisionary in this respect.

It is with regard to the truth of our *non-intrinsic goodness* judgments that the realist puts some real pressure on the nihilist. There is a promising line of argument—one that dates back perhaps to Aristotle, and is ubiquitous in ethics textbooks—that aims to

⁹¹ Recall that in the last chapter I argued against those non-cognitivist views of non-moral goodness that would deny that goodness judgments have truth-values. As such, I am assuming that at least some of our goodness judgments are true. Some non-cognitivists (Simon Blackburn (1993), for instance), can accept that some of our goodness judgments are true because they are deflationists about truth. Nothing I say here is incompatible with a deflationist view of truth.

show that we need there to be something that is intrinsically good in order to explain the truth of our non-intrinsic goodness judgments. I call this argument the ‘Argument from Need,’ or ‘AFN.’ Ambitiously, the realist might employ AFN in support of the *Strong Realist Claim*, the claim that the truth of *all*, or perhaps just *most*, of our non-intrinsic goodness judgments requires the existence of intrinsic goodness. This would certainly be a big bullet for the nihilist to bite! If the Strong Realist Claim were true, then any form of nihilism would be much too revisionary to be a plausible account of goodness. I will show, however, that the Argument from Need is deeply flawed, and does not support the Strong Realist Claim.

My objections leave open a less ambitious position to which the realist might retreat. She might still argue in favor of the *Weak Realist Claim*, the claim that the truth of *some* of our goodness judgments requires the existence of intrinsic goodness. Depending on the centrality of these judgments to our practical thinking, and our confidence that they are true, even this weaker realist claim has the potential to show that nihilism is too revisionary. I will argue that even this Weak Realist Claim is false.

a) The truth of intrinsic goodness judgments

The realist would say that the following sentences express paradigmatic examples of intrinsic goodness judgments:

- (A) Friendship is good.
- (B) Pleasure is good for its own sake.
- (C) A good will is good in itself.
- (D) Knowledge is intrinsically good.

The realist insists that at least one of the above sentences (or one very much like them) is true. The realist might claim that the sentence is true in virtue of expressing a true judgment. The corresponding goodness judgment is true when the thing that is referred to

by the subject term of the sentence has the property of intrinsic goodness. Were there no such thing as intrinsic goodness, these judgments and sentences would not be true. These sentences and judgments are true, says the realist, thus there must be something that is intrinsically good.

Given that the nihilist believes that no intrinsic goodness judgments are true, what should she say about sentences (A) – (D), and other sentences like them? Some nihilists have suggested that these sentences are meaningless.⁹² They are meaningless because they ascribe intrinsic goodness to things, when there is no such property. Such a view is hard to square with the prevalence of intrinsic goodness judgments. As mentioned above, the history of moral philosophy is full of very smart people arguing over the truth of sentences like (A) – (D). Philosophers have, for centuries, debated which things are intrinsically good. William Frankena compiled a list of all the things that philosophers have judged to be intrinsically good, and the list is daunting. He cites:

Life, consciousness, and activity; health and strength; pleasures and satisfactions of all or certain kinds; happiness, beatitude, contentment, etc.; truth; knowledge and true opinions of various kinds, understanding, wisdom; beauty, harmony, proportion in objects contemplated; aesthetic experience; morally good dispositions or virtues; mutual affection, love, friendship, cooperation; just distribution of goods and evils; harmony and proportion in one's own life; power and experiences of achievement; self-expression; freedom; peace, security; adventure and novelty; and good reputation, honor, esteem, etc (William Frankena, *Ethics*, 1973, pp.87-8).

And it is not just philosophers who trade in claims about which things are intrinsically good. As defenders of intrinsic goodness like to point out, millions of people have accepted the claim, appearing in the book of Genesis, that God looked at the world after

⁹² Judith Jarvis Thomson (2008) p.17 and Peter Geach (1956) p.41 defend this claim.

creating it, and “saw that it was good.”⁹³ How could so many people not see that what they were saying was incoherent or meaningless? Are we to assume that they were not competent users of the concept of goodness? That is a hard view to accept.

A more plausible thing for a nihilist to say about (A) – (D) (and other such sentences) is that all of them are technically false when we understand them as expressing judgments that ascribe intrinsic goodness. On the other hand, some of these sentences can be true when we understand them as expressing judgments that ascribe some form of *non-intrinsic* goodness. An end-relational theory of goodness, for instance, can explain the truth of these sentences quite nicely while at the same time explaining why so many people have mistakenly believed these sentences to be about intrinsic goodness. What do I have in mind?

The idea is that some sentences like (A) – (D) actually ascribe end-relational goodness to things even though the surface grammar of the sentence makes it appear as if they do not. Thus, when people take these sentences as expressing intrinsic goodness judgments, they are mistaken. The reason that the sentences do not mention ends or goals explicitly is because the relevant goal is taken as a given in the context, and conversational maxims prescribe against explicitly stating information that everyone takes as a given.⁹⁴ When we go see a financial adviser, she need not continue to dot her recommendations with the qualifier: ‘If you wish to be financially stable...’ This goal is

⁹³ Michael J. Zimmerman (2001) p.29, for instance, references the ‘pithy comment’ by Panayot Butchavarov: “Nevertheless, millions have thought they understood Genesis 1:31: ‘And God saw every thing he had made, and behold, it was very good.’” Butchavarov (1989) p.17.

⁹⁴ Finlay discusses this point with respect to moral judgments in general, that is, not merely with respect to goodness judgments (2008) p.353.

taken as a given, and thus does not need to be restated explicitly.⁹⁵ I believe a similar thing occurs in true sentences that appear to ascribe intrinsic goodness to things.

If I am correct, then the mistake that people make when they judge these sentences as if they are ascribing intrinsic goodness is the same kind of mistake people make when they think that sentences about the weight of objects ascribe an intrinsic property to objects. An object's *mass* is intrinsic, but its weight is a measurement of that mass being subjected to some magnitude of gravitational force. Weight is a relational property. When you change the gravitational force, you change the object's weight. However, because it is extremely rare that we are interested in gravitational contexts other than that on the Earth's surface, it is often unnecessary to explicitly qualify our sentences about the weight of objects with 'in an environment of 1 G of gravitational force'. So although the sentence, 'My cat weighs 14 pounds' does not appear to introduce a property that is relative to gravitational context, it does. Similarly, when a sentence of the form 'Friendship is good' is true, it introduces a property that is relative to a goal, even though it does not appear to.

To demonstrate how this would work, let us look at an example. For instance, let's return to the example of the creation story in Genesis. As you recall, God creates the Earth, looks upon it, then pronounces: "It is good."⁹⁶ If we only focus on the surface structure of God's utterance, it appears to be a judgment that ascribes intrinsic goodness to the Earth. However, I believe that the surface structure of that sentence disguises the fact that the judgment that God makes is one that ascribes end-relational goodness. And

⁹⁵ Finlay uses the example of a rugby captain, speaking about which plays the team ought to run, qualifying each with 'In order to win...' (2008) p.353.

⁹⁶ To be more precise, God "sees that [the Earth] is good." For the ease of discussion, I will pretend that God went on to vocalize this judgment.

while I do not wish to presume too much about God's motivations for creating the world, the following seems like a plausible story: God was interested in creating a habitable and pleasant place for human beings (and other animals) to live happy and fulfilling lives.

Thus, when he looks upon the Earth and judges it to be good, he is judging it as satisfying this goal of creating a habitable world. He judges the Earth as end-rationally good relative to this goal. To further support this interpretation, we should note that the Earth would not be end-rationally good relative to just any goal. Imagine that God sets out to make a harsh, ugly, and inhospitable planet on which to imprison and punish the worst sinners. He then creates the Earth as we know it. Would he still look upon it and judge it to be good? Or might he instead think: "This would be a good world for humans beings to inhabit, but it is a lousy world to use to punish sinners"? I simply cannot imagine God judging the Earth to be good in this scenario. I have a similarly hard time imagining God creating the Earth with no purpose in mind at all, looking upon the planet, and then judging it to be good. But if God did have a purpose in mind for the Earth, his judgment very well could have been one that ascribes end-relational goodness.

A similar story can be given for sentences (A) – (D). At first, it might be strange to think of friendship as end-rationally good—it just seems good on its own. The reason that 'Friendship is good' does not seem to us to ascribe end-relational goodness is due to the fact that the goal that friendship promotes is taken for a given. Friendship, it might be said, contributes to having a fulfilling and meaningful life.⁹⁷ This goal is so pervasive we can generally assume that speakers and audiences alike share this goal in just about every context. It does not need to be made explicit. But observe what happens if we construct a

⁹⁷ Finlay agrees that friendship is end-rationally good for living a fulfilling life, although he includes 'happy' where I include 'meaningful' (2014) p.34.

scenario where this goal is not contextually relevant. Suppose, for instance, we have captured a terrible warlord whom has brutalized and terrorized thousands of people. The international courts find him guilty of war crimes—but do not plan to execute him. They decide a more fitting punishment is that he lives out his life in a tiny cell in one of the many prisons where he and his followers once locked away political dissenters. One day, a few years into the warlord’s sentence, a guard intercepts a note that this warlord was passing to another prisoner in a nearby cell. The note is friendly, and reveals that the warlord and criminal have developed a friendship. The guard takes it to the warden, who grumbles, “A friendship!? Oh no, no, no! Friendship is a very bad thing!” and immediately begins plans to squash this friendship by moving the two prisoners further apart. Friendship generally promotes the goal of living a fulfilling and meaningful life. The warden does not want the warlord to have such a life. He wants the warlord to be punished and to suffer. Thus, in this context, friendship is a bad thing.⁹⁸ I suspect that sentences the most obvious uses of (B) & (D) also ascribe the end-relational property of promoting a fulfilling and meaningful life. (C) has a particularly moral feel to it—and I will discuss how to accommodate moral goodness into an end-relational theory of goodness below.

In summary, the nihilist is committed to saying that any sentence that ascribes intrinsic goodness to things is not true. However, she need not say that any sentence that *appears* to ascribe intrinsic goodness is not true. She can claim that these sentences implicitly introduce a non-intrinsic notion of goodness, and thus, can be true. The nihilist can thus agree with the realist about which things are good and which are not. She merely

⁹⁸ This example is reminiscent of Finlay’s discussion that although it might be good for a serial killer or child molester to experience pleasure, it might not be good *that* those villains experience pleasure (2014) pp. 27-28.

disagrees about the type of goodness that they have. Nihilism about intrinsic goodness is not as revisionary as it may at first seem.

b) The truth of non-intrinsic goodness judgments

Turning now to the topic of our non-intrinsic goodness judgments, the realist has a promising argument that suggests that even some of our non-intrinsic goodness judgments require the existence of intrinsic goodness if they are to be true. This argument should sound familiar—it is both ubiquitous in ethics textbooks, and has a long and venerable history. (Its first appearance may have been in Aristotle’s *Nicomachean Ethics*).⁹⁹ Consider the following paradigmatic contemporary statement of the argument, provided by Shelly Kagan. He writes:

Many objects are valued merely as means to other objects – they are valuable solely by virtue of the fact that they will produce (or help produce) those other objects. Those things valued as a means in this way possess ‘instrumental’ value. But what about the objects that the instrumentally valuable objects are means to? In some cases, of course, objects may possess instrumental value by virtue of being means to objects that are themselves of no more than instrumental value (as means to still other objects). But eventually – or so the thought goes – we must reach objects that are valuable as ‘ends’ or ‘for their own sake’ (Shelly Kagan, “Rethinking Intrinsic Value,” 1998, pp.278-9)

Following this line of argumentation, to show that there exists something that is intrinsically good the realist needs only to point to an example of something that is instrumentally good. If there exists something that is instrumentally good, then we need

⁹⁹ Aristotle writes: “If, then, there is some end of the things we do, which we desire for its own sake (everything else being desired for the sake of this), and if we do not choose everything for the sake of something else (for at that rate the process would go on to infinity, so that our desire would be empty and vain), clearly this must be the good and the chief good” (*EN* I.2.1094a18-22). Because this argument is couched in psychological terms (about what we desire), and not in metaphysical terms (about what is good or valuable), it might only be a precursor of AFN, and not a statement of AFN itself. However, it is clearly the inspiration for contemporary statements of AFN.

there to be something that is intrinsically good from which it derives its goodness. And clearly, if our judgments about the non-intrinsic goodness of things are to be true, then there must be something that is instrumentally good. More formally:

Argument from Need

(1) Assume that for some X, X is good.

(2) Either X is good for its own sake, or it is good for the sake of some further thing, Y.

(3) If X is good for its own sake, then it is intrinsically good.

(Sub Conclusion 1): Thus, if X is good for its own sake, then there exists something that is intrinsically good (namely, X).

(4) If, on the other hand, X is good for the sake of some Y (distinct from X), then it is instrumentally good.

(5) And while Y might itself be merely instrumentally good towards some further good, Z, this chain of goods cannot go on like that forever.

(6) At some point, a chain of instrumental goods must end with something that is intrinsically good, from which the instrumental goods on that chain derive their goodness.

(Sub Conclusion 2): Thus, if X is good for the sake of something else, then there exists something that is intrinsically good.

(Sub Conclusion 3): Thus, if X is either good for its own sake, or good for the sake of something else, then there must exist something that is intrinsically good.

(Conclusion): Therefore, there exists something that is intrinsically good.

Since the nihilist will agree that there is an X such that it is the case that X is good, she will assent to premise (1). Suppose, for instance, that she believes that medicine is good. The AFN would say that since medicine is not good for its own sake, it must be good for the sake of something else. That is, medicine is instrumentally good. And while it might be the case that medicine is a means to something else that is instrumentally good (for instance, remedying a particular illness), this chain of goods

cannot go on forever. There must be something at the end of the chain that is good for its own sake (e.g. health, the realist might say) that explains the goodness of medicine.

Thus, if the nihilist agrees that it is true that medicine is good, then she is committed to the existence of something else that is intrinsically good.

The Argument from Need thus puts the nihilist in a very uncomfortable position. It purports to show that if X is either good for its own sake, or for the sake of something else, then there must be something that is intrinsically good. Given that we have said nothing to limit what X can be, this argument threatens to show *any* sentence of the form ‘X is good’ is true only if there is something that is intrinsically good. That is, AFN appears to support the Strong Realist Claim:

The truth of *all*, or perhaps just *most*, of our non-intrinsic goodness judgments requires the existence of intrinsic goodness

If the Strong Realist Claim were true, this would be a powerful reason to abandon nihilism about intrinsic goodness. If the AFN is sound, then a nihilist would have to accept a wide-reaching error theory—one that commits her to claiming that even the judgment expressed by the sentence ‘Medicine is good’ would be false. Lucky for the nihilist, the AFN is unsound and the Strong Realist Claim is false.

c) My Objection to the Argument from Need

The Argument from Need has struck many as a plausible and promising argument in favor of the existence of intrinsic goodness.¹⁰⁰ I believe that it is actually deeply flawed. The success of the Argument from Need hinges on how the realist understands ‘instrumentally good.’ As I shall go onto explain, there is no single notion of instrumental

¹⁰⁰ For instance: Beardsley (1965) p.66, Audi (1997) pp.250-1, Hurka (2010) p.5, and Zimmerman (2010).

goodness that makes both premises (4) and (6) true at the same time without begging the question against the nihilist. The AFN thus fails to show that we need intrinsic goodness.

Let us begin with the standard definition of ‘instrumental goodness.’

Definition 1 – Instrumental goodness – X is instrumentally good if and only if it is good for the sake of something else, and it ultimately derives that goodness from something that is intrinsically good.

Clearly, if one is a nihilist, she will also be committed to denying that anything is instrumentally good in the sense given by Definition 1. After all, according to this definition, it is a necessary condition of a thing’s being instrumentally good that it derives its goodness from something that is intrinsically good.¹⁰¹ I argue that the inclusion of this necessary condition that requires that instrumentally good things be linked to intrinsically good things makes premise (4) of the AFN susceptible to counterexample. According to premise (4), if X is good for the sake of Y, then X is instrumentally good. However, if we accept *Definition 1*, end-relational goods are all counterexamples to premise (4).

For instance, when we judge that a parking space is good, or a steak knife, or an umbrella—we are not judging them to be good for their own sakes.¹⁰² These are all examples of end-relational goods—they are all good for the sake of something else. So they meet the first necessary condition of *Definition 1*. But should we go on to conclude that they meet the second necessary condition, too? That is, should we conclude that good

¹⁰¹ Some might worry that the AFN’s focus on *instrumental* goodness excludes other types of derivative value (for instance, Michael J. Zimmerman describes a variety of ‘extrinsic’ values, such as *inherent value*, *signatory value*, and *contributory value*) (2001) pp.251-8. If the definitions of these notions of goodness include the necessary condition that they derive their goodness from something intrinsically good, then my arguments in this section would equally apply to versions of the AFN that include these types of goodness in the place of ‘instrumental goodness’. If these types of goodness do not necessarily derive their value from their relation to something that is intrinsically good, then they are types of goodness a nihilist can accept.

¹⁰² Geach discusses the example of a ‘good steak knife’ (1956) p.37, while Thomson discusses good umbrellas (2008) p.35.

parking spaces, knives, umbrellas, etc derive their goodness from their relation to something that is intrinsically good? Consider what this would commit us to.

Presumably, when I judge a steak knife to be good, I judge that it will easily cut cooked meat into bite-sized morsels. In short, I judge that it adequately promotes the goal of cutting meat. But the goal of cutting meat into bite-sized morsels is not good for its own sake—it is good for aiding in the consumption of meat. The consumption of meat is also not good for its own sake. Now if we were forced to find something that is intrinsically good at the end of this chain of goods, the most plausible answer would be that consuming meat is instrumentally good towards one's having a pleasurable meal, and that pleasure is intrinsically good. Now while I will admit that this is a *possible* answer—it very well could be that in rare instances, people judge that steak knives are instrumentally good because of how they relate to pleasure—more often than not, however, people do not have such a grandiose or involved story behind their judgments about the goodness of steak knives. Their judgment that a steak knife is good is merely the judgment that it is, or would be, useful for achieving the end of cutting up meat; their judgment does not require any further judgment about the goodness of the goal of cutting up meat. A vegetarian who derives no pleasure from eating meat, and who believes that eating meat is a morally despicable goal, can still judge a steak knife to be good.

In summary, when we judge that X is end-relationally good, we need not make any further judgment about the value of the goal, G, relative to which we judge X to be good. G does not need to be intrinsically good. In fact, G could be an odious, downright evil goal. We can agree that an undetectable swarm of nano-bots would be a good murder

weapon, without thinking that murder is a good or worthwhile goal.¹⁰³ There is thus a whole class of goodness judgments that a nihilist may make without committing herself to the existence of intrinsic goodness. She can say that X is end-relationally good, without saying anything about the value of the goal relative to which it derives its goodness. Premise (4) is thus false if we interpret it in light of Definition 1. There are some things that are good for the sake of something else, but that are not instrumentally good.

My objection to AFN thus shows that the Strong Realist Claim is false. The Strong Realist Claim says that unless there exists at least one thing that is intrinsically good, all (or many) of our goodness judgments would be false. However, many of our goodness judgments ascribe end-relational goodness to things. As noted above, even some sentences that do not seem to be about end-relational goodness might, in fact, ascribe end-relational goodness. We do not need there to be anything that is intrinsically good to explain why they are true. The judgments that parking spaces, knives, murder weapons, etc are good can be true without there being anything that is intrinsically good. Thus, the nihilist can still believe that many of her judgments about non-intrinsic goodness are true without having to posit something that is intrinsically good.

To avoid my counterexamples, the realist might suggest that AFN uses a weaker notion of ‘instrumental goodness’ than the one we have up until now been considering.

She might offer us the following:

Definition 2 – Instrumental Goodness – X is instrumentally good if and only if it is good for the sake of something else.

¹⁰³ Given what he says about good thieves and good torture instruments, it is reasonable to assume that Richard Kraut denies this. He says: “If we believe that no one should ever steal, because of the great harm theft always does, then we should say there are no good thieves” (2007) p.270.

According to *Definition 2*, end-relational goodness is a kind of instrumental goodness. So end-relationally good things are no longer counterexamples to premise (4). They are by this definition, ‘instrumentally good.’ However, something that is end-relationally good need not be related to intrinsic goodness in the way required by premise (6). There need not be any intrinsic good that we must appeal to in order to explain our judgment that a steak knife is good, for instance. Thus, under *Definition 2*, premise (4) is true, but premise (6) is false. *Definition 2* is far too weak for the realist’s purposes.

Given the role that putative examples of instrumental goodness are supposed to play in the AFN, I think it is clear that Definition 1 best captures how the realist defines ‘instrumental goodness.’ That is, it is part of the definition of ‘instrumental goodness’ that things that are instrumentally good are, ultimately, good for the sake of things that are intrinsically good. For the rest of the paper, I will be using ‘instrumental goodness’ in this way. It is thus, worth mentioning that on this definition, a nihilist about intrinsic goodness is also a nihilist about instrumental goodness. Since my end-relational account of goodness commits me to saying there is no intrinsic goodness, it thereby also commits me to saying there is no instrumental goodness (in the sense captured by Definition 1).

I have shown that AFN does not support the Strong Realist Claim. However, it might still support a more modest claim. All that my arguments have shown so far is that for some X’s, it can be true that X is good, and yet, we do not need to appeal to the intrinsic goodness of some other thing, Y, in order to explain the goodness of X. The realist could accept this point, while arguing that the AFN might still support the Modest Realist Claim:

There is at least one true goodness judgment whose truth cannot be explained without appealing to intrinsic goodness.

“So forget steak knives and umbrellas” the realist might say, “let’s return to the goodness of medicine. Surely medicine is instrumentally good, and thus, there must exist something else that is intrinsically good (perhaps health) from which medicine derives its goodness. That is, we *need* there to be something intrinsically good to explain why it is true that medicine is good.”

Unfortunately for the realist, this example fails to show that there must be something intrinsically good—and so would any other. After all, the nihilist can agree that medicine is good, but doing so does not commit her to the existence of something with intrinsic goodness. The nihilist can claim that medicine is merely end-relationally good towards promoting the goal of curing a person of illness. And as you recall, when we say of something that it is end-relationally good, we do not need to make any further claim about the goodness of the goal that it promotes. Consequently, the nihilist does not need to say that the goal of curing people is a good goal, in order to explain the goodness of medicine. It is open for the nihilist to respond to any of the realist’s alleged examples of instrumental goodness, that they are just instances of end-relational goodness.

The realist will want to insist that medicine is not *merely* end-relationally good, but that it is also instrumentally good (in the sense captured by *Definition 1*) because it derives its goodness from the intrinsic goodness of the goal of curing people of illness. However, this insistence is mere question begging; whether or not anything is intrinsically good is the very issue at hand! So long as it remains open to the nihilist to claim that any alleged instance of instrumental goodness is merely end-relationally good, the AFN fails.

IV. Objection: It is not possible that all goodness is end-relational

The realist might accept everything I have argued to this point. She might concede that any sentence that appears to ascribe intrinsic or instrumental goodness can be interpreted as ascribing end-relational goodness instead. Thus, any alleged example of intrinsic or instrumental goodness that she points to, the nihilist can simply insist that it is merely end-relationally good. But even if it is possible in this way to give each goodness judgment an end-relational reading, it might still be implausible to give *all* goodness judgments end-relational readings; it is not plausible that *all* goodness is end-relational.

Robert Audi, in his book, *Moral Knowledge and Ethical Character*, expresses doubts about the plausibility of the view that all goodness is end-relational.¹⁰⁴ When discussing chains of goods, or as he calls them “valuation chains,” Audi notes that they can take four possible structures: (a) They can end in something intrinsically good, (b) they can go on forever, (c) they can be circular, or (d) they can end in something that is not intrinsically good.¹⁰⁵ Chains with this last structure he dubs “terminatingly good.” My proposal is that no chains terminate in something intrinsically good, but that all valuation chains are, in fact, terminatingly good. And while Audi does not deny that some chains will be terminatingly good, he argues that it is implausible that *all* chains are. He writes:

It would be widely agreed, however, that if X is terminatingly good (hence merely instrumentally good)¹⁰⁶ then it grounds no basic practical

¹⁰⁴ Audi (1997)

¹⁰⁵ Those familiar with a similar debate in epistemology over justification will also be familiar with this menu of options.

¹⁰⁶ Audi is here using the weaker definition of ‘instrumental goodness’ that I gave as Definition 2 above. He says: “Consider the existence of instrumental goods—something virtually anyone will grant, since it simply implies that some things are efficient in bringing about others” (1997), p.250. So on this view, end-relational goodness is a form of instrumental goodness. Again, I am using the term ‘instrumental goodness’ in the stronger sense I gave above in Definition 1.

reasons... On arriving at the termination, one could not say anything positive about the goodness of X as the terminal element, since X is by hypothesis, not good in any sense (Audi, *Moral Knowledge and Ethical Character*, 1997, p.250).

Audi points to the second role that intrinsic goodness is said to play in our practical lives. It grounds basic reasons. Without intrinsic goodness, we would have no basic practical reasons, so we need to posit intrinsic goodness.

While I do not make anything of it, it is worth noting that even if X is the terminal element on a terminatively good valuation chain it does not follow that X is “not good in any sense.” X would not be good for the sake of some further thing *on that particular valuation chain*, but X would likely belong to other valuation chains, and it might still be good for the sake of something on one of those other chains. This is merely a quibble with Audi’s phrasing. I take it that the more important claim he is making is that the terminal element on a terminatively good valuation chain is not good in any relevant sense, and thus, does not ground basic practical reasons. If this claim is correct, then it is a consequence of my end-relational account of goodness that we have no basic practical reasons. This bullet is far too big for me to bite; a theory of goodness that did not allow for basic practical reasons would be wildly implausible! Fortunately for me, Audi has overstated his case here.

Audi’s argument only succeeds if we assume that intrinsic goodness is the only ground for basic practical reasons. But this assumption is false. Suppose that X is beautiful, or just, or right. Wouldn’t the beauty of X, or the justice of X, or the rightness of X ground practical reasons? T.M. Scanlon thinks so. He claims

According to that definition, terminatively good valuation chains do not contain instrumentally good things—only those chains that end with something intrinsically good do.

that even though justice (etc) is good, it is not the *goodness* of the justice of X that explains why we ought to favor X. In fact, it would be odd, Scanlon argues, to list both the justice of X and the goodness of the justice of X as reasons to favor X.¹⁰⁷ According to Scanlon, it is the justice of X that gives us reasons to favor it, while the goodness of X merely ‘passes the buck.’¹⁰⁸ That is, saying that X is good just is saying that X has other properties that recommend it (such as beauty, justice, etc). As such, if a just state of affairs were the terminal element of a terminatingly good chain, then the justice of that state of affairs would ground basic practical reasons, and explain why we ought to favor the end-relationally good things that make up that valuation chain. Audi is thus wrong to claim that no terminatingly good chain can ground practical reasons.

There are two weaknesses of this response. First, it relies on our accepting Scanlon’s ‘buck-passing’ account of value—a view that some philosophers remain unconvinced by. Second, the realist does not need to follow Audi in making the exceptionally strong claim that no terminatingly good chain grounds practical reasons. In order to show that we need there to be intrinsic goodness, she need only claim that there are *some* basic practical reasons that are grounded in intrinsic goodness. If the realist can identify certain basic practical reasons that we cannot account for by only using valuation chains that do not end with an intrinsic good, that would show that there must be at least one valuation chain that ends in something that is intrinsically good. Let us call this the Modest Realist Claim about Reasons. If this Modest Claim about Reasons were true there would be an

¹⁰⁷ Scanlon (1998) p. 97

¹⁰⁸ *Ibid*, pp. 96-8

aspect of our practical lives that we cannot explain without appealing to the existence of something that is intrinsically good. In the next section, I will consider a promising argument that could be used to support the claim that an end-relational theory of goodness cannot account for all basic practical reasons.

V. Support for the Modest Realist Claim about Reasons

Consider the following example that I will call ‘*Scrooge*.’ Imagine a man named Scrooge who is remarkably reminiscent of the famous character of the same name from Charles Dickens’s classic, *A Christmas Carol*. He is a bitter, grouchy misanthrope whose only motivation in life is the acquisition of money. Suppose that Scrooge makes his money by charging excessive interest on ‘payday loans’ he makes to extremely poor people who are in desperate need of (and have no other means of acquiring) cash. Scrooge feels no mercy for the poor, and takes no greater pleasure than collecting what is owed to him. In fact, his business is so important to him that he rarely thinks about anything else. He has no friends, family, or loved ones to speak of. This does not bother Scrooge, though, as interpersonal relationships like those would only get in the way of his making more money.

I suspect that none of us would claim that Scrooge is living an admirable life. In fact, most of us would likely judge his life to be downright deplorable. If asked, we would likely point to two very serious practical mistakes that we believe Scrooge to be making. First, his attitudes and behaviors (what we might call his ‘responses’) towards money indicate that he values money much more than it deserves to be valued. His

responses to money are not fitting. Second, his attitudes and behaviors towards people and towards interpersonal relationships indicate that he values them much less than they deserve to be valued. His responses to people and interpersonal relationships are not fitting.

The realist would here point out that if we were to assume that there is both intrinsic and instrumental goodness, we could give a very simple and natural explanation of Scrooge's mistakes. The realist's explanation might look something like the following: Money is something that is merely instrumentally good. As such, there is a set of attitudes and behaviors relative to it that are appropriate or fitting. For instance, we might think that money should only be valued as much as, and never more than, the intrinsically good end or ends that it promotes. Thus, when a person like Scrooge dedicates all his time and energy into acquiring money, this indicates that he values money much more than the ends from which it derives its goodness. He has basic practical reasons not to do this; it is a serious practical mistake. Further, we might also think that things that are instrumentally good should only be valued *hypothetically*—that is, we should only value them on the condition that they actually promote an intrinsically good end. We should cease pursuing them when their acquisition will no longer promote the end or ends from which they derive their goodness. At some point, once Scrooge has accumulated a significant amount of wealth, acquiring more money will not significantly promote any of the ends relative to which an interest in money is justified. At that point, it would be a mistake for Scrooge to continue to spend time and energy acquiring money. Let us stipulate that Scrooge, however, is unwilling to adjust his responses to money. He would continue spending his time and energy collecting debts even when he can obtain no

further significant benefits from having the additional money. He has reason not to do this; this too is a serious practical mistake.

On the other hand, claims the realist, people and the interpersonal relationships one can have with them have intrinsic value. We might think that things with intrinsic value deserve a central place in our practical lives. We ought to value them *categorically*—that is, we ought to pursue them regardless of what other ends we have, and we ought not be willing to cease valuing people or their friendship for the sake of other ends (unless these other ends are at least equally valuable). They should be valued for their own sake. Scrooge's response to people and interpersonal relationships demonstrates that he values them merely hypothetically (if he even values them at all). This too, is a serious practical mistake.

The realist thus has a very plausible and natural explanation for the kinds of practical mistakes that we encountered in *Scrooge*. Money is merely instrumentally good, and thus, it is unfitting for Scrooge to value it for its own sake. Friendship is intrinsically good, and thus, it is unfitting for Scrooge to fail to value it (or to value it only for the sake of something else). Presumably, the realist believes that the nihilist's ontology is too sparse to account for these practical mistakes. If it is fitting to value some things for the sake of other things, and some things for their own sakes, don't we need there to be (at least) two types of goodness: instrumental and intrinsic?

VI. A Plausible Reply

I believe that the nihilist can give a plausible explanation for the practical mistakes exhibited in *Scrooge*—one that need not posit any other kind of goodness other

than end-relational goodness.¹⁰⁹ My claim is that an end-relational theory of goodness can fully account for the practical mistakes that the realist believes forces us to posit both instrumental and intrinsic goodness. If I am correct, then the nihilist's theory is simpler than the realist's, and has all the explanatory power as well. We thus ought to favor the nihilist's view.

Initially, it seems unlikely that a theory that only posits one type of goodness could explain why there are (at least) *two* different sets of responses that are fitting to have towards good things. It seems natural to think there will be a distinct type of goodness corresponding to each set of responses. However, I suggest that to explain the mistakes that Scrooge makes we do not need to distinguish two types of goodness—but we need to distinguish two types of *goals*. By distinguishing between (a) foundational goals that we as agents cannot help but take an interest in, and those (b) subsidiary goals that we can choose to take an interest in or not, the end-relational theorist can explain the fittingness of the two types of responses characterized in *Scrooge*. Briefly, the account I will be defending is that those things that *directly* promote the achievement of the first set of goals deserve a central place in our practical lives, and it is only fitting that we value them categorically. As I define it, X directly promotes a goal, G, if and only if X is the penultimate thing on a chain of end-relational goods that terminates with G. Those things that do not directly promote the achievement of foundational goals, or, that are only good (either directly or not) relative to our subsidiary goals, on the other hand, do not deserve such lofty treatment. It is fitting to value them hypothetically.

¹⁰⁹ Many actual nihilists, like Thomson and Kraut, are not end-relational theorists. In Chapter Four, I argue that their views can be subsumed under an end-relational theory of goodness like the ones that Stephen Finlay and I defend.

Unfortunately, I do not have space here to develop the view in much detail. I believe that the following sketch will adequately support my claim that an end-relational theory of goodness can explain the relevant practical mistakes seen in *Scrooge*. First, let us begin by identifying a plausible candidate for a foundational goal that we as agents cannot help but care about.¹¹⁰ Consider the goal of living a fulfilling and meaningful life. I will admit that this notion is a bit murky, so let me elaborate on it some. By ‘meaningful’ I do not mean ‘good’ or ‘valuable.’ Were I to mean that, my nihilistic account would be a realist account in disguise. Instead, what I have in mind here when I say ‘meaningful’ is the kind of thing that gives one’s life a purpose—a life plan, so to speak. A person who dedicates her life to helping others would be living a meaningful life, but so would be the person who is merely working to make ends meet. This notion of meaningfulness is thus, very normatively thin. My claim is reminiscent of David Velleman’s claim that an agent needs her actions to make sense to her—if she cannot say why she is acting as she is she will cease the activity until it is meaningful to her.¹¹¹ Similarly, if a person were to believe that her life was meaningless or without purpose, she would be distraught and would look for some purpose around which she can orient her life.

So suppose (and I think it is quite plausible) that the goal of living a full and meaningful life is one of these fundamental goals that we cannot help but care about.¹¹²

¹¹⁰ It seems quite possible to me that there could be a plurality of these types of goals, however, it does remain a possibility as well that there is really just one basic and fundamental goal. While I suspect that there is a plurality, nothing that I say here will require taking a stance either way.

¹¹¹ David Velleman (2007) pp.15-46

¹¹² Harry Frankfurt claims there are things we cannot help but care about in “The Importance of What We Care About” (2007) pp.80-94. There, he seems to be focused on things that the individual cannot help but care about, and it is thus not clear if he believes that there are any

It is quite plausible that, *in general*, friendship directly promotes the goal of having a full and meaningful life—thus, friendship is end-relationally good relative to this goal.

Insofar as we care about having such a life, we ought to care about friendship as well.

When I say that friendship ‘in general’ directly promotes this goal, what I mean is that nearly all life plans that people might adopt are such that having friends would contribute to the advancement of that plan. However, imagine a person who is single-mindedly dedicated to a form of art—say, writing literature or painting. She is a talented artist, and the work she produces is excellent (and appreciated by many). Let us call her

‘Dickinson.’ When asked about her process, Dickinson mentions that she intentionally cultivates a lonely life in order to channel her intense feelings of sadness and isolation into her art. Let us stipulate that if Dickinson were not so lonely, her art would suffer, and she would also be less motivated to work. I believe that Dickinson’s life plan is one that would not be promoted by friendship. Life might be more enjoyable or pleasant for her if she were to have friends, but it would come at the cost of what gives her life meaning.

To be clear, I do not wish to make the strong claim that she would be making a mistake if she were to develop some friendly relationships; I do not believe that she ought not cultivate friendships. I am, however, making the claim that she is not necessarily making a mistake by not developing those relationships.

I contend that life plans like that of Dickinson are quite rare, so it is still correct to say that, in general, friendship directly promotes the goal of living a full and meaningful life. Most lives will be made more full or meaningful with the addition of some loving relationships. So what should we say about Scrooge? Scrooge thinks that his life plan will

things that people, in general, cannot help but care about. I am making this point as a reasonable, and friendly, extension of his view.

not be advanced by the addition of friends. He is singularly interested in his business, and he believes that friends will detract from that. And to a certain extent, he is correct. Maintaining close relationships would require some of his time and energy that he might otherwise dedicate to his business. Further, such relationships might sensitize him to the sufferings of other people—lessening his enjoyment of collecting debt from the poor. In this way, he is like Dickinson. His life plan is characterized by his singular drive to work, and developing friendships will lessen the quality and quantity of work he will be able to produce. He would insist that having friends would come at the cost of some meaning in his life: so why do we believe Scrooge is making a mistake by not having friends, when we do not believe Dickinson is¹¹³?

The answer is simple. Dickinson is correct when she concludes that having friends would frustrate her having a full and meaningful life; Scrooge is not. Notice that Scrooge's life plan is ultimately shaped by two desires. First is the desire for money for its own sake. I will discuss this desire in some detail below. Second, however, is the desire for pleasure. He works so much because he *enjoys* collecting debts—it is fun to him. Given that part of his motivation for being so dedicated to his work is that it is a means of enjoying pleasure, we can recommend to Scrooge that he would better meet this goal by diversifying the types of pleasure he pursues. Having loving relationships, hobbies, travelling, etc are all other ways of adding the enjoyment of pleasure into one's life—and Scrooge is mistaken if he believes that these things won't advance his life plan.¹¹⁴ Dickinson, on the other hand, is not motivated to create her art because it serves

¹¹³ Assuming, that is, that you share my judgment about her case: that Dickinson is not making a mistake by failing to cultivate friendships.

¹¹⁴ Because we are recommending that Scrooge pursue friendship to advance his goal of enjoying pleasure, which in turn, advances his goal of having a meaningful life—we are not recommending

as a means of enjoyment for her. Unlike Scrooge, she knows that having friends would contribute to her pleasure; that is precisely why she avoids it! Her life's meaning comes from producing art for its own sake. She is thus not mistaken when she believes that having friends would come at the cost of meaning in her life. Scrooge is mistaken insofar as part of the reason he collects debts is for the sake of pleasure. There are better means of achieving that goal, and consequently, he would better achieve the goal of having a full and meaningful life if he were to pursue other types of pleasures. But what should we say about Scrooge's pursuit of money for its own sake?

Money is, in general, good for promoting the goal of living a full and meaningful life. In general, people need certain goods and services to live such a life, and money is good for obtaining those goods. But notice that money does not directly promote our basic goals. Of course money can be used to obtain things that directly promote our basic goals, and thus, in those cases we have reason to value and pursue the acquisition of money. However, were we able to receive those goods without first obtaining money, then money would cease to be end-rationally valuable towards our basic goals. Having a shelter to sleep in might be necessary for having a full and meaningful life, but having the money to buy a shelter is not. If we can get what we need to live a full and meaningful life without money, then it would be a mistake to continue to value and pursue the acquisition of money. We should value and pursue money only as much as, and to the extent that, it promotes our goals. It should only be valued hypothetically. Scrooge is thus mistaken for valuing money for its own sake. This error is all the worse,

friendship as something that *directly* promotes Scrooge's goal of having a meaningful life. This does not change the fact that, in general, having friends does directly promote having a meaningful life.

in cases like Scrooge's, when it leads to disvaluing the things that do promote our basic goals.

Returning to the language introduced by Audi, what I am suggesting is that all valuation chains terminate in goals. Some of these goals are basic and we cannot help but care about them. Others are not. Recall the discussion of the steak knife. In that example, we noted that the goal of cutting steak into bite-sized morsels was not an intrinsically good goal, nor was it a means to some intrinsically good goal. Putting it that way assumes realism. Now, however, we can recast that observation in terms that a nihilist can accept. We can say that the goal of cutting up meat is not a fundamental goal that we cannot help but care about, nor does it directly promote a goal we cannot help but care about. For those who do care about cutting up meat, the goodness of a knife relative to that goal would give that person reasons to care about or favor the knife. For those of us who do not care about cutting up meat, the goodness of a steak knife gives us no reasons to care about or favor the knife. On the other hand, there are some fundamental goals that we all cannot help but care about. Living a full and meaningful life might be such a goal. When some X is good relative to one of these fundamental goals, we all have reasons to care about or favor X (but only as much as, and to the extent that, it promotes that goal).

Notice I do not say that we have reasons to favor the fundamental *goal* that terminates the valuation chain. I say that we have reasons to favor the penultimate thing that is good relative to that goal. Since we all care about those fundamental goals, we consequently all have reasons to pursue them. In this sense, it isn't literally true that we ought to favor anything categorically. As defined above, if we ought to favor some X categorically, then we ought to pursue X regardless of what ends we have. Strictly

speaking, on the story I just gave, we do not literally value friendship (or anything for that matter) categorically then. The reason we ought to pursue friendship is that it promotes a goal that we cannot help but care about. We can take it as a given that people share that goal, and will continue to share that goal, so those things that directly promote it are special. So even though it is not literally true that friendship ought to be valued categorically—there remains some truth behind the distinction between categorical and hypothetical goods. Some goods are such that they promote a goal that people might take an interest in. We only have hypothetical reasons to pursue those things; that is, if we were to take interest in the goals, then we would have reasons to pursue the things that promote those goals. Other goods are such that they promote fundamental goals that people do care about (and cannot help but care about). We still only have hypothetical reasons to pursue those things if we had interest in the goals they promote. The fact of the matter is, however, that we do care about those goals, and always will. Consequently, our reasons for pursuing these goods are far more steady (because we cannot stop caring about the goals) and deeper (since the goals are foundational) than the hypothetical reasons we have to pursue the goods that promote those subsidiary goals we might or might not take interest in. That is to say, we can preserve something very much like the categorical/hypothetical reason distinction without positing two kinds of goodness.

In summary, we do have reasons to favor and respond to things that are end- relationally good when those things promote fundamental ends that we cannot help but care about. Audi was thus wrong to claim that there are basic reasons only if there are intrinsic goods. Further, when something directly promotes a fundamental end that we

cannot help but care about, we have deep, steady, basic reasons to favor that thing¹¹⁵. It should be a central part of our practical lives, and it would be a mistake to fail to favor it. When something (a) indirectly promotes an end, or (b) directly promotes ends that are not fundamental, or that we can help caring about, then we ought to care about it a different manner. For instance, it is a mistake to structure our practical lives around obtaining such a thing. These two types of practical mistakes are the ones that the realist claims that we could not explain without appeal to intrinsic and instrumental goodness. By showing that we can explain these mistakes, while still being a nihilist about intrinsic goodness, I have shown that we do not need intrinsic goodness.

VII. End-Relational Moral Goodness

At the outset of this chapter, I promised that I would show that an end-relational account of goodness, including moral goodness, would not be too revisionary. Above, I argued that given any purported list of intrinsic goods that a realist points to, the nihilist can agree that those things are good in a privileged way that makes it a practical mistake to not favor or behave towards them in some relevant way. While the end-relational theory of goodness commits us to the claim that nothing is good for its own sake, this turns out to be not too revisionary of a commitment.

As noted above, many moral philosophers have either equated intrinsic goodness with moral goodness, or reduced moral goodness to intrinsic goodness. Consider the Classical Utilitarian who says that happiness is the only thing that is intrinsically good. The Classical Utilitarian then tells us that it is morally required that we maximize

¹¹⁵ Again, we do not literally have categorical reasons—but we can nonetheless preserve something like the hypothetical/categorical distinction with respect to reasons.

intrinsic goodness. According to Classical Utilitarianism, intrinsic goodness is the only kind of goodness that is morally relevant. Or consider the Kantian who believes that rational agency is the only thing that is intrinsically good, and then tells us we are morally required to respect rational agents. I could list other moral theories with this structure, but the general schema is this: (a) Identify some thing or things that are intrinsically good, (b) declare that some set of responses to that thing or things that is morally required. If my end-relational theory of goodness is correct, then all moral theories that can be characterized by this schema are mistaken. Since nothing is intrinsically good, we can never satisfy the first condition (a), and thus, no moral theories of this type are true. Given that there are many deontological, consequentialist, and virtue theories that follow this schema, the end-relational theory of goodness seems to entail that a vast variety of moral theories are false. That is a pretty revisionary consequence of the view.

An end-relational theorist, however, can defend moral theories that are nearly identical to the ones mentioned above. She simply needs to construct the correlating theories according to a different schema—one that does not posit intrinsic goodness.

The starting point for an end-relational theorist is to identify at least one goal that is a *moral* goal. I can see two strategies here. First, one could try to identify some fundamental goal or goals that morality itself is meant to promote. For instance, one might insist that morality is required for a society of people with different interests to cooperate and live together in harmony. This is a plausible claim made by moral philosophers independent of the discussion of nihilism/realism of intrinsic goodness. David Copp, for instance, defends this claim in support of his Society-Centered Theory

of moral justification.¹¹⁶ Insofar as morality is necessary for social harmony and cooperation, and insofar as most of us are pursuing life plans that require our living in proximity to—and cooperating harmoniously with—other people, we could say that morality directly promotes the fundamental goal of living a meaningful life. This view would thus explain the appropriateness of our giving moral goals a privileged position in our practical lives. We cannot help but care about having a meaningful life, and since morality directly promotes that goal, we have reasons to favor it categorically. Some, however, might find this interpretation of morality as ‘too instrumental.’ They might insist that the goal of treating others equally, for instance, is not merely good because it promotes a harmonious society and thus, meaningful lives. They might think that the goal of treating others equally needs no further goal to justify it. Philosophers who reject this first end-relational account of morality might instead favor the view that there is a set of goals that is the set of moral goals, and it is just a matter of brute fact which goals belong to that set.¹¹⁷

Having pursued one of the above strategies, the end-relational theorist now has a list of one or more goals that are the *moral* goals. She can now begin to identify which things help promote those goals, and are thus, morally good. Let us apply this end-relational schema to our earlier example of Classical Utilitarianism. The end-relational utilitarian can either say that it is a brute fact that maximizing happiness is the fundamental moral goal, or that maximizing happiness directly promotes some further fundamental goal such as living a meaningful life. She can then make all the same

¹¹⁶ Though, he does not strictly speaking think a moral code is required. He believes that societies with a moral code would be better able to meet their needs Copp (1995) pp.195-6

¹¹⁷ Moral theorists of this sort are free to claim that the set of moral goals contains just one goal, or several goals.

prescriptions that the realist Classical Utilitarian makes about which actions or rules maximize happiness. For instance, she might insist that the morally right thing to do is to switch the runaway trolley from its current path, thus killing one person to save five others.¹¹⁸ Both the end-relational and realist Classical Utilitarians will say that the outcome of having one dead instead of five is one that is morally better than having five dead instead of one. They will both say that it is better because it maximizes utility. They will simply disagree about what makes it true that the maximization of utility is good.

The most revisionary aspect of an end-relational theory of moral goodness, then, is not the resulting moral theory. An end-relational theorist can defend a moral theory that looks nearly identical to, and agrees with her favored realist moral theory on all claims (except those that ascribe intrinsic goodness). The most revisionary aspect of an end-relational theory of goodness is what it says about the relation between the right and the good. The classical schema I introduced above involved first identifying what is intrinsically good, and then defining moral rightness in terms of its relation to the good. This is a very common procedure in the history of moral theorizing. However, this schema is closed off to the end-relational theorist. She needs to instead start by identifying what the moral goals are—that is, which goals are morally required, prohibited, etc—and then defining what is good in relation to those goals. Thus understood, an end-relational moral theory would not define moral rightness in terms of the good; it instead defines moral goodness in terms of the right. Defining the good in terms of the right is a less common, though not unheard of, way of relating the right and the good.

¹¹⁸ Foot (1967) introduced the ‘Trolley Car Example’.

In summary, though an end-relational theorist believes that nothing is intrinsically good, and thus, any moral theory that posits intrinsic goodness is mistaken, she can still defend moral theories that are nearly identical to those defended by the realist about intrinsic goodness.¹¹⁹ Her argument for her preferred moral theory will have to start with claims about which goals are the moral goals instead of with claims about which things are intrinsically good, but this is hardly a reason for rejecting end-relational theories of goodness! I conclude that the end-relational analysis of non-moral goodness does not have any implausible consequences when applied to moral goods. We thus have no good reason to reject the Commonality Thesis, or the extension of the end-relational analysis to moral goodness.¹²⁰

VIII. Objections

The view I just provided was just a sketch. As such, it will likely invite several misinterpretations. Such misinterpretations often give rise to plausible-sounding, but ultimately misguided objections. I shall now respond to some of these objections, in the hope of clarifying the view I just presented.

According to the view I just sketched, basic reasons for acting are grounded in what we care about. But there are people out there who care about a lot of terrible, deplorable things. So, an objector might claim, according to my view those people have

¹¹⁹ Finlay believes that his end-relational theory of goodness is compatible with the existence of intrinsic goodness (2014) pp.197-206.

¹²⁰ In Chapter Five, I will evaluate to what extent an end-relational moral theory can accommodate other intuitive constraints on what we expect from a moral theory. The goal of this section was just to show that an end-relational theory is not too revisionary with respect to what kinds of moral theory it can defend.

reasons to do those terrible things. Surely people do not have reasons to do such terrible things, thus, my view is mistaken.

This putative reductio of my view fails, however, since it is *not* a consequence of my view that terrible people have *basic* reasons to do terrible things. My claim was that for all things that we have basic reasons to do or to care about, the reason to do so would be grounded in things we cannot help but care about. I did not claim that for all things that we care about we have reasons to do those things. Such a view is far more revisionary than what I am after. I am trying to present a view that is extensionally no different than the realist's when it comes to what things deserve to be the center of our practical lives. The realist would say that having friends, enjoying harmless pleasures, appreciating beauty, etc. are all intrinsically good things, and deserve to be at the center of our practical lives. I agree that these things deserve to be at the center of our practical lives. I only disagree with the realist's explanation of why these things deserve a privileged spot in our practical lives. The realist thinks that the universe is a better place the more pleasure that exists. I deny that. I simply say that enjoying (harmless) pleasures, in general, contributes to the goal of having a meaningful life, and we all ought to seek out such pleasure to the extent that it promotes our having meaningful lives.

Similarly, the realist might object that my sketch of end-relational moral goodness makes our reasons to be moral depend on the motivations and goals of the agent. The resulting theory is thus, at its foundation, egoistic. It overlooks the reasons that are other-regarding. However, we tend to think of morality as, at its foundation, an other-regarding enterprise.

This objection applies only to one of the two types of end-relational moral theories I described above. This objection does not apply to those end-relational theorists who claim that it is a matter of brute fact which goals fall into the set of moral goals. On that view, it is perfectly reasonable to expect that a number of the moral goals will be essentially other-regarding, and thus, we will have moral reasons to respond in certain ways to the things that promote these other-regarding goals. The other type of end-relational moral theory, however, does suggest that moral reasons are derived from an agent's more fundamental goals. The goal that I have been focused on is the goal of having a full and meaningful life. I have focused on this goal because it really seems to me to be a foundational goal—I can imagine many other goals that promote it, but no further, more fundamental goals that it promotes. However, I do want to leave open the possibility that there might be other equally fundamental or even more fundamental goals that we cannot help but care about. Some of these goals might be essentially other-regarding. If so, then even this view of end-relational morality will have basic other-regarding reasons. However, if there were no such goals, then this type of theory does commit one to saying that moral reasons are derived from an agent's goal of living a meaningful life. Personally, I find this interpretation inoffensive. In fact, I find it easier to believe that my reasons for being moral depend upon my interest in having a meaningful life than the realist's story that my reasons for being moral are grounded in an obligation to make the universe an impersonally better place. My suggestion is that if my objector is unhappy with the apparently egoistic grounding of moral reasons that one is committed to on this version of end-relational morality, then she ought to favor the other type, wherein it is a matter of brute fact which goals the moral goals are.

An opponent might deny that there are any goals that we *all* care about. While it is plausible that most of us care about living a full and meaningful life, it is dubious that this is true of *all* of us. For instance, a teenaged slacker might not care about anything more than idling away his time on a sofa in his mother's basement. When pressed, he might admit that this is not a meaningful way to live, but also might insist that he has no interest in living a meaningful life. He would be a counterexample to the claim that we all care about having meaningful lives.

But this objection is due to a misunderstanding of the scope of the claim that I am making. I do not mean to make a universal claim when I say we cannot help but care about these basic goals. I mean to make a *general* claim—the sort that cannot be falsified by pointing out that it is not true of literally everyone. As I am interpreting things, the claim that “We care about living a full and meaningful life” is akin to the claim that “Human beings have two lungs” insofar as both are true, even though not every human being has two lungs, nor does every human being care about these fundamental goals. That humans have two lungs is a biological truth about the kind of creatures we are; that we all care about living a meaningful life is a fact that might be grounded in the nature of personhood. People are concerned with living meaningful lives. This is true in general, so the slacker is no counterexample.

Further, I am tempted to say that the slacker probably does actually care about living a meaningful life. Perhaps he cares about living a meaningful life, but does not believe any such life is open to him (I doubt that he could sustain such a pessimistic view for long—either he would find some meaning to attach to his life, or he would end his life). Or perhaps, deep down, being close to home and his family are things that are

important to him (he might not even be consciously aware that his living at home in the basement is his way of pursuing a life of meaning). If so, then living in the family basement is conducive to a meaningful life. Of course, we would recommend to the slacker that he do more with himself than idle his days away in the basement. This is because we believe that the slacker might be mistaken about what things give our lives meaning. People can be mistaken about this. Scrooge, for instance, thought that cultivating friendships would decrease his ability to have a meaningful life. I argued that he is mistaken—that his life plan would be promoted were he to develop some loving relationships. Similarly, the slacker might simply not be aware of, or fully appreciate, how gainful employment, hobbies, travel, etc. would promote his having a meaningful life. To be honest, I have little insight into the slacker’s psychology, so I will admit that it is possible that he, at least currently, does not care about living a meaningful life. Again, that would be a counterexample if I meant to ground basic reasons in goals that are shared by literally everyone—but that is not what I meant.

The realist might further complain that I am sneaking in a notion of intrinsic goodness by way of my description of basic goals that we care about. Am I not just saying that some goals are just plain good?

One way in which this accusation might be apt would be if I had said that we *ought* to care about these basic goals, or that we have *reason* to care about these basic goals. Some philosophers believe that intrinsic goodness is reducible to what we ought to do or favor, and thus, they would translate claims about what we ought to do or favor into claims about intrinsic goodness.¹²¹ Were I to say we ought to care about these

¹²¹ For instance: A.C. Ewing (1947) and Michael Zimmerman (2001).

fundamental goals, then I would be, by their lights, sneaking in a notion of intrinsic goodness.

I have not relied on the claim that we ought to care about these fundamental goals. My thinking was that in a certain sense, we do not have a choice. We do care about them, and we cannot help but care about them. Because whether or not we care is not under our control, it is not something that we could have a reason for. Just as it does not make sense for me to ask if I ought to (or have reason to) be over 6 feet tall, or to have been born to my parents, it does not make sense for me to ask if I ought to (or have reason to) care about having a meaningful life. I cannot help but have these traits; and I cannot but help care about certain goals.

Finally, the realist might argue that I am muddling up the issue, and getting things backwards. It is not the fact that we care about certain fundamental goals, as I suggest, that grounds our reasons for favoring those things that promote those goals. It is the intrinsic goodness of those goals that grounds our practical reasons. It is true that we care about these goals, but we care about them because they are good.

The objection notes that there are (at least) two contenders for what grounds our basic practical reasons: the fact that we care about certain fundamental goals, and the fact that these goals are good. The realist is correct in claiming that I have said very little to support the claim that it is the fact that we care about these goals, and not the goodness of those goals, that grounds reasons. But my goal has not been to argue that the realist's view is incoherent or untenable! I have rather been offering the end-relational theory of goodness as an alternative to the orthodox view in order to show that we need not posit instrumental and intrinsic goodness to make sense of certain aspects of our practical lives.

And while it does not settle the issue, I take it to be a significant point in favor of the end-relational theory that it is simpler without losing any explanatory power.

IX. Conclusion

In this chapter, I have been arguing that there is nothing implausible about extending the end-relational analysis of non-moral goodness to moral goodness. In other words, there are no good reasons for rejecting the Thesis of Commonality. I focused my discussion on one admittedly revisionary consequence of an end-relational theory of goodness—that on this view, nothing is good for its own sake, or ‘intrinsically good.’ I examined one line of reasoning that could lead us to conclude that nihilism about intrinsic goodness is an unacceptable consequence, namely: that there are certain roles that intrinsic goodness plays in our practical lives, and that if we do not posit intrinsic goodness, these roles cannot be filled. I looked at two such roles. First, that intrinsic goodness helps explain the truth of at least some of our goodness judgments. I suggested that a nihilist could agree that ‘X is good’ is true for any X that the realist points to. She can claim that any judgment that purports to ascribe intrinsic or instrumental goodness in fact ascribes end-relational goodness. Giving up intrinsic goodness does not commit us to radically altering which goodness judgments we believe are true.

The second role that intrinsic goodness is said to play is the grounding of basic practical reasons—and especially the grounding of moral reasons. The realist claims that we have basic reasons to adopt certain responses to intrinsic goods, and we have derivative reasons to adopt a different set of responses to things that are merely instrumentally good. I suggested that the nihilist could preserve the view that some goods

deserve to be favored categorically, while others only deserve to be favored hypothetically, if she distinguishes different types of goals. There are fundamental goals that we cannot help but care about, and we have basic reasons to favor the things that directly promote these goals. The things that indirectly promote these goals, or that promote less fundamental goals, are ones that we have derivative reasons to favor.

I conclude by stressing that if the nihilistic view that I have sketched in this chapter is plausible, then because it has the same explanatory power as the realist view with fewer ontological posits, it is the simpler view, and ought be preferred. Admittedly, the devil is in the details, and whether or not an end-relational theory of goodness is truly preferable will rest on how the view is spelled out. I will spell out my theory in Chapter Five. For now, I will only claim to have shown, more modestly, that the reasons in favor of realism are not persuasive. The realist is wrong; we do not *need* intrinsic goodness.

Chapter Four – Other Relational Accounts of Goodness

I. Introduction

In the last chapter, I defended the Thesis of Commonality—one of the key premises of my master argument that all goodness is end-relational. Another key premise of that argument is that an end-relational account is the *best* account of non-moral goodness. In Chapter Two, my argument in support of this premise concentrated on demonstrating that an end-relational account of goodness gives us a very natural, and plausible way of understanding our everyday judgments about the goodness of parking spots, hammers, umbrellas, steak knives, etc. I further argued against some non-relational accounts of non-moral goodness. This of course, left open the possibility of someone defending a relational account of non-moral goodness that was not *end*-relational. In this chapter, I will discuss the other relational accounts of goodness that have been offered in the literature. If any of these other views were as plausible as the end-relational view that I offer, then there would be equal reason to substitute those views for my end-relational view into my master argument. That is to say, for all I have said so far, the master argument given in Chapter Two could still be used to show that all goodness is relational to something other than to ends. To show that my argument supports end-relational goodness, then, I need to show that these other relational theories are less plausible than an end-relational theory.

To assess the plausibility of these competing views, I will evaluate these theories of goodness in two respects. First, how plausible are they as an account of our judgments about non-moral goodness? Second, and more importantly, could we plausibly substitute

one of these types of theories of goodness into my master argument? That is, could we give a plausible account of *all* goodness, even moral goodness, using one of these alternative relational views?¹²² To answer these questions, I will test how well each of these types of account fits with certain pre-philosophical intuitions we have about goodness. I conclude that these other relational theories of goodness all conflict in important way with our intuitions.

And although I lean heavily on our intuitions about goodness, my thought is not that these intuitions are infallible; I simply believe that they track our commonsense concept of goodness and thus constitute significant constraints on what a successful or plausible account of goodness could be like. For this reason, when I show that these other theories of goodness conflict with some of our commonsense intuitions, I should not be interpreted as attempting to offer fatal counterexamples to those theories; I am not aiming to give ‘knock-down’ arguments against these other relational theories! The thesis of this chapter is that when you consider all the virtues and vices of each type of relational theory, an end-relational theory of goodness emerges as the most plausible.

II. Kind-Relational Goodness

a) The View and its Proponents

Historically, the first type of relational view that was defended was a *kind* relational theory of goodness. According to a kind-relational theory of goodness, there is no such thing as goodness simpliciter, but only goodness relative to *kinds*. In discussing

¹²² In the course of defending the claim that his end-relational account of goodness is the best, Finlay also argues against interest-relational and kind-relational accounts of goodness. His arguments focusing on showing that sentences that appear to introduce interest or kind-relational goodness really introduce end-relational goodness (2014) p.32-4 & 36-8.

the plausibility of this view, I will be drawing from the writings of Peter Geach, Judith Jarvis Thomson, and Philippa Foot. Before doing so, a word needs to be said about the extent to which each of these philosophers can appropriately be categorized as kind-relational theorists.

Of the three, Geach is the one who says the least by way of a positive account of goodness. In his seminal article, “Good and Evil,” Geach’s project is mostly a negative one.¹²³ That is, he argues against certain views of goodness (namely, non-naturalist views, like those of G.E. Moore and W.D. Ross, and non-cognitivist views, like those of A.J. Ayer and R.M. Hare). As a result, he says very little about what he thinks goodness *is*. As such, it might be unfair of us to categorize Geach as a kind-relational theorist when he presents little more than a sketch of a relational view. On the other hand, the few positive claims Geach makes in that article provide the chief argument and motivation for thinking that goodness is kind-relational. So even if Geach would resist the claim that all goodness is relative to kinds, much of what he says is relevant to our project of evaluating the plausibility of a kind-relational theory.

One might also be wary to categorize Foot as a kind-relational theorist. It is certainly true that in her book, *Natural Goodness*, Foot focuses on a particular type of goodness that is kind-relational: so-called ‘natural goodness.’ Briefly, natural goodness is the kind of goodness we attribute to a living creature when we claim that the creature is a good member of its species given the form of life for creatures of that species. A species’ form of life is determined by how that species obtains nourishment, develops,

¹²³ Geach (1956)

defends itself, and reproduces.¹²⁴ So for instance, deer need to have keen eyesight, sensitive ears, and be swift if they are to evade predators. A deer that has these features is a good deer (at least, in respect to its ability to defend itself); a deer that is defective in one of these features lacks natural goodness (to some extent). So while natural goodness is obviously kind-relational, one might wonder if Foot would accept the claim that *all* goodness is. For instance, would Foot accept that the goodness of non-living things like toasters or steak knives is also relative to kinds? Foot's answer is brief, but telling. When we evaluate non-living things in the natural world (such as soil or weather) or artifacts (such as toasters or steak knives), the goodness we ascribe to them is:

What I should like to call secondary goodness. It is in this derivative way that we speak of the goodness of, for example, soil or weather, as such things are related to plants, to animals, or to us... By contrast, 'natural' goodness, as I define it, which is attributable only to living things... depends directly on the relation of an individual to the 'life form' of its species. On barren Mars there is no natural goodness, and even secondary goodness can be attributed to things on that planet only by relating them to our own lives, or to living things elsewhere. (Philippa Foot, *Natural Goodness*, 2001, pp.26-27).

If I understand her correctly, Foot believes that good weather is derivatively good because it helps living creatures to pursue their form of life and achieve natural goodness. Steak knives are good, then, when they too contribute to humans achieving natural goodness. I will not discuss the plausibility of this claim in any detail. I will recall, however, my complaint from Chapter Three about a similar interpretation of the goodness of steak knives. There, I claimed that once we identify that a steak knife cuts meat well we could truthfully call it a good steak knife; we need not make any further claim about the goodness of the goal of cutting meat. Similarly, I see no reason why artifacts, etc must contribute to the natural goodness of some form of life or another in

¹²⁴ Foot (2001) pp. 33-34

order to be good. This is, nevertheless, Foot's claim. So, because secondary goodness is derivative of natural goodness, and natural goodness is kind-relational, we can confidently categorize Foot as a kind-relational theorist.

On the other hand, it is probably inappropriate to label Judith Thomson as a kind-relational theorist since she does not explicitly claim that *all* goodness is relative to kinds. In *Normativity*, she defends a theory of goodness that contains at least four different types of goodness properties. Three of these types of properties are fundamentally kind-relational. The fourth type is what she calls "being good-modified." This is the type of goodness we attribute to something when we judge, for instance, that:

- (1) S is good at doing crosswords.
- (2) X is good for England.
- (3) S is good in *Hamlet*.
- (4) S is good with children.¹²⁵

Roughly, according to Thomson, any time the word 'good' is followed by a preposition, this use of 'good' refers to a distinct goodness property. Importantly for our purposes here, at least some of these types of goodness properties are *not* kind-relational.

Thomson even leaves it open that her list of goodness properties is not exhaustive. She writes: "The word 'good' shows up in other constructions, so we should allow that there are still other goodness properties."¹²⁶ Given that Thomson recognizes at least one type of goodness that is not relative to kinds, it is not *strictly-speaking* correct to call her a kind-relational theorist.

On the other hand, Thomson does not limit her discussion only to natural kinds—she also discusses kinds that are quite complex or gerrymandered. For instance, one kind

¹²⁵ Thomson (2008) p.27

¹²⁶ *Ibid*, p. 31

that she discusses is the kind: typewriter made in 1900.¹²⁷ Accordingly, it is open to her to suggest that the above sentences (1) – (4) express propositions that introduce complicated kind-relational forms of goodness. “S is good at doing crosswords” might mean something like S is good relative to the kind: crossword-player. My suspicion is that the main reason that Thomson introduces this fourth type of goodness—being good-modified—is that she is too preoccupied with the grammatical surface structure of sentences, when she should really be concerned with the propositions that those sentences express. The surface grammar of the sentence: ‘S is good at doing crossword puzzles’ is not explicitly about kind-relational goodness, but neither is the sentence ‘That toaster is good.’ However, ‘That toaster is good,’ and ‘That is a good toaster’ intuitively express the same proposition,¹²⁸ and the latter sentence ascribes kind-relational goodness (if, in fact, any sentences do). Why not accept that even sentences that do not explicitly modify ‘good’ with a kind term might nonetheless ascribe kind-relational goodness? I conclude that Thomson is too quick to dismiss this possibility, and that sentences like (1) – (4) can be reasonably interpreted as ascribing kind-relational goodness. Were she to accept this friendly amendment to her view (that is, that the things that are allegedly ‘good-modified’ are actually just good relative to complicated kinds), she would be a kind-relational theorist.

But let’s not worry too much whether or not we can alter Thomson’s stated view so that we can label her a kind-relational theorist. Since three of the four goodness properties she discusses are kind-relational, and since they do most of the work in her

¹²⁷ *Ibid*, p.42

¹²⁸ In certain contexts, these two sentences might express different propositions. Later, I will discuss in more detail context’s influence on the ways we should interpret sentences that express judgments of goodness.

theory, Thomson ends up developing a view that can largely be taken up by a kind-relational theorist. And because her view is well developed and supported by arguments, in much of the discussion below I will be focusing on aspects of Thomson's theory of goodness. To be clear, my doing so should not, however, be mistaken for the claim that Thomson accepts that all goodness is kind-relational. That is not her view as she states it.

b) Motivation for the View

Almost every philosopher who advances any type of relational account of goodness appeals at some point to Geach's distinction between logically attributive and predicative adjectives.¹²⁹ As you will recall from Chapter One, Geach took the linguistic fact that 'good' operates as a logically attributive adjective as good reason to reject the claim that there is such a thing as goodness simpliciter. If 'good' were a predicative adjective, we would expect the proposition expressed by the sentence "Kyle is a good tennis player" to entail both that Kyle is a tennis player, and that Kyle is good. But it does not have this latter entailment, so 'good' must be an attributive adjective. We should understand this example sentence as evaluating Kyle's goodness *qua tennis player*. That is, he is good relative to the kind tennis player. He is not, however, just plain good.

This seems like a natural way of reading sentences like "Kyle is a good tennis player," where the word 'good' is explicitly modified by a kind term, but how should we understand sentences that lack such a modification, such as "Friendship is good"? Geach suggests that even these sentences are to be understood as expressing propositions that contain a reference to a kind. He writes:

¹²⁹ While Geach is cited more frequently, W.D. Ross discussed this distinction when applied to 'good' a few decades prior in *The Right and the Good*, (1930), pp.65-7

Even when ‘good’ or ‘bad’ stands by itself as a predicate, and is thus grammatically predicative, some substantive has to be understood; there is no such thing as being just good or bad, there is only being a good or bad so-and-so (Peter Geach, “Good and Evil,” 1956, p.34).

Thus, according to Geach, all goodness is relative to *kinds*. Much in the same way you cannot be just plain tall, but only (for example) tall for a kindergartner or tall for a basketball player, you cannot be just plain good, but only (for example) a good tennis player or a good person.

c) Virtues and Vices of the View

A kind-relational theory of goodness is attractive for two reasons. First, it gives a straightforward explanation of the inappropriateness of concluding that Kyle is good simpliciter, even if it is true that Kyle is a good tennis player. That is, it explains why ‘good’ operates as an attributive predicate. Again, the explanation is simply that all goodness is relative to kinds, so even if one and the same person, Kyle, is good as a tennis player, this does not ensure that he is good relative to any of the other kinds he belongs to. This explanation is not only straightforward, but it is also *prima facie* quite plausible.

Second, when paired with the Thesis of Commonality, a kind-relational theory of goodness yields a natural and elegant account of *all* goodness—including moral goodness (Thomson, for instance, develops an account of moral goodness that is, at its foundation, kind-relational).¹³⁰ Notice that the kind-realist is uniquely poised to explain two notions that are often treated as fundamental or central to moral theorizing: the notions of a good life, and of a good person. The ‘good’ in both of these notions appears to be kind-relational. A moral theory rooted in an account of what it is to have a good life, or to be a

¹³⁰ Thomson (2008) Chapter XII

good person, could very easily be a moral theory that is rooted in kind-relational goodness. If this is correct, and one could give an elegant account of kind-relational moral goodness, then one could easily replace my appeals to end-relational goodness with appeals to kind-relational goodness into my master argument from Chapter Two, and it would still go through.

The view is, however, open to a number of objections. The first objection starts with the observation that it sounds odd to ascribe goodness to certain kinds of things. Granted, for some kinds, it is very natural to talk about good members of that kind. We not only feel that it is appropriate to judge the goodness of steak knives, for instance, but we even have some idea what it means to say that a steak knife is good. For other kinds, however, it sounds very strange to talk about good members of that kind. Pebbles, swamps, Gila Monsters, and corpses do not seem like they belong to kinds that one can evaluate for goodness.^{131,132} We are simply in the dark about what it would mean for there to be a good member of these kinds. For now, let us call our intuition that it sounds odd to talk of good members of certain kinds the *Odd Sounding Judgment Intuition*.

Now suppose that a kind-relational theorist were to claim that for every kind there are members of that kind that are good. Were she to claim this, she would thus be committed to there being good pebbles *qua pebbles*, good swamps *qua swamps*, etc. This would stretch our ordinary concept of goodness beyond the extension we believe it to have. Intuitively, there are no good pebbles *qua pebbles*. The objection to a kind-

¹³¹ These are some of Paul Ziff's examples (1960) pp. 210-1.

¹³² Foot would object to the inclusion of Gila Monsters on the list of things to which it sounds strange to attribute goodness. According to Foot, Gila Monsters (and indeed, all living creatures with a form of life) can have natural goodness (2001) p.26. I discuss this claim in more detail below.

relational theory of this sort, then, is that it allows too many things to fall under the extension of goodness. Let us call this the *Too Much Goodness Objection*.

Of course, a kind-relational theorist need not be committed to the claim that there are good members of that kind for *every* possible kind. Both Geach and Thomson deny this explicitly. For instance, Geach claims that there are no good events *qua event*, nor things *qua thing*.¹³³ Thomson agrees with these two claims and adds that neither are there good states of affairs, facts, or possible worlds.¹³⁴ These are all kinds such that it does not make sense to speak of good members of that kind. In this way, a kind-relational theorist can deny that there are good pebbles *qua pebbles*. Thus, the view need not be susceptible to the *Too Much Goodness Objection* (at least not for this reason).

Once the kind-relational theorist admits that there are some K such that nothing is a good member of K, however, one might reasonably wonder how to identify which kinds have good members, and which do not. Is there any principled way to demarcate these two groups? Or, to return to the example above: Given that so many philosophers believe that we can meaningfully speak of good events, states of affairs, things, and possible worlds, what can Geach or Thomson say to convince them otherwise?

Here is what Geach says to support the claim that there are no good events *qua* events:

We cannot sensibly speak of a good or bad event, a good or bad thing to happen. 'Event', like 'thing', is too empty a word to convey either a criterion of identity or a standard of goodness... Caesar's murder was a bad thing to happen to a living organism, a good fate for a man who wanted divine worship for himself, and again a good or bad act on the part of his murderers; to ask whether it was a good or bad event would be senseless (Peter Geach, "Good and Evil," 1956, p.41).

¹³³ Geach (1956) p.41

¹³⁴ Thomson (2008) pp.25-26

According to Geach, sentences of the form ‘X is a good thing’ or ‘Y is a good event’ are ‘senseless’ since the words ‘thing’ and ‘event’ are too *empty*. But how should we understand this notion of ‘emptiness’? Given what Geach goes on to say in this passage, we can interpret his notion of emptiness as involving either (a) a lack of criteria of identity, (b) a lack of a standard of goodness, or (c) both. In regards to (a), Geach seems to have in mind that terms like ‘event’ and ‘thing’ are so broad or general that there are few to no restrictions on what properties that X must have to be properly categorized as a member of that kind. Any X can properly be categorized as a thing, no matter what properties it has. However, lacking identity criteria cannot be a necessary condition for a K failing to have good members. After all, terms like ‘pebble,’ ‘corpse,’ or ‘swamp’ do have identity criteria—there are restrictions on what properties a thing must have to qualify as a pebble or a corpse, etc. But ‘That is a good pebble’ still elicits in us the *Odd Sounding Judgment Intuition*. The kind, pebble, has identity criteria, but does not have good members. Thus (a) is not necessary for a K to lack good members. And if (a) is not necessary for lacking good members, (c) cannot be necessary either.

We should thus focus our attention on (b): that certain K lack ‘standards of goodness.’ This condition is certainly more in line with the explanation that Thomson gives for why certain K lack standards of goodness. The distinction between kinds that have standards of goodness and those that do not is central to Thomson’s discussion of what she calls ‘*goodness-fixing kinds*’. Something is a goodness-fixing kind if “what being a K *is* itself sets the standards that a K has to meet if it is to be good *qua* K.”¹³⁵ Functional kinds like toaster or umbrella are obvious examples of goodness-fixing kinds.

¹³⁵ Thomson (2008) p.21

These kinds have essential functions—and to be a good member of that kind is to perform that function well. If you know what an umbrella is, then you know what standards it must meet to be a good umbrella. But it is not just functional kinds that are goodness-fixing, according to Thomson. She claims that the kinds: beefsteak tomato, tiger, and even human being, are goodness-fixing, even though they do not have essential functions.¹³⁶

Thomson argues that it is only when a thing belongs to a goodness-fixing kind that it can be good qua member of that kind. So for instance, the kind toaster is a goodness-fixing kind, so a particular toaster can be good *qua toaster*. But because there is nothing in what it is to be a possible world, event, thing, state of affairs, etc that sets the standards for what a good member of that kind would be, there is no such thing as a good possible world *qua possible world*. The same goes for the kinds: pebble, swamp, and corpse. There are no good pebbles qua pebble, and thus, we have the *Odd Sounding Judgment Intuition* about the judgment, ‘That’s a good pebble’.

Suppose that we grant that some kinds are goodness-fixing kinds, and others are not (this does not seem like a huge concession—in fact, it seems true). I contend that even granting her this fact, Thomson has not yet answered the question with which we began this discussion. That is, Thomson has yet to tell us how we can distinguish those goodness-fixing kinds from non-goodness-fixing kinds. Her ‘answer’ is that only goodness-fixing kinds have standards of goodness that are set by what it is to be a member of that kind. Further, you can tell that a kind is such that it does not set the standards of goodness for that kind if you have the *Odd Sounding Judgment Intuition*

¹³⁶ *Ibid*, p.20

about judgments that attribute goodness to members of that kind. Why do those judgments give us that intuition? Because the relevant kind is not goodness-fixing. This ‘explanation’ is thus circular and uninformative.

Consider an analogous ‘explanation.’ Some substances dissolve in water, while others do not. We call the former set of substances ‘water-soluble’. Now if we were interested in if a certain substance, S, was water-soluble or not, there would be a way to answer that question: namely, we could drop S in water and see if it dissolves. But suppose we want to know *why* things like S dissolve in water. It would not be a very illuminating answer to our question that S dissolves in water because it is water-soluble. We are asking for a deeper answer than that. We want to know *what makes it the case* that things like salt are water-soluble, while things like copper are not. Similarly, Thomson tells us that toasters belong to a goodness-fixing kind because what it is to be a toaster sets the standards of goodness for toasters. But we might still inquire why toasters have this feature, and pebbles do not. Thomson does not give us an answer to this further question, and it is hard to imagine what answer she could give other than that it is simply a matter of brute fact which kinds are goodness-fixing and which are not.

Notice that it is fairly easy to answer this question if we were to give up on being kind-relationalists. One answer that suggests itself, for instance, is that goodness-fixing kinds are the ones that have salient interests or ends associated with them. To be a good member of a goodness-fixing kind is just to promote that salient interest or end. If a kind lacks a salient end (as the kinds pebble, corpse, and swamp do), then that kind is not goodness-fixing. As we saw, the paradigm examples of goodness-fixing kinds are functional kinds like toaster, umbrella, steak knife, etc. These kinds have essential

functions—and it is thus obvious that when we judge their goodness we are evaluating them for how well they satisfy those functions. Their essential functions are the salient interest or end we judge them by—and these ends or interests ground the standards by which we judge members of that kind.

What should we say about beefsteak tomatoes, tigers, and human beings?

Thomson claims these are goodness-fixing kinds, but that they clearly lack an essential function. The view I am suggesting does not require that a kind have an *essential* function for there to be good members of that kind; it requires that the kind must have a salient interest or end associated with it. So, for instance, presumably we would judge the goodness of beefsteak tomatoes relative to how well they would satisfy certain interests we have in cooking and eating. Admittedly, it is a little harder to find salient ends or interests that good tigers or good human beings promote. Of course, we can simply disagree with Thomson that these are goodness-fixing kinds. I for one, get the Odd Sounding Judgment Intuition with regards to the judgment ‘That’s a good tiger.’ To my ears, this sentence only sounds acceptable if we imagine that photographers or zookeepers are looking to capture a tiger that can serve as an exemplary model of its species. So perhaps the salient end that a good tiger serves is being a representative tiger. Notice that this interpretation could possibly subsume Foot’s notion of natural goodness. Foot claimed that a good tiger would be one that has all the features that help it flourish in the form of life of a tiger. Such a tiger would certainly be an exemplary model of its species. So, one might argue that Foot’s natural goodness is not, deep down, a type of kind-relational goodness; it is actually a type of end or interest-relational goodness,

where the salient end makes reference to a kind.¹³⁷ Again, we need not agree with Thomson or Foot that there are good tigers, qua tigers. But if we did agree with them, it is plausible that our judgment that a tiger is good could be saying of that tiger that it promotes the end of being representative of its species.

It is highly unlikely, however, that when we judge that someone is a good human being, that we are judging her as representative of her species. We might, for instance, believe that most humans are petty, awful people, so a representative human being would also be petty and awful. Even if all people are awful, it makes sense to claim that a good human being would be someone who is *not* awful; a good human being would be the rare exception to the norm. We should thus not interpret ‘good human being’ as roughly meaning ‘representative human’. So what does it mean to be a good human being? It strikes me that ‘Kyle is a good human being,’ has a particularly *moral* sound to it. To avoid favoring one view of morality over another, we might simply say that the relevant end or interest that human beings promote, in virtue of which we judge them to be good, is the goal of *being moral*. This neutral way of referring to the salient end allows that we can cash out what it is to be moral in terms of having virtues, doing one’s duty, maximizing utility, engaging in fulfilling loving relationships, etc. I will return to this point in a moment.

For now, it is worth summarizing the discussion so far. I began by raising a worry for kind-relational views: that they cannot give a deep explanation for why we have the Odd Sounding Judgment Intuition when we call members of some kinds of things good. Thomson gave the start of an answer by distinguishing between goodness-fixing kinds

¹³⁷ I will further motivate this claim below.

and other kinds. Only goodness-fixing kinds, she claimed, have standards of evaluation built into what it means to be a member of that K, and thus, it is only relative to goodness-fixing kinds that something can be good qua K. For all other K, we will get the Oddness Sounding Judgment Intuition when we purport to attribute goodness qua K to members of that K. I argued that this answer does not *really* get to the heart of the issue, since it seems that the only way to determine if a kind has standards of evaluation built into what it means to be a member of that K is to see whether or not we have the Oddness Sounding Judgment Intuition when we judge that something is good qua K. The ‘explanation’ for why we get the Oddness Sounding Judgment Intuition is, thus, circular.

I proposed a natural, plausible alternative, where we distinguish goodness-fixing kinds from non-goodness fixing kinds by whether or not they have a salient end or interest associated with them. If they do, then the end or interest generates a set of standards of evaluation, and it will thus sound natural to speak of something as being good qua that kind. The problem for a kind-relational theorist, of course, is that this interpretation makes the goodness of members of goodness-fixing kinds a type of interest or end-relational goodness. It thus gives up on giving a unified, kind-relational theory of goodness. Someone who wished to substitute kind-relational goodness into my master argument could not accept this way of distinguishing goodness-fixing kinds. Worse yet for the kind-relational theorist: there does not appear to be any easy, yet truly kind-relational answer to the puzzles surrounding the Oddness Sounding Judgment Intuition (other than simply claiming that some kinds are such that they have standards of evaluation built into them, and this is just a brute fact that admits of no further explanation). Again, this is not meant to be a devastating objection to kind-relational

theories of goodness. After all, there might yet be an undiscovered way to resolve this puzzle that is open to the kind-relational theorist. Other things being equal, however, we ought to prefer a theory that *does* resolve this puzzle. Interest-relational and end-relational theories both resolve this puzzle in a way that is both compelling and plausible.

d. Kind-relational ethics?

Before moving on to look at interest-relational theories of goodness, I wish to evaluate the plausibility of extending the kind-relational analysis to cover *moral* goodness. I suspect there are a variety of very ingenious and nuanced ways of constructing a moral theory that is compatible with a kind-relational view; I cannot possibly address them all. Instead, I will focus on what I take to be the most natural and fitting forms of moral theory for a kind-relationalist to defend (and also, not coincidentally, the forms of moral theory that actual kind-relationalists *have* defended). I do so fully realizing that, of course, someone might come along and offer a more complex, gerrymandered kind-relational moral theory that sidesteps the criticisms of this section. Were someone to concoct such a view, it would not affect my overall argument. Again, I am only interested in defending the relative claim that a kind-relational account of goodness is, all things considered, less plausible than my preferred end-relational account. In the next chapter, I will show that my account lends itself quite naturally to a simple, and elegant account of morality. If the only way the kind-relationalist can avoid the criticisms I offer in this section is by providing a complicated, baroque moral theory, then all the more support for my claim that my view is more plausible.

As I noted above, because kind-relational views say that goodness is relative to kinds, they stand in a privileged position to make use of the notions of a ‘good person’ or

a ‘good life.’ These two notions are often central to virtue theories of morality, so kind-relational accounts of goodness lend themselves quite naturally to virtue theories of morality. Further, since kind-relationalists often deny that there is such thing as a good possible world, good state of affairs, or good thing (to happen), they generally must deny the plausibility of other types of moral theories. Both Thomson and Foot, for instance, argue against Utilitarianism on the grounds that it is incoherent.¹³⁸ I will thus focus my discussion on the prospects for a kind-relational virtue theory.

As a reminder, since we are only considering kind-relational views of goodness that might replace my end-relational theory of goodness in my master argument, there are even some forms of virtue theory that the kind-relationalist cannot defend—namely, those theories that require the positing of intrinsic goodness. A kind-relationalist can say, for instance, that certain traits are characteristic of a good human or person. She cannot, however, say that those traits, the possession of those traits, or people with those traits, are intrinsically good. For this reason, a kind-relationalist could not support the Perfectionist moral theory that Thomas Hurka defends. One of the constraints that Hurka places on his search for the “best or most defensible perfectionism,” is that a “perfectionist concept of [human] nature assigns intrinsic value to certain properties, and these must on their own seem morally worthy of developing.”¹³⁹ Similarly, a kind-relationalist can talk about what constitutes a good life, but she could not say that the good life is one that is *intrinsically* good—as Aristotle does. Aristotle says that eudemonia is the only thing that is good for its own sake, and the cultivation and exercise

¹³⁸ See Thomson, “Goodness and Utilitarianism,” (1994) and Foot, “Utilitarianism and the Virtues,” (1985).

¹³⁹ Hurka (1993) p.9

of the virtues are constitutive of a life of eudemonia.¹⁴⁰ So, while a kind-relationalist can agree with Aristotle that the cultivation and exercise of the virtues are constitutive of a good life, she cannot agree with his further claims about the intrinsic value of a good life. In short, there is a fairly narrow class of moral theories open to the kind-relationalist.

Whether a kind-relationalist grounds her moral theory with an account of a good person, or with an account of a good life, she will be faced with two very difficult questions. First, why think that a good life/person has the particular features that you identify? Second, having settled the first question, why should we care about having a life like that, or being a person like that? Again, the easy answer here would be that the features of a good life/person are themselves intrinsically good, and thus, having a good life/being a good person is intrinsically good—but this answer is not open to the kind-relationalist. I will argue that, in fact, the only plausible responses to these questions require giving up the thesis that *all* goodness is kind-relational. Thus, a kind-relationalist must either give up on offering a unified kind-relational account of goodness, or, she must admit that she cannot give an explanation of how we identify the virtues, or why we should care about virtues. Either alternative counts against the plausibility of extending the kind-relational theory of goodness to apply to moral goodness.

Philippa Foot grounds her moral theory in the notion of a good human life. Just as a tiger has a ‘form of life’ that is natural to creatures of its species, and a tiger that is not well-suited for its form of life is a bad tiger, a similar thing can be said for humans. The details of the transition from talk of the natural goodness of tigers to the natural goodness of humans are a bit murky. Foot seems to have in mind here that the natural form of life

¹⁴⁰ Aristotle, *Nicomachean Ethics*, 1097^a 15-22, & 1099^b 25-33.

for a human is one that requires a lot of cooperation with other humans. We do not have any natural power to bind people into cooperating with us, so our pursuit of a naturally good life requires that we be able to trust that others will cooperate with us when they say they will. This for instance, explains why it is prima facie wrong to lie. If we lie, then we are not fully cooperating with other humans, and it is part of our natural form of life that we cooperate. The badness of lying, for instance, is thus explained in terms of the naturally good human life. A person who is disposed to lying is thus not a virtuous person.¹⁴¹

That explains Foot's answer to the first question: Why think that a good person has the particular characteristics that you identify? The more troubling question, for Foot, is the question: Why should I care about being a good person—so described? One particularly forceful way of pursuing this question is to draw attention to something that Philippa Foot overlooked. Foot defines the natural goodness of an animal in relation to how fitting it is to its form of life. She states “that it is the particular life form of a species of plant or animal that determines how an animal should be.”¹⁴² It seems to me, however, that there are two ways to fail to fit within a form of life. Foot focuses entirely on one: where a creature is weak or defective and thus not a good member of its kind. But there is clearly another way a creature can fail to fit the form of life for its species—by being different in ways that make that creature flourish.

It is difficult to come up with examples of what I have in mind, partly because most creatures are well adapted to their habitats, and deviating from the natural way of life for that creature does result in that creature's failing to flourish. However, it is

¹⁴¹ Foot (2001) pp.45-6

¹⁴² Ibid, p.32

possible that for some species, the form of life of that species is not conducive to the flourishing of the individual creatures of that species (nor to the species itself). Recall my discussion of the Giant Panda from Chapter Two. Giant Pandas eat only bamboo—a food source from which they barely receive the nutrition they need to get through the day. It is not possible to eat enough bamboo to build up a layer of fat that would sustain a panda if it wishes to hibernate through a cold winter. Bamboo does not flourish in the cold either. So the panda's diet requires that the panda stay somewhere that is fairly warm. As humans encroach upon the panda's natural habitat, it is being pushed into colder areas in which it is not well suited to survive. These pandas natural form of life—the exclusive eating of bamboo—means they will be hungry and cold. Further, pandas rarely mate—and female pandas are actually only fertile for a very short time out of every year. When panda numbers diminish, it becomes rarer and rarer for pandas to meet. It thus becomes even more rare for mating to occur, and for it to occur at a time that can result in an offspring. For all of these reasons, the form of life of the panda is currently an important factor in why pandas are an endangered species. Clearly it would be good for the individual pandas if they were more flexible about what they eat. It would be better for the panda species if they wanted to mate more often, and if they were fertile for longer periods of time. Suppose then that there was a panda that liked to eat more diverse and nutritious foods, and who was fertile and willing to mate more frequently. This panda would not be a naturally good panda in Foot's sense. However, it strikes me that this is indeed a good panda! A good panda is one that flourishes, even if that means living quite differently from its form of life.

My diagnosis of the last example is that Foot's natural goodness is really a form of end or interest-relational goodness. There is an end—the flourishing of the animal or species—and it is constituted, in part, by the procurement of food, mates, and the avoidance of predators. In nearly every case, species have evolved so that a naturally good member of that species is also good at accomplishing these tasks, but as the panda example shows, these two ways of being good can come apart. When they do come apart, I do not see why we should favor the creature that fits the form of life for creatures of its kind, instead of the creature that is better at obtaining food, wooing mates, and avoiding predators.

The worry, of course, is not about animals—but about people. An honest person is virtuous, according to Foot, because lying does not fit our natural form of life. She thus says that dishonesty is a defect in humans. I have suggested a way we might agree with her prior claim (that lying does not fit our natural form of life), but that nonetheless, it is wrong to call dishonesty a 'defect'. Perhaps being dishonest can lead to a person's flourishing more efficiently than being honest and adhering to the natural life of a human. Many politicians surely do well due to their dishonesty! My intuition is that there is no good reason to prefer pursuit of a natural life to pursuit of a life of flourishing (when the life of flourishing does not fit the natural form of life for members of that species). However, if we identify the virtues as those traits that promote flourishing, then we have abandoned a kind-relational moral theory, and instead, proposed an end or interest-relational theory.

Judith Thomson's moral theory equally struggles to answer the question of why we should care about having the traits she identifies as the virtues, without giving up on offering a purely kind-relational theory.

According to Thomson, a feature, F, is a virtue in a kind, K, only if:

- (i) K is a goodness-fixing kind
- (ii) A K is as good a K as a K can be only if it has F
- (iii) It is possible for there to be a K that lacks F, and
- (iv) It is not nomologically impossible for there to be a K that has F.

A virtue is a *moral* virtue in a K if the goodness that we are concerned with in condition (ii) is moral goodness. That is, if F is a moral virtue, then (ii) reads: a K is as morally good a K as a K can be only if it has F.¹⁴³

Let us consider what Thomson would thus say about the moral virtue of honesty. First, according to Thomson, human being is a goodness-fixing kind, so condition (i) is met. Earlier, I suggested that there are reasons to be suspicious of this claim. Let us set those worries aside. Second, it looks as if it is possible for there to be humans that are not honest, and it is not nomologically impossible for a human to be honest (even the cynic will grant that!). So conditions (iii) and (iv) are met. Thus, whether honesty is a virtue hinges on whether or not it is true that a human is as morally good as a human can be only if that human is honest.

The problem with this account is that for any particular human, that human will fall under many different kinds. Larry, for instance, is a person, a gangster, a golfer, a husband, a friend, a mammal, etc. What counts as a virtue relative to some of these kinds will not count as a virtue relative to others. Being empathetic and supportive of others is a virtue in a person, but not one in a gangster. Suppose that Larry is considering having a

¹⁴³ Thomson (2008) p.73

surgery performed on his brain that will make him much less empathetic. We want to be able to say that he ought not have that surgery. However, Larry might note that the surgery will make him more virtuous, qua gangster, but less virtuous, qua person. On what grounds could we insist that he ought to favor being virtuous qua person?

Thomson predicts this worry, and offers the following solution. What we ought to do is avoid defect. When our having a trait would render us defective qua K_1 , and lacking that trait would render us defective qua K_2 then what we ought to do is settled by appealing to whichever K is the more general, 'super-kind'. K_1 is a super-kind of K_2 if all members of K_2 are necessarily members of K_1 .¹⁴⁴ So in Larry's case, since all gangsters are human,¹⁴⁵ but not all humans are gangsters, gangster is the sub-kind, and human is the super-kind. Larry thus ought to avoid defect qua human, even if that means being a defective gangster.

This reply is not successful; given some of the other moving parts of Thomson's theory, it is easy to reconstruct the theoretical impasse. Her theory, after all, allowed for the introduction of complicated, even gerrymandered kinds. Let me introduce one such kind: Machiavellian-Politician-in-Training (or MPIT).¹⁴⁶ I stipulate that (i) an MPIT is a goodness-fixing kind, (ii) it is a virtue in a MPIT that she be dishonest, and (iii) all persons are MPIT's. Some might doubt that it is permissible for me to stipulate (iii),¹⁴⁷ but here is what I have in mind: To some extent, we all have the drive for power, and as we live our lives, we are all exposed to the fact that being ruthless and dishonest is a very

¹⁴⁴ *Ibid*, p.212

¹⁴⁵ Of course, for it to be the case that all gangsters are *necessarily* human, we have to claim that it is impossible for there to be non-human gangsters. Consequently, it would be impossible for an alien like Jabba the Hut to exist. That is clearly false! We should thus understand Thomson as introducing a unique technical kind gangster* where only humans can be gangsters*.

¹⁴⁶ I am here rehearsing an objection given by David Copp (2010) pp.2-3.

¹⁴⁷ I have slightly amended Copp's counterexample to make this stipulation more palatable.

efficient way to gain power. We are thus in training with respect to becoming Machiavellian politicians. Some of us are not as successful at cultivating our drive for power, or at learning how to do so most efficiently; some of us are so unsuccessful that we do not even aspire to become Machiavellian politicians! No matter, we are all nonetheless receiving the training. And since the standards for becoming a Machiavellian are impossibly high, even the most ruthless politicians are still in training (practice makes perfect, after all). Thus, all people are MPIT's.

The kind MPIT is not a super-kind relative to the kind person, but neither is the kind person a super-kind of MPIT. All people are MPITs, and all MPITs are people. By stipulation, an honest MPIT is a defective one. Suppose that yet again, Larry is considering brain surgery, this time that will make him less honest. This would make him a less virtuous person, but a more virtuous MPIT. On what grounds could we insist that he ought to favor being virtuous qua person? We are back at our theoretical impasse.

Thomson might reply that while dishonesty is a virtue in a MPIT, it is not a *moral* virtue in a MPIT. Recall that to be a moral virtue, F must be such that a K cannot be as morally good a K as it can be unless it has F. Larry would not be as morally good a MPIT as he could be if he were dishonest, so honesty is a moral virtue both for a person, and for a MPIT. This reply, however, threatens to give up on the kind-relational commitments of this theory. Nothing about what it is to be a MPIT explains why it would be a moral virtue *for a MPIT*, to be honest. If we are inclined to agree that a morally virtuous MPIT is an honest MPIT, this says more about our estimation of honesty than it does about how we understand MPIT's. The plausibility of Thomson's moral theory would thus rest on our prior commitments to the goodness of honesty—and deep down, would not be

grounded in the nature of kinds. If on the other hand, we deny that honesty is a moral virtue in MPIT's, we are left with the question: Why should Larry—or any of us for that matter—aim at being a morally virtuous person, instead of aiming at being a virtuous Machiavellian Politician In Training? Thomson's view either cannot answer this question, or in answering it, must give up on offering a purely kind-relational theory.

e.) Final Assessment

I noted two obstacles for being a kind-relational theorist. First, kind-relational views cannot explain why some goodness judgments give us the Odd Sounding Judgment Intuition. They leave it a mystery why some kinds are 'goodness-fixing,' while others are not. The most plausible ways to solve this mystery involve appealing to interests or goals. These answers amount to giving up on having a *kind*-relational view, since ultimately, it is the goals or interests, and not the kinds, that are doing the theoretical work. Second, I argued that a kind-relationalist would have difficulty developing a moral theory without abandoning her commitment to a kind-relational theory of goodness. More specifically, she would have trouble (i) justifying any substantive list of virtues, and (ii) motivating the claim that we ought to aim at developing those virtues. These considerations, while not fatal, do warrant our looking elsewhere for a more plausible account of goodness.

III. Interest-Relational Goodness

a) The View and its Proponents

According to an interest-relational theory of goodness, a thing is good if and only if it promotes or is conducive to achieving some interest. Because the term 'interest' is ambiguous, there is more than one way to be an interest-relationalist. My discussion of

these views will focus on the writings of Paul Ziff, John Mackie, and Richard Kraut. But before going into any detail, let us disambiguate (at least some of) the notions that are expressed with the term ‘interest’.¹⁴⁸

On one understanding of ‘interest,’ it refers to a thing that a person in fact takes an interest in. This is the sense of the term used in the sentence: “Woodworking is one of my interests.” For lack of a better term, I will call this the *subjective sense* of ‘interest.’ It is subjective because whether or not a sentence of the form “X is one of my interests” is true will be wholly determined by facts about the speaker’s attitudes. If we were to know all the speaker’s attitudes, we would also know all the interests she has. On another understanding of ‘interest,’ the term refers to a thing that is an advantage or benefit for a person. This is the sense of the term used in sentences like: “It is in your child’s interest that she learn to play a musical instrument.” Whether this sentence is true or not does not turn on what the child likes; we might think that learning to play a musical instrument is in a child’s interest, even if the child hates everything about learning to play that instrument. I will call this the *objective sense* of ‘interest,’ because whether or not a sentence of the form “X is in my interest” is true will not be determined merely by facts about that speaker’s attitudes.^{149,150}

Because there are (at least) these two senses of ‘interest’—there are (at least) two corresponding interest-relational theories of goodness. The *Subjective Interest Theory* says that something is good if and only if it promotes something that a person is actually

¹⁴⁸ Finlay coined the term ‘interest-relational’ in (2001)—however, the view he advanced under that label was actually an end-relational account of goodness.

¹⁴⁹ Several other philosophers make use of this distinction. See for instance: Ziff (1960) pp.219-20, Mackie (1977) p.58, Railton (1986) pp.173-5, Finlay (2004) p.215. The labels ‘subjective’ and ‘objective’ come from Railton.

¹⁵⁰ A person’s objective interest is often partially determined by her attitudes (for instance, in general, it is in a person’s objective interest to be able to pursue her subjective interests).

is interested in. The *Objective Interest Theory* says that something is good if and only if it promotes a person's well-being. Mackie appears to be a proponent of the subjective interest theory; Kraut is a proponent of the objective interest theory. I will also consider a *Hybrid Interest Theory*—in the style of Peter Railton—that says that something is good if and only if it promotes something that a fully informed version of the agent would want the less informed agent to be interested in. Ultimately, I will argue that it is only the Objective Interest Theory that is a plausible alternative to an end-relational theory of goodness.

It is a bit difficult to determine which banner we ought to place Ziff under. He clearly states that he takes goodness to be relative to interests.¹⁵¹ However, he also notes that 'interest' is multiply ambiguous. He writes:

Something must be said about my use of the word 'interest'. I mean to be using that word in an ordinary way. I shall assume you know what that is, that you are familiar with the word. Since the word has been used in extraordinary ways by philosophers it is, I suppose, necessary for me to disassociate myself from the tradition. I take it that interests, motives, wants, wishes, hopes, cravings, longings, likings, hankerings, and so on are all different (Paul Ziff, *Semantic Analysis*, 1960, p.219).

He also says that by 'interest' he does not mean the same thing as 'end' either.¹⁵² I must confess that I am not sure what Ziff means by 'interest'; both the subjective and the objective sense of interest that I discussed above seem to capture 'ordinary ways' in which we use the term. It is unclear which of these, if either, Ziff has in mind. However, I do not think it is vital that we resolve this issue here. What Ziff says about goodness can be equally used to motivate either the subjective or the objective interest-relational view.

¹⁵¹ Ziff (1960) p. 215

¹⁵² *Ibid*, p. 218

So let us simply bracket the question of which kind of theorist Ziff is, and focus only on what he says that can help us motivate the more plausible type of interest-relational view.

b) The Motivation

Ziff was extremely bothered by what I called the Odd Sounding Judgment Intuition. He noted that it sounded odd to speak of good pebbles, corpses, swamps, or Gila Monsters. On the other hand, there is nothing odd about speaking about good crystals, cadavers, pastures, or cats. This is very puzzling because, after all, the words ‘cadaver’ and ‘corpse’ are two ways of picking out the same object—a dead body. How could the way we describe something change whether or not that thing is good?

He concluded that for something to be good is for it to promote or be conducive to an *interest*. Cadavers can be used for dissection. When we describe a dead body a ‘cadaver,’ we indicate that we are thinking about the body in terms of how it can be used for the purpose of dissection. When we describe that same body as a ‘corpse,’ we do not indicate that we are thinking about the body in terms of any uses or interests at all.¹⁵³

An interest-theory can explain many phenomena that are quite puzzling on the kind-relationalist’s view. First, it can explain why some kinds are, in Thomson’s language, ‘goodness-fixing.’ When the kind term that we use to describe X makes salient an interest that we might promote with X, then that kind term names a goodness-fixing kind. It is the interest, not the kind, relative to which X is good. Further, this explains why so-called ‘functional kinds’ are the most obvious examples of goodness-fixing kinds. Functional kinds have essential functions, and those functions are the salient interest relative to which we judge their goodness. Second, an interest-relational view predicts that if the kind term that we use to describe X does *not* indicate some interest we might

¹⁵³ Ziff (1960) p. 211

promote with X, then we will get the Odd Sounding Judgment Intuition. Since there is no salient interest that pebbles promote, it sounds odd to say ‘That is a good pebble.’ Yet, the interest-relational theory also predicts that we can discharge any oddness from how the judgment, ‘That is a good pebble’ sounds by first contextually indicating an interest that the pebble would serve. As even Thomson admits, if we stipulate that someone is looking for a pebble to fill a hole in the bottom of a flower pot, then there is nothing odd-sounding if she says, upon discovery of a pebble, ‘Aha! Now this is a good pebble!’¹⁵⁴

The most interesting prediction of the interest-relational theory is that even for goodness-fixing kinds, we should be able to contextually shift the standards relative to which we judge members of that kind. As it turns out, we can contextually shift the standards of evaluation for goodness-fixing kinds. Consider the kind: razor.¹⁵⁵ I assume this is a goodness-fixing kind since it is a functional kind, and what it is to be a razor appears to set the standards of evaluation for the kind: a good razor should be sharp. However, suppose that your local theater group is staging a production of *Sweeney Todd, the Demon Barber of Fleet Street*, and it is your job to acquire all the necessary props for the play. One prop that you must acquire is the razor with which Sweeney Todd slits the throats of several of his barbershop’s customers. While shopping for the razor, someone might show you a very sharp razor, and you might truthfully exclaim: “No, no, no! That razor is no good!” Within normal contexts, a good razor is sharp, and the razor that was just shown to you would be good. Within the context of putting on the play, where presumably you do not want to put your actors at risk, a good razor is not sharp.

¹⁵⁴ Thomson (2008) p.22

¹⁵⁵ ‘Razor’ appear on von Wright’s list of examples of ‘instrumental goods’ (1963) p.8.

In Chapter One, I claimed that ‘good’ was a context-sensitive term like ‘tall.’ I suggested that both terms introduce standards of evaluation, and that these standards of evaluation were in some way supplied by context. The interest-relational view explains *how* we can use context to switch the standards of evaluation—we contextually indicate a new interest. What is really exciting about this is that the puzzle of goodness-fixing kinds and of the Odd Sounding Judgment Intuition both turn out to be consequences of the context sensitivity of the word ‘good.’ The interest-relational view thus has quite a lot of explanatory power! We certainly ought to prefer an interest-relational view to a kind-relational one for this reason.

c) Which Interest Theory?

Having now shown that we ought to prefer an interest-relational theory to a kind-relational one, we can now turn to the question: *Which interest theory is most plausible?* As I hope to now show, the Subjective Interest Theory is untenable, and the Hybrid Interest view is really just an Objective Interest Theory in disguise. Consequently, the Objective Interest Theory is the only actually viable alternative to my end-relational theory. Objective theories, however, either conflict with our commonsense judgments about which things are good, or collapse into end-relational views. Either way, they are not more plausible than my end-relational theory of goodness.

My argument against Subjective Interest Theories is brief. The Subjective view makes goodness too dependent on the contingent tastes of people. The following two scenarios seem possible to me, but cannot be accounted for on the Subjective view. First, everyone on Earth could lose interest in things that we would intuitively consider good. It is not hard to imagine a dystopian future where people have lost all interest in

philosophy, literature, and art. According to the Subjective Interest Theory, then, philosophy, literature, and art would thus cease to be good things. That seems plainly false. Philosophy is a good thing, even if no one is in fact interested in it. Second, we can also imagine a future wherein everyone on Earth takes interest in something that is intuitively not good. Suppose that instead of pursuing philosophy, art, etc, people spend their time counting blades of grass.¹⁵⁶ In that world, counting grass would be a good thing to do, according to the Subjective Interest Theory. But again, this seems clearly false. Intuitively, what is good and bad does not depend on the actual interests of people.

The problem with the Subjective Theory is that it makes peoples' attitudes sovereign—they cannot be mistaken about what is good because they determine what is good. Intuitively, however, they *can* be mistaken about what is good. To accommodate this intuition, we might decide to pursue a Hybrid Interest Theory.

According to a Hybrid Interest Theory, a thing, X, is in a person, S's interest (and thus good) if and only if a fully informed version of S would want the less informed version of herself to want X. This is, roughly, the definition for non-moral goodness offered by Peter Railton.¹⁵⁷ The idea is this. A person can make mistakes about what is in her interests due to ignorance about (a) her situation, (b) herself, and (c) the world. To borrow Railton's example, suppose that Sandra is travelling and is not feeling great. She feels like drinking a glass of milk to settle her stomach. The Subjective Interest Theorist would say that her drinking the milk would be good, since doing so satisfies one of her interests. But suppose that what is making Sandra feel ill is dehydration. She does not know this, and she also does not know that drinking milk would aggravate her stomach

¹⁵⁶ This is Rawls' example (1971) pp.432-3.

¹⁵⁷ Railton, (1986) p.176

further. Were she to know these facts, says Railton, she would no longer want to act on her desire to drink milk. A fully informed version of Sandra would not want her uninformed self to want to drink milk. Drinking milk is thus, not in her interest, and on this view, thus not good.¹⁵⁸

I worry that this view is secretly an objective-interest theory in disguise. To motivate this suspicion, consider a slightly altered hybrid view. I call it the Anti-Nemesis View. According to the Anti-Nemesis View, X is in S's interest, and thus good, if and only if, a fully informed nemesis of S would *not* want S to want X. I think that by contrasting Railton's view with my Anti-Nemesis View we will notice a feature of Railton's view that was not explicit, but that was doing a lot of work in his theory. In structure, our views are very similar. They both suggest that we can define goodness in relation to what a fully informed agent would choose for us. However, my view says we should avoid the things that the fully informed agent would choose for us, and Railton's theory says we should favor the things that the fully informed agent would choose for us. Why this difference? Clearly, it is because our nemesis is not looking after our interest—in fact, she wants to frustrate our interests. This reveals a background assumption of Railton's view: that the fully informed version of yourself has your interest at heart. This raises a worry—how do we characterize the difference in preferences between Railton's benevolent fully informed agent, and my malevolent fully informed nemesis? I just put it in terms of 'interest'—but that would render the theory viciously circular (X is in your interest if and only if a fully informed person with your interest at heart would want you to want it).

¹⁵⁸ *Ibid*, pp.174-5

We could, instead, say that my fully informed self wants what's good for me, and my nemesis does not—but this requires that facts about what is good for me are independent and prior to what the informed agent chooses. It is unclear, then, what role is left for the informed agent to play in this theory. It appears that all she does is (a) use her full information to recognize what is good for me, and due to her motivation to promote my interests, (b) she wants that I would want those things that are good for me. We could eliminate the middleman and simply say that there is a list of things that are in my interests, and something is good if and only if it promotes one of those interests. This is precisely what the Objective Interest Theory of goodness says!

I conclude that the Hybrid Theory is really an objective interest theory in disguise. I will now consider the virtues and vices of the Objective Interest Theory.

d) Virtues and Vices of the View

The Objective Interest Theory states that if X is good, this is because it promotes a creature's objective interests, or well-being. Consequently, this view very handily accounts for the connection between goodness and reasons. Nothing fanciful must be said to show that I have reasons to favor X when X promotes my well-being! And insofar as I am motivated to promote my own well-being, learning that X is good would motivate me to favor X. This gives us a very simple account of the connection between goodness and motivation. I think these are the strongest of the interest-relational theory's virtues.

Another virtue of the interest-relational view is that it would be difficult to pin the Too Much Goodness Objection on it. The range of possible good things is actually quite restricted on the interest-relational view. One might wonder, however, if that range is too

restricted. Are there things that are good that are not related to a creature's well-being? If so, the interest-relational theory might fall prey to the *Too Little Goodness Objection*.

For instance, there are a variety of skills that one might possess that would make that person good at thieving. Intuitively, a person with those skills would make a good thief. But, we might notice, thieving does not promote anyone's well-being. So interest-relational theories cannot say there are good thieves. Similar things can be said about good murder weapons, good torture implements, and good doomsday devices.¹⁵⁹ None of these things contribute to creatures' well-beings, so the interest-relational theory must say there are no good murder weapons, etc. But clearly an undetectable swarm of nano-bots that can cause a person to have a stroke would be a good murder weapon. If an interest-relational theory denies that there are good murder weapons, thieves, etc, then it is susceptible to the Too Little Goodness Objection.

When I raised this worry to Kraut in personal correspondence, he agreed that he was too quick to say that there cannot be such a thing as a good thief.¹⁶⁰ He said:

Did I say that there is no such thing as a good thief? If so, I wish I hadn't. What I would like to say is this: There are times when theft is justified. A good thief needs to know when theft is justified, not merely how to accomplish an act of thievery. It may be that, by common standards, all it takes to be a good thief is to have the skills needed to accomplish acts of thievery. But I think our standards should be higher than that. A good thief would know when to thieve, and not merely how to.

Kraut is here distinguishing between a morally good thief, and an efficient thief. And while he says that our standards should be 'higher' than demanding that a good thief is merely an efficient thief, what he does not say is that the 'common standards' are

¹⁵⁹ Finlay also uses the examples of doomsday devices to argue against interest-relational views (2014) p.31. Kraut also discusses the goodness of instruments of torture, though as you will see, not as a counterexample to interest-relational views (2007) p.270.

¹⁶⁰ Kraut (2007) p.270

insufficient. There can be a good thief that is just someone who thieves well. Similar things might be said about murder weapons, doomsday devices, etc.¹⁶¹ There can be good murder weapons if they are efficient at achieving the goal or function of murder weapons. We can distinguish this from the claim that it is morally good that there are murders or murder weapons.

I find Kraut's response unsatisfactory for two reasons. First, pursuing this answer comes at a great theoretical cost. The interest-relational theorist who adopts this view is no longer offering a unified account of goodness.¹⁶² Now she has goodness relative to interests, and goodness relative to functions, without any obvious answer what it is that these two types of goodness share such that it makes sense to call them both 'good.' She thus gives up the Thesis of Commonality. Second, and more worrying for the interest-theorist, is that once we introduce a 'functional' sense of good (and stretch it to cover things like thieves), we now have the theoretical tools available to give a unified theory of goodness *that is no longer an interest-relative theory of goodness*. We might now adopt an end-relational theory of goodness that covers all the things that the interest theory did before. So, things that promote well-being (like health, art appreciation, friendship, etc) might be good because there is an end that is of flourishing, and they promote that end. That is to say, if we extend our notion of 'functional goodness' beyond the few problematic cases that it is meant to resolve –and I see no principled reason not

¹⁶¹ Kraut denies this. He claims that, "A torture instrument... can be a good one, only if torture can be justified in terms of the harm it avoids" (2007) p.270. Again, this sounds to me like it would be a *morally good* torture instrument, but that there are also end-rationally good torture instruments.

¹⁶² Finlay (2014) pp.32-3

to—interest-relational goodness simply becomes a subset of ‘functional goodness.’

Functional goodness is just end-relational goodness.¹⁶³

e) Final Assessment

Although it is hardly a decisive point against the interest-relational theorist, their inability to offer a unified theory of goodness that allows that there are good thieves, murder weapons, etc counts against the plausibility of the view. If a theory of goodness can explain all the cases that the interest theory does *and* can explain why there are good murder weapons, then, other things being equal, this theory would be more plausible than an interest-relational theory. I believe an end-relational theory is just such a theory.

IV. End-Relational Goodness

a) The View and its Proponents

According to an end-relational theory, something is good if and only if it promotes an end. As stated in the previous chapters, I favor an end-relational account of goodness. As such, my view is very similar to that of Stephen Finlay, who also defends an account of goodness where goodness is relative to *ends*. It is worth stressing that his view and mine disagree on some important points. For that reason, I dedicate the next chapter to spelling out the details of both our views, giving special emphasis to the points where they differ. My discussion here will be a fairly brief sketch of the motivation, virtues, and vices of end-relational views in general.

b) The Motivation for the View

¹⁶³ It should come as no surprise that Finlay and I agree on this point (2014) p.32. Anyone offering a unified view of goodness must claim that, deep down, judgments that appear to ascribe a different kind of goodness (relational or not) must *really* be ascribing that theorist’s preferred type of goodness.

End-relational views have all the explanatory power that interest-relational views have. Like interest-relational views, it has a simple way of answering the question that plagued kind-relational views, namely, why is it that certain kinds have good members, and others do not. The answer for the end-relationalist is that some kinds have salient ends associated with them, so good members of those kinds are ones that promote the salient end associated with the kind. A good hammer is one that is good at driving nails—unless we make some other end contextually salient. There are no salient ends for swamps, pebbles, and corpses—so there are no good swamps, pebbles or corpses. Like the interest-relational view, this view predicts that ‘good’ will be context sensitive—we can switch the standards by which someone is evaluating X by contextually indicating different ends relative to which we should judge X.

One issue that is thus central to developing an end-relational view is how we should define ‘ends.’ Finlay and I disagree upon this point. We both agree, however, that the notion of ends is broader than that of interests. Roughly, ends are things that people *could* be interested in.¹⁶⁴ As such, an end-relational theory can explain all the cases that an interest-relational theory can, and more. So for instance, both an objective interest theory and an end-relational theory would preserve our judgment that *Guernica* is a good painting even if everyone stopped being interested in viewing or producing paintings. Unlike the objective interest-relational theory, however, an end-relational theory of goodness also preserves our judgment that there are good murder weapons, thieves, and doomsday devices—even though we all agree that it is wrong to murder, steal, or bring

¹⁶⁴ Finlay (2014) p.32

about the end of the world. The main motivation for adopting an end-relational view, then, is that it is an extremely flexible theory that explains a wide range of cases.

c) Virtues and Vices

One potential objection to end-relational views is that they fall prey to the Too Much Goodness objection. The thought is this. If ends are simply things that people *could* be interested in, then ends are abundant and it will likely be possible to imagine an end relative to which even a plate of mud would be good.¹⁶⁵ But plates of mud are not good, and it is a big bullet to bite if an end-relational view says that they are. That is, even if it is the case that we can think of an end relative to which a plate of mud is good, it is just not the case that a plate of mud is truly *good*. There is a special type of goodness—and plates of mud are not good in this way. Let's call this the *Special Good Things Intuition*.

If my opponent means that a plate of mud is (a) not good for its own sake, or that it (b) does not give us any basic reasons for acting, then I would agree: a plate of mud is not good in those ways. My view would be absurd if it said that everything that is good gives us basic reasons for acting. A doomsday device might be good relative to the goal of destroying the world, but its goodness relative to that goal is no reason to build such a device! So, while there will be at least one goal relative to which a plate of mud is good, it is unlikely that this goal will be one that gives us reasons to act. Plates of mud (and doomsday devices) are just not good in that special way.

Nevertheless, as I argued in the last chapter, my view can accommodate the fact that we will of course expect our theory of goodness to be able to distinguish the good things that we have reasons to adopt, pursue, and favor from those that we do not. My

¹⁶⁵ I pursue this objection in more detail in the next chapter.

view handles this not by claiming that certain things are intrinsically good, but by saying that there is some special set of goals, and the things that are good relative to those goals deserve a privileged position in our practical lives. So, even though on my view all things are good, there is still room to preserve our intuition that there are special good things.

V. Conclusion

The question I have been considering in this chapter is the following: What is the most plausible unified account of goodness? I argue that kind-relational views have problems explaining why we have the Odd Sounding Judgment Intuition when we try to attribute goodness to certain kinds. Kind-relational views will also have difficulty constructing truly kind-relational moral theories. It is natural to think that a kind-relational moral theory will be constructed around the notion of a good person or a good life—which appear to be paradigm cases of kind-relational goods. I argued, however, that to do the work of grounding a plausible moral theory, we need to interpret ‘good person’ or ‘good life’ in non-kind-relational ways.

Interest-relational views have an answer to why we have the Odd Sounding Judgment Intuition when we attribute goodness to certain kinds, but ultimately, they are too narrow. The most plausible versions of interest-relational views of goodness are committed to saying that there are no good murder weapons or doomsday devices. End-relational views, I conclude, are thus the most plausible accounts. I hope to reinforce this conclusion in the next chapter, where I scrutinize two end-relational views.

Chapter Five – My End-Relational Theory of Goodness

I. Introduction

Having settled that an end-relational theory of goodness is the most plausible relational theory, the last task remaining for me to complete is to develop and defend my preferred end-relational account of goodness in greater detail. The most general statement of an end-relational account of goodness is that such a theory claims that when predicated of a thing, ‘good’ ascribes to that thing a particular relation that holds between it and an end.¹⁶⁶ To further elaborate and clarify the view, an end-relational theorist needs to answer the following three questions:

- a) What kind of things can stand in the goodness relation with ends?
- b) What are ends?
- c) What is the nature of the goodness relation?

In the process of answering these questions, I will be comparing my view with that of Stephen Finlay.¹⁶⁷ I focus on his view because it is the most developed end-relational theory that can be found in the literature. Although our views share many similarities, as I read him, Finlay disagrees with me over how we ought to answer each of the above questions. As such, Finlay offers a substantially different end-relational theory of

¹⁶⁶ For ease of discussion, I will simply call this “the goodness relation.”

¹⁶⁷ I had already developed much of my own view when I learned of Finlay’s then forthcoming book (2014) wherein he more fully develops an end-relational account of goodness. Lucky for me, he and I disagree about a few key aspects of how best to develop an end-relational theory. I am truly indebted to him, as his book served as a foil as I worked on the remaining details of my view.

goodness from the one that I prefer. In what follows, I will argue that where our theories disagree, it is my theory that is more plausible.

II. What things can be good?

The first aspect of my end-relational view of goodness that I wish to elucidate is what kind (or kinds) of things can stand in the goodness relation to ends. My inclination is to take seriously the vast variety of objects that feature in our goodness judgments. We judge all of the following as good: physical objects (e.g. knives, umbrellas, parking spaces, diamonds), people, other living things, places, actions, events, states of affairs, abstract entities such as properties or relations (e.g. beauty, justice, being loved), imaginary characters, possible worlds, etc. In short, there seems to be no restriction on what kind of things to which we can ascribe goodness. I thus see no reason for my theory of goodness to limit what types of things can stand in the goodness relation to ends. In short, I understand ‘things’ as broadly as possible when I say that ‘*good*’ *ascribes to a thing a relation that holds between it and an end*.

In fact, my account of goodness not only says that all kinds of things *can* be good, but says that all things *are in fact good* (in some way or another). This might seem counterintuitive, or worst still, it might invite the complaint that my account of goodness falls prey to the Too Much Goodness Objection discussed in the last chapter. Surely there are things that are not good! Some things, like plates of mud, do not seem to be good because they serve no known interests.¹⁶⁸ Other things, like doomsday devices, do not

¹⁶⁸ G.E.M. Anscombe introduced the example of a “saucer of mud” (1957) pp.70-1.

seem to be good in any way because they are contrary to all interests.¹⁶⁹ But we need to remember that end-relational goodness comes cheaply. If we can point to an end that a thing promotes, then that thing is good in relation to that end. Because end-relational goodness comes cheaply, accepting that all things are good in some way does not require biting any bullets. In fact, there is a quick and clever argument that proves that goodness does come cheaply, and thus all things are good: If it were true that something (a plate of mud or a doomsday device, for example) was not good in any way, it would be a good counterexample to the claim ‘All things are good in some way.’ So it would be good in a way. Thus, it is impossible for something to fail to be good in a way. That is, all things are good.¹⁷⁰

I take it that when someone claims that a plate of mud or doomsday device is not *really good*, this is not because she denies that there are ends that those things promote. She could grant me that.¹⁷¹ Her concern might remain: that there are some things that we ought to favor or promote because they are good, and plates of mud and torture devices are not good *in that way*.¹⁷² The Too Much Goodness Objection, if it is to have teeth, must concern itself with the special subset of the set of good things. It must say that it is objectionable if a theory of goodness is too permissive in regards to ascribing this special kind of goodness to things. My view does not commit me to the claim that everything is good in this special way. In fact, my view does not commit me to any definitive answer as to which things are good in this special way. At the end of Chapter Three, I sketched

¹⁶⁹ Finlay also uses this example to make this point (2014) p.31.

¹⁷⁰ Stephen Finlay, Judith Thomson, and I have each independently arrived at a version of this argument. See Finlay (2001), p.73 and Thomson (2008), p.10.

¹⁷¹ A person who knew enough about doomsday devices to say that they are not good in any way presumably understands what end doomsday devices serve. She just believes, correctly, that the destruction of all life on Earth is a horrible goal for one to pursue.

¹⁷² Richard Kraut raised the concern in this way in personal correspondence.

one way my view might account for the fact that some ends are worthy of our promoting them, while others are not. That sketch provided a framework from which we could discuss the substantive question of what is good in that special way. It did not, however, commit me to any substantive answers to the question. This is to say that my view does not fall prey to the Too Much Goodness Objection.

Finlay disagrees with me about what the bearers of goodness are. He argues that the bearers of goodness are not things in the broadest sense of the term, but *states of affairs*.¹⁷³ Finlay arrives at this conclusion because he is interested in finding a deep unifying syntax that is shared by all sentences that predicate goodness to something.¹⁷⁴ He notes that two of the most common types of sentences used to express goodness judgments, however, appear to be too different to share a unifying syntax:

(Type 1) X is good.

(Type 2) It is good that X.

The problem for a unified syntax is that X appears to be the subject of sentences of Type 1, but not of Type 2.¹⁷⁵

To quickly summarize Finlay's discussion, he argues that sentences of Type 1 are elliptical for "X is good for S_a to φ ," where 'S_a' refers to an agent, and ' φ ' represents an action.¹⁷⁶ This interpretation captures the intuitive idea that to say that X is good is to say that X is *good for someone to do something with*. Finlay proposes that the phrase 'For S_a to φ ' is itself elliptical for a complete sentence. It is elliptical because ' φ ' is an

¹⁷³ Finlay, (2014) pp.26-8

¹⁷⁴ *Ibid*, p.22

¹⁷⁵ *Ibid*, p.26

¹⁷⁶ *Ibid*, p.26

incomplete predicate that requires a noun phrase to complete it.¹⁷⁷ The noun phrase that completes the predicate, he argues, is X itself.¹⁷⁸ For instance, on this picture, “Chocolate is good” is elliptical for “Chocolate is good for a person to eat [chocolate].”¹⁷⁹ Eating is always the eating of something, but we do not normally say the word ‘chocolate’ twice because it is clear from context that the object of the action of eating is chocolate. So in general, Type 1 sentences can be analyzed as having the following form:

G1: X is good for S_a to φ X¹⁸⁰

Finally, he notes that many English sentences include a ‘dummy subject’ like ‘it’ (e.g. “It is raining”), which simply serves as a placeholder to render the sentence grammatical. The presence of the dummy subject does not contribute to the meaning of the sentence. For this reason, some English sentences feature subject-movement—where the subject of the sentence moves away from the front of the sentence and is replaced by a dummy subject—with no change in the meaning of those sentences.¹⁸¹ For instance, the sentence “Bill ate a sandwich” can be expressed as “It is the case that Bill ate a sandwich” with no change of meaning. He thus proposes that we insert such a dummy subject into the analysis of Type 1 sentences give above, resulting in:

G2: It is good for S_a to φ X¹⁸²

¹⁷⁷ *Ibid*, p.21

¹⁷⁸ *Ibid*, p.26

¹⁷⁹ One might notice that nothing about the target sentence, ‘Chocolate is good,’ indicates that we should substitute the word ‘eating’ in for ‘φ’. However, later in the book, Finlay gives a pragmatic explanation of why we can assume ‘Chocolate is good,’ is about *eating*, that is, unless some other clues in the context make another action (smearing on one’s face, for instance) more salient.

¹⁸⁰ *Ibid*, p.25

¹⁸¹ *Ibid*, p.27

¹⁸² *Ibid*, p.27

This analysis, he claims, applies to sentences of Type 2 as well—thus giving us a unifying syntax for both types of sentence.¹⁸³ But how well does it really handle sentences of Type 1? According to this analysis “Chocolate is good” has the same meaning as “It is good that one eats chocolate.” But that does not seem right. The target sentence, intuitively, ascribes goodness to *chocolate*. The proposed paraphrase, intuitively, ascribes goodness to an activity that people might do. The dummy subject appears to be altering the meaning of the sentence—indicating that the analysis does not adequately handle sentences of Type 1.

Finlay predicts this objection, and adds some more machinery to his analysis to handle it.¹⁸⁴ I will discuss this machinery, and Finlay’s reply, when I give his final analysis in the next section. What is important for the current discussion is that according to Finlay’s analysis, goodness is predicated not of X itself, but of the potential state of affairs of S_a’s φ -ing of X. According to Finlay, it is potential states of affairs, not things in general, that are the bearers of goodness.

I find it a bit surprising that Finlay claims that only states of affairs can be bearers of goodness, especially in light of the fact that he offers a version of the reductio argument I gave above to show that all things are good in some way. According to his view, it is not *literally* true that all things are good in way (because only states of affairs can be good). So it is odd that he offers a reductio argument to support the claim that all things are good in some way.

Perhaps because he recognizes this tension, Finlay addresses the putative goodness of things in an appendix. There, he argues that his view can “accommodat[e]

¹⁸³ *Ibid*, p.27

¹⁸⁴ *Ibid*, p.27

the common intuition that the goodness of objects is often prior to and explanatory of the goodness of states of affairs.”¹⁸⁵ His claim is, roughly, that that the goodness of a state of affairs, A, depends on how it relates to some relevant contrast class.¹⁸⁶ Often, there will be some choice over which dimensions of comparison it is relative to which we are evaluating A. Suppose, for instance, that we judge that (1) it is good that Jones uses that hammer, to drive those nails, through that kind of lumber. If we express judgment (1) with the sentence, “That is a good hammer,” we signal that the relevant contrast class is the one where Jones is still interested in driving those nails through that kind of lumber, but uses a tool other than that hammer to do it. Alternatively, we could express (1) with the sentence, “Those are good nails.” Our choosing to express (1) this way signals that we mean that it is good that Jones uses those nails, and not some other nails. In both these cases, it was the presence or the absence of an object that made the difference between the state of affairs that we are calling good and the contrast class. We can thus explain the difference between these states of affairs by appealing to features of the object—the so-called ‘goodness-making’ features. Appealing to those properties of the hammer thus explains why it is good that Jones uses that hammer. This, says Finlay, accommodates our intuition that it is the hammer, and not Jone’s use of the hammer, that is good.¹⁸⁷

I agree that Finlay’s discussion goes some way towards accommodating the intuition that objects can be bearers of goodness, but does it go far enough? We can actually use the original *reductio* argument from above to create a new argument against Finlay’s claim that only states of affairs can be good. Suppose, for the sake of a *reductio*, that it were the case that only states of affairs could be good. If that were so, then

¹⁸⁵ *Ibid*, p.46

¹⁸⁶ More on this point in Section IV.

¹⁸⁷ *Ibid*, pp.46-7

ordinary objects like plates of mud could not be good. It follows that a plate of mud would be a good counterexample to the claim that ‘All things are good in some way.’ Thus the plate of mud is good in a way: as a counterexample. The claim that only states of affairs can be good leads to a contradiction, and thus, should be rejected.

Finlay would likely reply that a plate of mud is not literally good, nor is it literally good as a counterexample. It is the state of affairs of my using a plate of mud as an example that we can evaluate as either good or not. But this reply blurs the distinction between the thing that is the counterexample, and the event of someone *giving* that counterexample. This distinction is important; there are times when you have a good counterexample, but it is not good for you to *give* that counterexample (for instance, in a quarrel with one’s significant other, even if one has a very effective counterexample to something her significant other said, often it will only aggravate the quarrel further if one gives that counterexample). A counterexample can be good, and so can the giving of a counterexample. On the assumption that only states of affairs are good, however, a plate of mud is literally not good in any way. It would thus be a good counterexample to the claim that ‘All things are good in some way,’ *even if no one ever gives it as a counterexample*. The contradiction remains, and thus, we ought to reject the claim that only states of affairs can be good.¹⁸⁸

This argument can be generalized. Suppose, for the sake of a *reductio*, that only some restricted class of things can be good. Pick a thing, X, that falls outside of that

¹⁸⁸ There is a popular view that claims that states of affairs are the only bearers of *intrinsic* goodness. See, for instance, Chisholm (1986), Audi (1997), and Zimmerman (2001). I deny that anything is intrinsically good (Chapter Three). However, if we understand their claims as being about a special kind of end-relational goodness that we should take interest in and promote, then my view is compatible with theirs. I do not say anything about what kinds of things can be the bearers of that special kind of goodness. I allow that states of affairs can be good, so it remains possible on my view that only states of affairs have this special kind of goodness.

class. Thus, X is not the kind of thing that can be good, so it is not good in way. If X is not good, then it is a good counterexample to the claim that ‘All things are good in some way.’ If X is a good counterexample, then X is good in some way. We have a contradiction, so we need to reject our initial assumption. That is, it is not the case that only some restricted class of things can be good.

I conclude that it is a mistake to claim that there are any restrictions on the types of things that can be good. Finlay’s account of end-relational goodness places restrictions on what things can be good; mine does not. This is one reason to favor my view over his.

III. What are ends?

It is with regard to how we define ends that Finlay’s and my accounts of goodness are most similar. Finlay uses ‘end’ as a “term of art for any proposition conceived as a potential outcome.”¹⁸⁹ To put it more succinctly, ends are potential states of affairs. I too think that ends are potential states of affairs. However, as I understand him, an end can be any state of affairs. This is not so for ends as I understand them. To illustrate the way our views come apart, I need to talk first talk a bit about two different kinds of attitudes: desires and intentions.

Both desires and intentions have objects—that is, a desire is always a desire *for something*, but an intention is always the intention *to do something*. Thus, intentions are different than desires. How so? Where desires differ from intentions, I contend, is in their satisfaction conditions; a desire is easier to satisfy than an intention. Consider for

¹⁸⁹ Finlay (2014), p.32

instance, that I have the desire that my friends Alex and Casey start dating, and I intend to bring it about that they start dating. My plan is to invite both to a dinner party, sit them next to one another, and steer the conversation towards topics I know they both share an interest in. Days before my party happens, though, Alex and Casey meet in the Science Fiction section of a used bookstore. They begin chatting about their favorite Sci-Fi authors, and eventually move to discussing things over coffee, and then dinner. After dinner, Alex gives Casey her phone number; they have started dating. My impression of this story is that this is a scenario where my desire, that Alex and Casey start dating, has been satisfied. My intention, however, has not. I planned on having a hand in bringing these two friends together—but I played no role in it.¹⁹⁰

It might be pointed out that the reason my intention is more difficult to satisfy is that my intention was that *I bring it about* that Casey and Alex start dating, while my desire was merely that they start dating. Had my intention been merely that they start dating, it would have been satisfied alongside my desire. This reply, however, threatens to erase any distinction between desires and intentions. Allow me to put it in another way: I have lots of desires, some of which are incompatible. For instance, I desire to finish this chapter, but I also desire to go out with friends. It is not possible that I do both (at least, not this evening), so I must choose which desire I will act on. Once I have chosen to finish this chapter that becomes my intention. It is my intention because I plan on acting in such a way as to bring about that the state of affairs that I desire obtains. I contend that intentions are always intentions to *bring about that P*, whereas a desire can

¹⁹⁰ I am drawing extensively from Michael Bratman's so-called 'planning theory of intentions.' According to Bratman, intentions are a distinct type of mental state that is used to direct future action. Having an intention to X involves having a plan (even a partial plan) for how one will go about doing X. This view is developed in (1987).

simply be for *P*. Thus, even if I say that “I intend that Casey and Alex start dating,” this should be understood as elliptical for, “I intend to *bring about* that Casey and Alex start dating.”

What should we say if instead we considered the case where I have both the intention and the desire to bring it about that Casey and Alex start dating? Again, the desire is easier to satisfy because the intention involves some plan of action. Returning to the example, suppose that both Casey and Alex went to the used bookstore due to my recommendation of the Sci-Fi section there. It is mere coincidence that they go at the same time, meet, start to chat, and (eventually) begin dating. My desire to bring it about that they start dating has been satisfied. Again, my intention has not. I intended to bring it about that they start dating *by way of a plan* (that they come over for dinner, etc). Intentions necessarily involve plans for how to bring about the object of the intention—desires do not.

I can now explain the difference between my account of ends and Finlay’s. On Finlay’s account, ends are states of affairs. Every end, in Finlay’s sense, can be the object of a desire. However, it is not the case that every end in Finlay’s sense can be the object of an intention. Some of the ends that he has in mind are simply outcomes—they are just that certain events occur. The ending of the war in Iraq is an end on Finlay’s account. As I wish to define them, ends are only those states of affairs that can be the object of an intention. My view would thus say that the ending of the war in Iraq is not an ‘end’; I cannot intend that the war in Iraq ends. I can, however, intend to do my part in bringing it about that the war in Iraq ends. On my view, my bringing it about that the war ends is an end, but the war’s ending is not. Perhaps this is why I prefer the term ‘goal’ over the

word ‘end’ as a way to refer to the relevant notion. On my view, ends involve plans for action—and ‘goal’ seems to better capture that. Regardless, I will continue to ‘goals’ and ‘ends’ interchangeably—what I mean by these terms is the possible object of an intention.

There is not much to say in this section about why we ought to think of ends as plans for action instead of mere outcomes. As we shall see, the way that I characterize ends does play an important role in explaining the nature of the goodness relation. For the most part, I believe that the reason to prefer my characterization of ends over Finlay’s is due to the way it supports my account of the nature of the goodness relation. There is, however, one minor point that I believe weighs in favor of my account of ends.

On Finlay’s view, since ends are just states of affairs, and since the bearers of goodness are states of affairs, it turns out that one and the same end can be good relative to itself. Goodness is sometimes a reflexive relation. My drinking a beer, for instance, is good relative to the goal of my drinking a beer. That is a strange consequence of Finlay’s characterization of ends.

Finlay argues that this consequence is not a bullet to bite. Quite to the contrary, the fact that e promotes e explains the locution, said of an end, that it is ‘good for its own sake.’ Of course, if that is what it means for something to be good for its own sake, we have the bizarre consequence that every e is good for its own sake! Why then, would we ever talk about something’s being good for its own sake (if it is trivial that everything is good for its own sake)?¹⁹¹ Finlay’s answer, again, is pragmatics. To say that an end is good for its own sake can signal that that end is one that you desire for its own sake, and not for the sake of some further end. We also sometimes talk of an end being good for its

¹⁹¹ Finlay discusses this objection on *Ibid*, p.202

own sake to *moralize*—to convince others that they too ought aim at this end.¹⁹² For the most part, I find these pragmatic explanations compelling—but I cannot help but shake the feeling that there is something odd about the fact that on Finlay’s view, an end can be good relative to itself.

For what it is worth, my view does not have that consequence. As I understand ends, they are abstract plans. Intuitively, an abstract plan cannot promote the achievement of itself. My having the plan, my acting on the plan, etc can bring about the achievement of the plan—but the plan itself cannot. So, on my view, nothing is good for the sake of itself in the way that Finlay discusses.

This is not to say that my view does not have its own trivial judgments about goodness. Say my goal is very simple—waving ‘Hello’ to a friend. The plan would consequently be simple as well: the best way to promote this goal is for me to simply wave hello to a friend. It sounds a bit silly to say that waving hello to your friend is good relative to the goal of waving hello to your friend, but we have to be clear about what this sentence really means. According to my view that sentence expresses the proposition: the activity of waving hello promotes (in the relevant respect) the achievement of the plan to wave to your friend. This is long-winded; but not so very absurd.

Again, the criticism of Finlay’s account of ends is mostly a quibble. Where our accounts of ends gain or lose plausibility, is in which answers they support to the question, ‘What is the nature of the goodness relation?’ I turn to that question now.

¹⁹² *Ibid*, pp.202-4

IV. What is the nature of the goodness relation?

The greatest difference between my view and Finlay's is how we characterize the goodness relation. I contend that his account of the goodness relation gets things wrong in certain scenarios. That is, his view commits him to (a) saying that some things are good relative to G, when intuitively, they are not good relative to G, and (b) saying that some things are not good relative to G, when intuitively, they are good relative to G. This is the primary reason I believe we ought to prefer my view to his.

That we disagree about the goodness relation should not be too surprising. There is a lot of room for disagreement on this issue, given that for any two relata there are many (innumerable?) relations that can hold between them. The end-relational theorist believes that one of these relations that can hold between two (or more) relata is the goodness relation, but which one? What is the nature of this relation? That is, what differentiates this relation from all the other relations up there in 'Plato's heaven'? I will consider three plausible candidates.

a) Achieving *e*

The first answer that might spring to one's mind, given that we are talking about a way in which things relate to ends, is that the goodness relation is that of *achieving some end*. According to this view, to say that X is good is to say that X achieves some end *e*.¹⁹³ This certainly seems to get things right in a large number of cases: A good steak knife is one that achieves the end of cutting meat into bite-sized pieces. A good paperweight is one that achieves the end of keeping papers on one's desk, even if a significant breeze blows over it. Examples abound.

¹⁹³ In earlier versions of his theory, Finlay defended an achievement view of goodness. He called it a 'satisfaction' view. See Finlay (2001), p.74 & (2004), p.214.

But for a variety of reasons, this relation is too strong to be the goodness relation. For one, there are many things that achieve ends, but are intuitively not good relative to those ends. I can cut a steak up with a chainsaw—but the result is not pretty. Similarly, I could, with enough time and patience, fell a redwood tree using only a butter knife. Intuitively, however, a butter knife is not good for cutting down redwood trees, nor is a chainsaw good for cutting steak.

Second, there are things that merely *contribute* to the achievement of ends that are not themselves sufficient for the achievement of those ends (partial causes, tools, background conditions, etc) that are nonetheless good relative to these ends. For instance, a mild, sunny day might be good relative to the end of our enjoying a bike ride—but it is certainly not sufficient. Without a bike our ride would not happen, no matter how pleasant the weather is!

Third, it also seems clear that many things are good, even if they never in fact achieve or help achieve *e*. The following story seems perfectly unobjectionable:

Samurai: A virtuous samurai was awarded, in recognition of his virtue, the finest sword in the world. It had all the virtues of a top-quality sword: it was extremely sharp, light, and could be swung swiftly with ease. The virtuous samurai, however, never engaged in battle, preferring instead to use well-reasoned arguments to convince his enemies of the errors of their evil ways. He was so successful that he never used the sword.

Despite the fact that the sword is never used in battle, it remains true that the sword is good. But if goodness is the relational property of achieving *e*, this sword is not good.

Similarly, suppose there is a sword very much like the one described in *Samurai*, except with one minor flaw. There is a tiny, imperceptible imperfection on the blade that, if struck just right with a very hard material, would snap the blade in two. The odds of this happening are extremely unlikely. In an overwhelming number of cases, then, this

blade would achieve *e*. Isn't this still a good sword? I believe that we should say that it is. If so, then we should also be willing to say that it is good, even if due to some outrageously bad luck someone were to strike the weak-spot in just the right way and split the sword in two. In that scenario, it would be a good sword that just so happens to fail to achieve *e*. But if goodness is the relational property of achieving *e*, this sword is not good.

Finally, goodness comes in degrees; achievement of *e* does not. We frequently compare the goodness of things, ranking one thing as better, worse, or equally good as another.¹⁹⁴ The achievement view can make sense of two things being equally good: they both achieve *e*. But what can this view say about the claim that X is better than Y? Either a thing achieves *e* or it does not. So if it were the case that X achieves *e*, and Y does not, then it would be the case that X is better than Y. After all, on the view we are considering, if X achieves *e*, it is good; if Y fails to achieve *e* then it is not good. Being good is better than being not good. But what if both X and Y achieve *e*? They would thus both be good, but how could we ever explain why one is *better*? Intuitively, a chainsaw is better than a butter knife for felling trees, even if both achieve the goal. The achievement view does not have the moving parts to explain this.

For all these reasons, I contend that we should not identify the goodness relation with that of achieving goals.

b) Promoting *e*

¹⁹⁴ In "The Possibility of Parity", Ruth Chang claims that there is a further relation, being on a par with, which is distinct from being equally good as (2002) p.661. Whether there is such a relation or not is neither here nor there for our purposes. An achievement view fails to explain how one thing can be better than another. That is surely reason enough to dismiss the achievement view without inquiring into how it handles the relation of being on a par with.

According to Finlay, the goodness relation can be described as the relation of promoting. He says:

Being good for e is approximately *being promotive of e* ... What does it mean to be good for or promote some end e ? An intuitive answer is that p is good for (or “promotes”) e if and only if p increases the probability that e . (Stephen Finlay, *Confusion of Tongues*, 2014 p.38)

Finlay admits that his account requires taking a stance on a philosophically vexing question: How best should we understand probability? I have no horse in that race, so I will not raise serious criticisms of his account for how it understands probability. I will, however, simply draw attention to this as a potential weakness of the view. If we can explain the nature of goodness without taking sides on other philosophically vexing issues, we probably ought do that.

According to Finlay, to promote e is to increase the probability that e . But increase it from what? We need a contrast class relative to which we determine whether the probability of e increases when p occurs. Finlay first suggests that the contrast classes be stated in terms of possible worlds. He suggests that if e obtains in a greater proportion of possible worlds where p occurs than it does in possible worlds where p does not occur, then p is good. If not, p is bad.¹⁹⁵ Of course, there are some difficulties here due to familiar features of possible worlds. There are an infinite number of possible worlds where both p and e obtain, but also, an infinite number of possible worlds where *not- p* and e obtain. How then, could we speak meaningfully of e occurring in a greater proportion of p worlds than in *not- p* worlds? Finlay’s response is to switch to discussing possible world *types*.¹⁹⁶ At first glance, this does not seem to resolve the issue. Aren’t there an infinite number of possible world types as well? For instance, there is a set of

¹⁹⁵ *Ibid*, p.39

¹⁹⁶ *Ibid*, p.40

possible worlds wherein I win the lottery and win \$10. Then there is another set where I win \$11. There is another where I win \$12, and so on. Presumably, what Finlay has in mind is that we should not distinguish types of possible worlds in such a fine-grained manner. If we categorize possible worlds in a fairly general way, there will not be an infinite number of possible world types. Consequently, we should be able to evaluate whether e obtains in a greater proportion of possible world types where p than where *not- p* .


It is also worth noting that the contrast class should not include *all* the possible world types where *not- p* occurs. In asking if an additional sheet of paper, stacked onto a pile of papers, would make a good paperweight, I am not interested in worlds that have radically different physical laws than ours, or where the Earth's gravitational force is 1,000 times stronger. On those worlds, sheets of paper might make excellent paperweights. But that is irrelevant whether a sheet of paper would make a good paperweight in the actual world. In order to restrict the *not- p* worlds to only the relevant ones, Finlay proposes that we only include *not- p* worlds that are consistent with some set of propositions, B , that is true of the actual world.¹⁹⁷

Let's look at one of Finlay's examples.¹⁹⁸ There is a baseball player, let's call him Casey, and he is at the plate. We want to know if it would be good, relative to Casey's team winning the game, if he drew a walk. So in this case, p is Casey's drawing a walk, and e is Casey's team winning the game. We now ask: Does p increase the probability of e ? Again, to answer this, we need to put some constraints on the contrast class of *not- p* . For surely there are possible worlds where the team's batting line-up is different, where

¹⁹⁷ *Ibid*, pp.40-1

¹⁹⁸ *Ibid*, p.40

the rules of baseball are different, or even where the rules of physics are different. These worlds will not help us in determining if p is good. We also would exclude worlds where the baseball game that has been played prior to Casey's turn at bat unfolded differently. Casey's taking a walk might be a good strategic move, but not in all situations. This is just to say that we want the worlds that we are looking at to be very 'close' to the actual world. To this end, we come up with some set of propositions, B , that specify some of the constraints we just noted (the laws of physics are such-and-such, the game played out in the following manner, etc). We then look only at the world types that are consistent with B , and where $not-p$. In some of those worlds, Casey will strike out. In others, he might hit a line drive or even a homerun. These are the worlds that are relevant when we are determining whether or not p would increase the probability of e . Having identified our contrast class, we can now draw a diagram like this one that follows:¹⁹⁹



Finlay explains that the shaded shape labeled 'e' represents the world types in which e obtains. One of the ovals in the large oval represents the worlds in which p obtains (the one labeled 'p'), while the others represent "various possible ways of realizing $not-p$ consistent with B"²⁰⁰ These ovals come in different sizes to represent that some types of worlds are more likely than others. It is more likely that Casey strikes out than that Casey

¹⁹⁹ The following figure is a slightly adapted version of FIGURE 2.1 from *Ibid*, p.45.

²⁰⁰ *Ibid*, p.45.

hits a home run. So perhaps the littlest oval (that is all grey) represents worlds where Casey hits a home run,²⁰¹ and the largest inner oval that is only slightly shaded would represent worlds where Casey strikes out. To determine whether p is good, we simply compare the shaded proportion of the p oval with the shaded proportion of the other ovals. In this case, roughly 75% of the p -type worlds are worlds where e obtains, and roughly only about 25% of the not- p type worlds are worlds where e obtains. So p increases the probability that e , and is thus, good.

One problem I have with this view is that to determining the sizes of those smaller ovals would push us right back into the problem of probability over worlds. How do we determine how much more likely world types are where Casey strikes out than world types where Casey hits a home run? Intuitively there is no problem here, but when you remember that there are an infinite number of B consistent worlds where Casey strikes out, and an infinite number of B consistent worlds where he hits a home run, it is hard to see how we can claim one type of world is more likely than the other.

Suppose that we grant Finlay that these difficulties regarding probability across possible worlds can be overcome. Even so, I think that he is mistaken to define ‘good’ in terms of increasing the probability of e . As I will now show, this view gives us the wrong answer in certain cases.

To set up my counterexample, first consider a case that I will call *Ball Bearing One*. Suppose that a certain e requires the use of a ball bearing. That is, if I do not use a ball bearing, the probability that e obtains is 0. I am handed a ball bearing, B_1 , and unbeknownst to me, microscopic imperfections in the surface of B_1 cause it to have a

²⁰¹ For the oval to be completely shaded would mean that Casey’s team is guaranteed to win in that type of world. Thus, it must be the case that the game is tied and it is the bottom of the ninth or of a later inning. Otherwise, a homerun would not guarantee Casey’s team winning.

15% chance of failing to bring about e . Is B_1 good relative to e ? Finlay's answer is that B_1 's goodness depends on whether using B_1 increases the probability that e obtains. Whether or not using B_1 increases the probability of e obtaining will depend on what the alternatives are to using B_1 , how likely those alternatives are, and in how many of those alternatives e obtains.

Since B_1 has a 15% failing rate, we can expect that e will obtain in 85% of the p worlds (that is, the worlds where I use B_1). Looking at that number in isolation, one might come to conclude that B_1 is good. It is very likely to get the job done—what more could you want? However, Finlay would insist that B_1 's goodness depends on what happens when I do not use B_1 . Thus, we do not yet have enough information to judge the goodness of B_1 ; we cannot evaluate the goodness of B_1 until we know if using it *increases* the probability of e . That is, we cannot know if B_1 is good until we consider the contrast class.

So let's add some details to the story. Suppose there is a bin with 9 other ball bearings in it that I could have chosen instead. Suppose further that all of them have the same surface imperfections as B_1 , and thus also have a 15% chance of failing. In evaluating whether B_1 is good or not, it seems that there are 10 relevant possible world types against which I must contrast the probability of e obtaining (the worlds where I use B_2 , the worlds where I use B_3 , ... B_{10} , and the worlds where I do not use any ball bearing at all). Nine of these world types are equally likely, and e will obtain in 85% of the these *not- p* worlds. So far, p worlds are doing equally well as *not- p* worlds. Because my intent is to bring about e , because I know that doing so requires a ball bearing, and because I have several ball bearings on hand, the tenth type of *not- p* world—the one where I do not

use any ball bearing at all—is extremely unlikely. However, that tenth type of possible world makes the difference. In those worlds, *e* never obtains. Thus, if *e* obtains in 85% of the *not-p* worlds where I use a ball bearing other than B_1 , and obtains in 0% of the possible worlds where I do not use a ball bearing at all, then *e* would obtain in just slightly less than 85% of the total *not-p* worlds. Recall that *e* obtained in exactly 85% of the *p* worlds. Thus, B_1 (barely) increases the probability that *e* will obtain, and is thus, good. That seems like the right result: B_1 is a good ball bearing.

But what if we change things slightly? Let us call the new case, *Ball Bearing Two*. This scenario is similar to *Ball Bearing One*, except that this time, the ball bearings in the bin have from a 1% to 30% fail rate. The distribution of these ball bearings is uneven, but the average fail rate of the bin is 14%. As before, *e* obtains in 85% of the worlds where I use B_1 . Also as before, *e* obtains in 0% of the worlds where I do not use any ball bearing at all. This type of world, however, is extremely unlikely. For the other *not-p* worlds, where I use ball bearings B_2 - B_{10} , the number of worlds where *e* obtains will vary. However, taken together, *e* obtains in 86% of those possible worlds. The probability of *e* in the *not-p* worlds is now slightly greater (just under 86%) than the probability of *e* in the worlds where I do use B_1 (85%). Finlay would have us conclude from this that B_1 is *not good relative to e*. But that is surely counterintuitive! B_1 only has a 15% fail rate!

Finlay's view is a sophisticated example of a broader type of account of the truth conditions for goodness judgments that I will call a *comparative account*. On this type of view, being good involves being better than some contrast class. The most naïve view of this sort says that X is good if it is better than most (or perhaps, merely 'enough') of the

other members of a kind of which X is also a member. X is a good ball bearing, the view would say, if it is better than the average ball bearing (or perhaps, merely better than enough other ball bearings).

A proponent of this view could agree that in *Ball Bearing Two*, B_1 is still a good ball bearing because there are some ball bearings that have a fail rate as high as 30%, and B_1 is better than them. On the other hand, were I to have stipulated that all the ball bearings in the bin had a fail rate of 1 - 15% then it would no longer be true that B_1 is a good ball bearing. It would be worst than enough of the others, so Finlay's account would be correct in claiming that it is a bad ball bearing.

But suppose that all the ball bearings in the bin are very fragile and have a 99% fail rate! B_1 , however, has a fail rate of 98%. Since B_1 is better than the average ball bearing, the naïve comparative account would have to say that B_1 is good. Finlay's account would say that this ball bearing is good as well, since the probability that e obtains is better (albeit, only slightly so) when I use B_1 , than when I do not. Even though we can agree that (i) B_1 is better than all the other ball bearings, and that (ii) B_1 raises the probability of e , intuitively, (iii) B_1 is not good. B_1 is extremely fragile and very likely to fail. The proper thing to say about this scenario is that this is a world where all the ball bearings are bad. A similar story can be given where all the ball bearings in the world are good, even if one ball bearing has a slightly higher fail rate than the rest (for instance, a world where all ball bearings have a 1% fail rate, except for B_1 , which has a 2% fail rate).

Comparative views go wrong because they confuse the relational property of being good relative to e with the relational property of being good for a K, (where 'K' stands for a kind). I think it might make sense to say that when B_1 has a 98% fail rate it

is good *for a ball bearing* (in the world where all the other ball bearings have a 99% fail rate). But if we want to accomplish an *e* that requires a non-failing ball bearing, then B_1 is not good relative to this goal. In fact, it is very bad. The fact that being a good *K* and being good relative to *e* are two distinct relations should be familiar turf for anyone who has discussed politics with a disillusioned voter.²⁰² Such voters will often say that they are voting for the ‘better of two evils.’ I take this expression to mean, roughly, that within the category of electable candidates there are only two options, and while one is a better candidate than the other, neither one is good. If such a complaint makes sense, then we must reject naïve comparative accounts of the truth conditions for goodness judgments.

In summary, my complaint is that comparative accounts, including Finlay’s, fail to capture the intuitive notion that we can judge *X* to be good, independently from a judgment about the goodness of other things (even other things of its kind). We do not need to know how well other things can accomplish *e* to know if *X* is good relative to *e*.

It is worth stressing that this objection to Finlay’s view is potentially quite damaging. Above, my objections were more or less that my view was more intuitive or natural than Finlay’s in certain respects. But, as noted at the outset of the chapter, because Finlay’s end-relational account of goodness is just one aspect of his larger end-relational theory of normativity, he can allow that certain aspects of the theory be implausible or counterintuitive. He can allow this so long as these costs are made up for by other theoretical virtues of the entire theory (that it is simpler, for instance). This kind of response is not as compelling to my objection from this section. My objection here shows

²⁰² Or who has been a disillusioned voter herself!

that Finlay's theory gets things wrong in a number of cases. If X is really lousy, but slightly less lousy than the other options, Finlay has to say that X is good. If Y is more than adequate at accomplishing *e*, but all the other options are sufficiently better, then Finlay has to say that Y is bad. Both these judgments are mistaken. Theoretical virtues like simplicity do not go far in outweighing the costs of accepting a theory that gets things wrong in many cases.

c) Sufficient to Achieve a Necessary Sub-Goal of G

I believe we should abandon discussion of contrast classes and increasing probabilities. Instead, we should return to the intuitive, but ultimately mistaken view, that the goodness relation is that of achieving *e*. We should see if we can salvage this view with the addition of some further moving parts. I believe we can.

Let us start by revisiting the problems with the achievement of *e* view. According to this view, X is good if it achieves *e*. I noted four kinds of cases where this gets it wrong: (1) Where X is not good relative to *e*, but X achieves *e*, (2) Where X cannot achieve *e*, but it is good relative to *e*, (3) Where X is good relative to *e*, but is never used to achieve *e*, and (4) Where X and Y both achieve *e*, but intuitively, one is better than the other. I believe that with a few adjustments, mostly to how we conceive ends, we can protect a kind of achievement view from these objections.

The first thing we need to do is to switch from talking about X actually achieving *e*, and instead talk about X being sufficient for achieving *e*. This move is needed given how I understand ends. As I describe them, ends are abstract. There are a plentitude of ends that we have never even thought about, and consequently, these ends have never been achieved. However, I want it to be the case that there are still truths about what

would be good relative to these ends. So just because no one has ever had to fight off a wave of (non-philosophical) zombies, it is still true that certain weapons would be good relative to this goal, and that others are bad. Some weapons are such that using them would prove successful against a horde of zombies—others are not. Weapons are good relative to fighting off zombies when those weapons are sufficient for getting the job done. They do not actually ever have to do the job for them to be good towards that job. Understood this way, we can respond to the third type of objection to achieving views. We can preserve the intuition that the virtuous samurai's sword is good, even if he never uses it in battle. It is good because if it were to be used in battle, it would get the job done.

Handling the other three objections requires noticing another feature of ends as I understand them. On my view, ends are the potential content of intentions—they are plans that we could decide to pursue. Now often, when talking about our ends, we do not bother stating all the details of the plan. We avoid doing so for two reasons. First, sometimes we simply do not have all the details in our mind. More often than not, our plans are rough sketches, and our plans are consequently somewhat flexible. It would thus be impossible to state all the details of the plan we have in mind. Second, and perhaps more importantly, it would be obnoxious and time consuming if everyone went around describing their goals by expressing all the excruciating details of how they plan on accomplishing those goals. No one has time for that! For these reasons, we often refer to our ends with short, pithy descriptions such as: I am going out to lunch.

My contention is that these pithy descriptions are just the tip of the iceberg when it comes to the actual content of our ends. Consider the goal of going out to lunch. When

I pursue this end, my plans tend to be fairly detailed. For instance, I generally want to eat lunch in a certain amount of time, I don't want to travel very far to get to a restaurant, I only want to spend a certain amount of money, I want to enjoy my food, I want to eat with certain people, I want to avoid certain conversation topics, I do not want to eat at the same place I ate yesterday, I want to get through the meal without dribbling food all over myself, I do not want to choke on my food and die... etc. Each of the items in this list deserves to be considered part of my goal of going out to lunch. I might not even be aware that my goal of going out to lunch contains all of these parts. However, it is worth noting that if I believed that one or more of these things would not occur, I might decide that I cannot achieve my goal, and I might cease intending to go out to lunch.

My claim is that each of these sub-goals is a part of the goal I am pursuing. Further, were I to have a different set of sub-goals in mind, I would be pursuing a different goal. Thus these sub-goals are necessary parts of the 'super-goal.' This might sound counterintuitive at first. On this view, for instance, if part of my plan for lunch is not eating at Chipotle because I ate there yesterday, and you do not have that as part of your plan for lunch, then we are not pursuing the same goal. We are, however, both going out to lunch. That sounds a bit odd.

Oddness aside, the goal of going out to lunch at any local restaurant *is* a different goal from going out to lunch at any local restaurant other than Chipotle. They have different satisfaction conditions. I will not accomplish my goal if we eat at Chipotle, but you will. However, we should not thus conclude that I failed to go out to lunch. Again, the label 'going out to lunch' can apply to a very wide variety of goals. If I am somehow coaxed into eating at Chipotle, it is true that I do not accomplish my original goal, but I

might simply switch to pursuing a very similar goal that can still be described as ‘going out to lunch’. I can still successfully go out to lunch, even if my original plan is foiled. Even though I still succeed at going out to lunch, I might nonetheless grumble as we sit down to lunch at Chipotle: “This isn’t exactly what I had in mind when I suggested we go out to lunch.”

Or consider two Olympic athletes about to compete in the 500-yard dash. They both have the goal of taking home the gold. One of them, however, is hit in the face by some garbage that a belligerent spectator flings onto the track. The athlete consequently stumbles and falls flat on her face. The other athlete goes on to win the gold. People might say to the winner, “Congratulations! You achieved your goal!” to which she might respond, “No, I didn’t,” or “Yes, but I did not intend it to happen like this.” If she were to say either of these things, we would come to learn that her goal was not simply to win a gold medal in the 500-yard dash. She intended to win it as the result of a fair race, not because her competitor was disadvantaged. Notice too, that if her goal is to win fairly, then the garbage hitting her competitor is not good relative to this goal. If, on the other hand, she wants to win no matter what, then the garbage hitting her competitor is good.

I hope that these examples suffice to show that the goals we adopt can be much more complicated than the pithy descriptions that we give them—they often include many sub-goals, the achievement of which are necessary for the achievement of the super-goal.²⁰³ If this is right, then we are now in a position to answer the remaining three objections to an achievement theory of goodness. First, it was noted that a butter knife can be used to cut down a redwood tree, but that intuitively, a butter knife is not good

²⁰³ Gilbert Harman makes the observation that our intentions are often more complicated than we describe them in (1976).

relative to that goal. If we note, however, that goals are more complicated than the labels we give them, this objection is easily handled.²⁰⁴ Cutting down a redwood tree with a butter knife requires a lot of time and effort. If my goal were simply to fell a redwood tree, then I could indeed achieve this goal with a butter knife. Odds are, however, that my goal of cutting down a redwood tree contains certain sub-goals about how much time and effort I plan to put into the endeavor. It is likely that a butter knife cannot cut down a redwood fast enough or with the amount of ease with which I aim to accomplish this task. A butter knife is not sufficient for achieving my goal. It is thus false that a butter knife is good for cutting down redwoods.

The fourth kind of objection noted that achievement of *G* is binary—it either occurs or it does not. However, goodness comes in degrees; thus goodness cannot merely be achievement of goals. My view can answer this objection in two ways. First, *X* and *Y* might both be good relative to *G*, but *X* might be better than *Y* if it achieves more of the sub-goals of *G*. If we have to choose between an *X* and a *Y*, and *X* does more to bring about *G* than *Y*, we can conclude that relative to *G*, *X* is better than *Y*. Better to ‘kill two birds with one stone’ than to have to use two stones to do so. The second way that we can respond to this objection is to note that sometimes it is the indeterminate nature of our goals that explains why *X* is better than *Y*. Suppose, for instance, that my goal is to cut down a tree quickly. A butter knife does not achieve this goal, but perhaps, one of those old-fashioned, two-person saws would. A chainsaw, however, would cut down the tree

²⁰⁴ Finlay builds these restrictions into the background (2014) p.60-1. Recall that the ‘background’ was a set of propositions *B*, that limited which possible world types were relevant when evaluating if *p* increases the probability of *e*. We only looked at world types that were consistent with *B*.

even quicker. The chainsaw is thus better, relative to cutting down the tree quickly, than the two-person saw.

This is not to say that more is always better. If I am looking to buy a new car which seats four, it does not follow that a large van that seats twelve is thus better. It depends on my goal. If my goal is to own a vehicle that can transport as many people as possible, the van is better than say, a Toyota Prius. If, however, my goal is just to be able to seat four people in the car, the fact that the Prius and the van both accomplish this goal shows that they are equally good (relative to my goal). My decision about which to buy would be made with respect to how the two vehicles compare in other respects. If my goal is to seat four, and no more than four, the van is worse than the Prius.

The second kind of objection noted that things that are not sufficient for achieving an end can still be good relative to that end. The example I gave above was that pleasant weather is good for a bicycle ride—but clearly it is not sufficient for a bicycle ride! To answer this objection, we need to fine-tune our achievement account one more time. Good weather is not sufficient for the achievement of a bike ride. However, good weather might be sufficient for a necessary sub-goal of my goal of going for a bike ride. Suppose, for instance, it is part of my goal that I do not want to ride in the rain. If the weather is sunny and pleasant, then my sub-goal of not riding in the rain will be accomplished. The sunny weather accomplishes this necessary sub-goal, and thus, is good relative to the super-goal of going for a bike ride. It is worth noting that I do not need to achieve the super-goal of going for a bike ride for it to be true that pleasant weather would be good for doing so. Remember, we are no longer talking about actual achievement of goals, but only of what is sufficient for the achievement of goals. Pleasant weather is sufficient for

achieving the goal of not riding in the rain—and thus good relative to that goal—whether I in fact climb on my bike and take advantage of the weather or not.

Finally, we need to add one more nuance to the view to block a certain kind of counterexample. Suppose I ate lunch at Chipotle yesterday, and I do not want to eat there again today. Part of my plan for going out to lunch is avoiding Chipotle. One way I could avoid Chipotle would be if I were suddenly and brutally murdered. My death is sufficient for achieving the sub-goal of not eating at Chipotle—but it is not, intuitively, good relative to the goal of going out to lunch today. The problem is that my death frustrates the achievement of many of the other sub-goals that are components of my super-goal of going out to lunch. I must add to my account of goodness, then, that X cannot be good relative to G if X frustrates a sub-goal of G. So while my death does ensure that I do not eat at Chipotle, it is nonetheless bad relative to my goal of going out to eat.

What should we say, then, about the following case: Again, I am going out to eat and I do not want to eat at Chipotle for the second day in a row. However, as luck should have it, a massive fire starts in the kitchen at Chipotle, burning the restaurant to the ground. Now of course the destruction of Chipotle is sufficient for achieving my sub-goal of not eating at Chipotle. Further, Chipotle's destruction does not frustrate any of the other sub-goals of my going to eat out. Should we conclude that the destruction of Chipotle is good?

Perhaps counter-intuitively, I believe the destruction of Chipotle is good. I say this only because I understand that claim as restricted relative to my goal of eating out at any place other than Chipotle. Its destruction helps achieve that goal. However, I nonetheless might not ever wish for the destruction of Chipotle, and I might even declare

that it is a bad thing that the building burned down. Chipotle's destruction will be bad relative to other goals—some that I might even care about more deeply than I do my getting lunch today at somewhere other than Chipotle. For instance, due to the fire, people were put in danger and likely even injured. Those are things that I would prefer did not happen! I might thus, all things considered, lament the destruction of Chipotle even though its destruction is good relative to a goal I am interested in pursuing.

In summary, the view I prefer states that *X is end-relationally good relative to G if and only if X is sufficient for the achievement a goal G or for a sub-goal of G, while at the same time not frustrating any of the other sub-goals of G*. This view does not fall prey to the sorts of objections that a naïve achievement view does. Further, it avoids both having to take a controversial stance on how to understand probability, and the problems that arise when for the view that the truth conditions of our goodness judgments be determined in relation to a comparison class. For instance, on my view, we could have a world where all steak knives are good (because they are all sufficient for achieving the goal (or a sub-goal) of cutting steak into bite-sized pieces). It is also possible on my view for there to be a world where all steak knives are bad (they fail at achieving this goal, or frustrate some further sub-goal). I conclude that my account of the nature of the goodness relation is more plausible than any of the alternatives.

V. End-Relational Moral Goodness?

Way back in the *Introduction* to this dissertation, I noted that for most philosophers, the interesting questions about the nature of goodness are in regards to *moral* goodness. In this last section, I wish to briefly discuss the plausibility of an ethics

that only posits end-relational goodness. Of course, fully developing an end-relational ethics would be its own book-length project. For the time being, I simply wish to gesture at what such a theory might look like. In this way, I wish to ward off the worry that an ethical theory that only posits end-relational goodness would lack some of the features that we expect of an ethical theory. I will consider three such features: (1) Universality, (2) Categoricality, and (3) The Priority of the Moral. If an end-relational account of goodness cannot preserve these features of a morality, then an end-relational account of moral goodness would be a failure.

Before looking at that, let me recall a few things that fall out of accepting an end-relational theory of goodness. First, as I argued in Chapter Two, there is no conceptual connection between goodness and people's desires, pro-attitudes, etc. Goals, as I understand them, are merely possible intentional plans. There are facts about what is good relative to these goals independent of whether or not anyone even has these goals. For this reason, my theory of goodness rules out non-cognitivism about goodness. However, my view also explains why non-cognitivism has seemed compelling to so many philosophers. According to my view, to judge something to be good is to judge that it promotes a goal. Whenever someone makes a particular goodness judgment, then, she does so with some goal in mind. It is very rare for people to think (and speak) about goals that they do not have pro-attitudes about, so generally speaking, it is rare for people to make goodness judgments on topics they do not care about. It is easy to confuse this near constant conjunction of pro-attitude and goodness judgment for a necessary connection.

Second, as I discussed in Chapter Three, since there is no such thing as intrinsic goodness, a very common way of characterizing one's moral theory is precluded. Briefly,

the schema for this characterization is: (a) Identify some thing or things that are intrinsically good, (b) declare that some set of responses to that thing or things that is morally required. If my end-relational theory of goodness is correct, then this is not an adequate way to characterize one's moral theory. Since nothing is intrinsically good, we can never satisfy the first condition (a), and thus, no moral theories of this type can be true. Given that there are many deontological, consequentialist, and virtue theories that follow this schema, the end-relational theory of goodness seems to entail that a vast variety of moral theories are false. That is a pretty big bullet to bite.

An end-relational theorist, however, can defend moral theories that are nearly identical to the ones mentioned above. She simply needs to construct the correlating theories according to a different schema—one that does not posit intrinsic goodness. The idea was to interpret talk of 'intrinsic goodness' as actually being about goodness relative to some special set of goals. If a compelling argument can be given, for instance, that pursuing happiness is a special kind of goal, then we can reconstruct a consequentialist ethical theory without appealing to the intrinsic goodness of pursuing happiness. For instance, one might claim that the pursuit of happiness is the only morally relevant goal. If so, then only things that promote happiness are morally good. This sounds, in principle, like an acceptable way to get a consequentialist theory going. It thus seems that recognizing that goodness is relational is not quite the silver bullet to consequentialism that Foot and Thomson thought it was.²⁰⁵

But could such an end-relational ethics preserve the fact that moral reasons are categorical? That depends on how we understand categoricity. According to one sense of 'categorical,' a reason is categorical if it holds of someone, regardless of what ends they

²⁰⁵ Foot (1985), Thomson (1994)

have. This notion of categoricity is contrasted with hypothetical reasons—reasons that hold of you only if you have certain ends. The fact that a certain kind of lumber is good for furniture building does not give you reasons unless you want to build furniture. The reasons for using that lumber are thus hypothetical. Moral reasons are said to be categorical. It does not matter what goals you have, if there is a moral reason to X, then that reason holds of you.

End-relational moral goodness does not ground categorical reasons in this sense. According to an end-relational account of moral goodness, there are certain goals that are the moral goals, and for X to be morally good is for X to promote one of those moral goals. The fact that X is morally good only gives you reason if you have the moral goal. The reason it gives you is hypothetical, since it would not hold of you *regardless* of what ends you have—you need to hold the moral goal for X’s moral goodness to give you reason. Even if the end-relationalist argues that we all *in fact* have the relevant moral goals, the reasons that X’s moral goodness generates would still be hypothetical in the sense we are currently discussing.

Of course, it is open to the end-relationalist to challenge this notion of categoricity; there are, after all, other ways to characterize categoricity. For instance, on one conception of categoricity, a reason is categorical if it is, in some sense, inescapable.²⁰⁶ The thought is that when it comes to hypothetical reasons, like that a certain kind of lumber is good for building furniture, I can ‘escape’ this reason by not having the goal of building furniture. There is no similar way to escape moral reasons, so they are categorical.

²⁰⁶ Finlay considers this interpretation of ‘categoricity’ in (2014) p.177.

Understood in this way, moral reasons grounded in end-relational goodness can be categorical. Here is a sketch about how reasons grounded in end-relational moral goodness can be inescapable. Start with the goal of being able to pursue one's goals. This goal is one that we all care about, and one that we cannot help but care about. Perhaps it is grounded in human nature, or perhaps it is grounded in the nature of agency, but one thing that we all aim at is to be able to pursue goals. This goal is thus, 'inescapable'. Now suppose, as I think it is quite right to, it is a necessary condition for our being able to pursue our goals that we have others' cooperation. Notice that even if my goal does not directly require that anyone help me—for instance, say I want to go for a solitary walk through the woods—it does require *indirect* cooperation; it requires that no one interferes with my goal. I cannot take a solitary walk in the woods if someone decides to always walk two steps behind me! We now have the claim that cooperation from others is necessary for my achieving my inescapable goal of pursuing my goals. Finally, we can claim that one role that acting in accordance with moral rules plays is that it makes possible our shared cooperation with one another. Our cooperation is dependent on being able to trust one another. Without some set of moral rules stipulating how we must act towards one another, and unless we can expect that others will act in accordance with those rules, we cannot trust in one another's cooperation. If that is right, then the moral rules are necessary for cooperation,²⁰⁷ and cooperation is necessary for our pursuing our goals, and it is inescapable that we aim at pursuing our goals. We thus have inescapable

²⁰⁷ Many moral philosophers have discussed the relation of cooperation and morality. See, for instance: Hobbes (1651), Locke (1689), Rousseau (1762), Kant (1793), Rawls (1971), Gauthier (1986), and Copp (1997) (to name a few).

reasons to act in accordance with the rules of morality. Again, this is a merely a sketch, and would take some time to develop in more detail.²⁰⁸

If something like this sketch is correct, we can also explain why moral reasons are universal—that is, why they hold of everyone—and why we give priority to moral considerations. My claim was that *everyone* (or at least, all moral agents) aim at pursuing their goals. If doing so requires the cooperation of others, and the cooperation of others requires that we obey the rules of morality, then we all have reason to obey the rules of morality. Further, since the ability to pursue one’s goals is necessary no matter what goal it is that one opts to pursue, the ability to pursue one’s goal is good relative to all the other goals we want to pursue. If our behaving morally is a necessary condition for the preservation of our ability to pursue the goals we opt to pursue, then our behaving morally is also necessarily good relative to all the other goals we wish to pursue. Behaving morally is a necessary prerequisite for our ability to pursue any other kind of goal that we opt to. For that reason, we give moral considerations priority over others.

The point of this section was not to convince you that this is the correct view of morality. That would require much more space than this! The point of this section was to convince you that there could be a plausible account of morality that only posits end-relational goodness and yet preserves the categoricity, universality, and priority of morality. Developing that view in any detail must be left for future work.

²⁰⁸ As it stands, it seems susceptible to free-rider problems, or to the problem of people merely appearing to be moral, when they really aren’t. I could ensure the cooperation of others by merely pretending to be moral. It is thus, not necessary that I behave morally.

VI. Conclusion

The arguments of this chapter support the claim that my theory of end-relational goodness is more plausible than the one offered by Stephen Finlay. He and I disagree over (i) what things can stand in the goodness relation to ends, (ii) what ends are, and (iii) what the nature of the goodness relation is like. I claimed that in each of these disagreements, my view was more plausible than his. The conclusion of this chapter, then, is that *if* we are going to defend an end-relational theory of goodness, it should be one where a thing (broadly construed) is good if and only if it is sufficient for the achievement of a goal, G (understood as abstract plan), or for the achievement of one of G's sub-goals, while at the same time, not frustrating any other sub-goals of G.

The arguments of the previous chapters have supported the claim that we should, in fact, defend an end-relational theory of goodness. We should do so because an end-relational theory of goodness is the most plausible account of non-moral goodness. Kind-relational and Interest-relational views of goodness are, for a variety of reasons, just not as plausible. If this is correct, then our use of 'good' in our non-moral judgments ascribes to things an end-relational property. Further, the word 'good' would ascribe the same attribute in moral judgments as it does in non-moral judgments, unless it is ambiguous. It is not ambiguous, so it does not ascribe different attributes in moral judgments than it does in non-moral ones. It ascribes end-relational goodness in non-moral judgments, thus, it ascribes end-relational goodness in moral judgments as well.

I conceded that if adopting an end-relational theory of goodness conflicted with too many of our intuitions about what we should expect from a theory of goodness, then all the worse for an end-relational theory of goodness! In this case, we would have reason

to abandon either our end-relational account of non-moral goodness, or the claim that ‘good’ ascribes the same attribute in moral judgments as it does in non-moral judgments. I argued, however, that an end-relational theory of goodness does not conflict with too many of our commonsense intuitions. For instance, even though on the end-relational theory that I favor there is no such thing as intrinsic goodness, we do not have to give up on all the roles that intrinsic goodness plays in our theoretical and practical lives. We can explain, for instance, why we have basic reasons to pursue some goals rather than others, without positing intrinsic goodness.

The theory I favor is simple, elegant, and has a great deal of explanatory power. Further, it opens some new and interesting ways of thinking about ethics, and how moral goodness relates to the goodness of parking spaces, umbrellas, etc. In short, my theory is a very good theory of goodness.

Bibliography

- Anscombe, G.E.M. (1957). *Intention*. Cambridge, MA: Harvard University Press.
- Aristotle. *Nicomachean Ethics*. In *The Basic Works of Aristotle*, ed. Richard McKeon (1941). New York: Random House.
- Audi, Robert (1997). *Moral Knowledge and Ethical Character*. Oxford: Oxford University Press.
- Ayer, A.J. (1936). *Language, Truth and Logic*. London: V. Gollancz, Ltd.
- Beardsley, Monroe, C. (1965/2005). "Intrinsic Value" Reprinted in *Recent Work on Intrinsic Value*. Eds. T. Rønnow-Rasmussen and Michael J. Zimmerman. Netherlands: Springer: 61-75
- Bentham, Jeremy (1907). *An Introduction to the Principles of Morals and Legislation*. Oxford: Clarendon Press.
- Blackburn, Simon (1993). *Essays on Quasi-Realism*. New York: Oxford University Press.
- Bratman, Michael (1987). *Intention, Plans, and Practical Reason*. Cambridge, MA: Harvard University Press.
- Brink, David O. (1986). "Externalist Moral Realism." *Southern Journal of Philosophy*, 24 (Supplement): 23–40.
- Butchavarov, Panayot (1989). *Skepticism in Ethics*. Indianapolis: Indiana University Press.
- Chang, Ruth (2002). "The Possibility of Parity." *Ethics* 122: 659-88.
- Chisholm, Roderick M. (1986). *Brentano and Intrinsic Value*. Cambridge: Cambridge University Press.
- Copp, David (1995). *Morality, Normativity, & Society*. New York: Oxford University Press.
- Copp, David (2001). "Realist-Expressivism: A Neglected Option for Moral Realism." *Social Philosophy and Policy* 18(2): 1-43

- Copp, David (2011). "Book Review: Thomson, Judith Jarvis *Normativity*." *Australasian Journal of Philosophy* 88: 1-3
- Enoch, David (2011). *Taking Morality Seriously*. Oxford: Oxford University Press.
- Ewing, A.C. (1947). *The Definition of Good*. London: Routledge.
- Finlay, Stephen (2001). "What Does Value Matter? The Interest-Relational Theory of the Semantics and Metaphysics of Value." Ph.D. dissertation, University of Illinois, Urbana-Champaign.
- Finlay, Stephen (2004). "The Conversational Practicality of Value Judgment." *Journal of Ethics* 8(3): 205-223.
- Finlay, Stephen (2005). "Value and Implicature" *Philosophers' Imprint* 5(4): 1-20.
- Finlay, Stephen (2006). "The Reasons that Matter." *Australasian Journal of Philosophy* 84(1): 1-20.
- Finlay, Stephen (2008). "The Error in Error Theory." *Australasian Journal of Philosophy* 86(3): 347-69.
- Finlay, Stephen (2014). *Confusion of Tongues: A Theory of Normative Language*. New York: Oxford University Press.
- Foot, Philippa (1967). "The Problem of Abortion and the Doctrine of Double Effect." *Oxford Review* 5: 5-15
- Foot, Philippa (1985). "Utilitarianism and the Virtues." *Mind* 94: 196-209.
- Foot, Philippa (2001). *Natural Goodness*. Oxford: Clarendon Press.
- Frankena, William K. (1973). *Ethics*. Englewood Cliffs: Prentice Hall.
- Frankfurt, Harry (1982/2007). "The importance of what we care about." Reprinted in *The Importance of What We Care About*. New York: Cambridge University Press: 80-94.
- Gauthier, David (1987). *Morals by Agreement*. Oxford: Oxford University Press.
- Geach, Peter (1956). "Good and Evil." *Analysis* 17: 33-42.
- Grice, H.P. (1975/1989). "Logic and Conversation." Reprinted in *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.

- Hare, R.M. (1952). *The Language of Morals*. Oxford: The Clarendon Press.
- Hare, R.M. (1957). "Geach: Good and Evil." *Analysis* 18: 103-12
- Harman, Gilbert (1976). "Practical Reasoning." *The Review of Metaphysics* 29(3): 431-63.
- Hobbes, Thomas (1651/1994). *Leviathan*. Indianapolis, IN: Hackett.
- Hurka, Thomas (1987). "'Good' and 'Good for'." *Mind* 96(381): 71-73.
- Hurka, Thomas (1993). *Perfectionism*. New York: Oxford University Press.
- Hurka, Thomas (2010). *The Best Things in Life: A Guide to What Really Matters*. New York: Oxford University Press.
- Jackson, Frank (2000). *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Oxford University Press.
- Kagan, Shelly (1998). "Rethinking Intrinsic Value." *Journal of Ethics*, 2(4), Intrinsic Value: 277-297.
- Kant, Immanuel (1785/1999). *Groundwork for the Metaphysics of Morals*. In *The Cambridge Edition of the Works of Immanuel Kant: Practical Philosophy*. Ed. Mary J. Gregor. Cambridge University Press: 37-108
- Kant, Immanuel (1793/1999) "Theory and Practice" In *The Cambridge Edition of the Works of Immanuel Kant: Practical Philosophy*. Ed. Mary J. Gregor. Cambridge University Press, (1996), pp. 273-310.
- Korsgaard, Christine M. (1983). "Two Distinctions in Goodness." *Philosophical Review* 92(2): 169-195.
- Korsgaard, Christine M. (2013). "The Relational Nature of the Good." *Oxford Studies in Metaethics* 8(1).
- Kraut, Richard (2007). *What is Good and Why: the Ethics of Well-Being*. Cambridge, MA: Harvard University Press.
- Kraut, Richard (2011). *Against Absolute Goodness*. New York: Oxford University Press.
- Mackie, J.L (1990). *Ethics: Inventing Right and Wrong*. London: Penguin Books.
- Mill, John Stuart (1906/1991). *Utilitarianism*. Reprinted in *On Liberty and Other Essays*. Ed. John Gray. New York: Oxford University Press.

- Moore, G. E. (1903/1993). *Principia Ethica*. Edited by Thomas Baldwin. Rev. ed. Cambridge, England: Cambridge University Press.
- Parfit, Derek (2011). *On What Matters*. Volume 1 & 2. New York: Oxford University Press.
- Perry, R.B. (1926). *General Theory of Value*. New York: Longmans, Green & Co.
- Rabinowicz, Wlodek & Rønnow-Rasmussen, Toni (2000). "A Distinction in Value: Intrinsic and for its Own Sake." *Proceedings of the Aristotelian Society* 100(1): 33-51.
- Rabinowicz, Wlodek & Rønnow-Rasmussen, Toni (2003). "Tropic of Value" *Philosophy and Phenomenological Research* 66(2): 389-403.
- Railton, Peter (1986). "Moral Realism." *The Philosophical Review* 95(2): 163-207.
- Rawls, John (1971/1999). *A Theory of Justice: Revised Edition*. Cambridge, MA: Belknap Press.
- Ross, W. D. (1930/2009) *The Right and the Good*. Edited by Philip Stratton-Lake. Oxford: The Clarendon Press.
- Scanlon, T.M. (1998). *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.
- Schroeder, Mark (2010). *Noncognitivism in Ethics*. London: Routledge.
- Sennet, Adam (2011). "Ambiguity", *The Stanford Encyclopedia of Philosophy* Ed. Edward N. Zalta. URL=
<<http://plato.stanford.edu/archives/sum2011/entries/ambiguity/>>.
- Shanklin, Robert (2011). "On Good and 'Good'." Ph.D. dissertation, University of Southern California.
- Suits, Bernard (1978). *The Grasshopper: Games, Life and Utopia*. Toronto: University of Toronto Press.
- Sumner, L.W. (1996). *Welfare, Happiness, & Ethics*. Oxford: Oxford University Press.
- Thomson, Judith Jarvis (1992). "On Some Ways in Which a Thing Can Be Good." *The Good Life and the Human Good*. Ed. Fred Miller, E.F. Paul, and Jeffery Paul. New York: Cambridge University Press.
- Thomson, Judith Jarvis (1993). "Goodness and Utilitarianism." *Proceedings and Addresses of the American Philosophical Association* 67(2): 145-59.

- Thomson, Judith Jarvis (1997). "The Right and the Good." *The Journal of Philosophy* 94(6): 273-98.
- Thomson, Judith Jarvis (2003). "The Legacy of Principia." *Southern Journal of Philosophy*, Spindell Conference, Supplement
- Thomson, Judith Jarvis (2008). *Normativity*. Chicago: Open Court.
- Thomson, Judith Jarvis, Philip Fisher, and Amy Gutmann (2001). *Goodness and Advice*, The University Center for Human Values Series. Princeton, NJ: Princeton University Press.
- Velleman, David (2007). *Practical Reflection*. The David Hume Series: Philosophy and Cognitive Science Reissues. Stanford, CA: CSLI Publications.
- Wittgenstein, Ludwig (1953/2001). *Philosophical Investigations*. Blackwell Publishing
- von Wright, Georg Henrik (1963). *The Varieties of Goodness*. London: Routledge & Kegan Paul.
- Ziff, Paul (1960). *Semantic Analysis*. Ithaca, NY: Cornell University Press.
- Zimmerman, Michael J. (2001). *The Nature of Intrinsic Value*. Lanham, MD: Rowman & Littlefield.
- Zimmerman, Michael J. (2004). "Judith Jarvis Thomson, Goodness and Advice." *Noûs* 38(3): 534-52.
- Zimmerman, Michael J. (2007). "The Good and the Right." *Utilitas* 19(3): 326-53.
- Zimmerman, Michael J. (2010). "Intrinsic vs. Extrinsic Value." *The Stanford Encyclopedia of Philosophy*. Ed. Edward N. Zalta. URL = <http://plato.stanford.edu/archives/win2010/entries/value-intrinsic-extrinsic/>.