

THE BEST MORAL THEORY EVER:
THE MERITS AND METHODOLOGY OF MORAL THEORIZING

by

Jason Brennan

Copyright © Jason Brennan 2007

A Dissertation Submitted to the Faculty of the

DEPARTMENT OF PHILOSOPHY

In Partial Fulfillment of the Requirements
For the Degree of

DOCTOR OF PHILOSOPHY

In the Graduate College

THE UNIVERSITY OF ARIZONA

2007

THE UNIVERSITY OF ARIZONA
GRADUATE COLLEGE

As members of the Dissertation Committee, we certify that we have read the dissertation prepared by Jason Brennan entitled *The Best Moral Theory Ever: The Merits and Methods of Moral Theory* and recommend that it be accepted as fulfilling the dissertation requirement for the Degree of Doctor of Philosophy

Date: 1/04/07
David Schmitz

Date: 1/04/07
Mark Timmons

Date: 1/04/07
Michael Gill

Final approval and acceptance of this dissertation is contingent upon the candidate's submission of the final copies of the dissertation to the Graduate College.

I hereby certify that I have read this dissertation prepared under my direction and recommend that it be accepted as fulfilling the dissertation requirement.

Date: 1/04/07
Dissertation Director: David Schmitz

STATEMENT BY AUTHOR

This dissertation has been submitted in partial fulfillment of requirements for an advanced degree at the University of Arizona and is deposited in the University Library to be made available under rules of the Library.

Brief quotations from this dissertation are allowable without special permission, provided that accurate acknowledgement of source is made. Requests for permission for extended quotation or reproduction of this manuscript in whole or in part may be granted by the copyright holder.

SIGNED: Jason Brennan

ACKNOWLEDGEMENTS

Staring at a blank page with the title "Acknowledgements" confronts me with the debt I owe to others. I hope that this work begins the process of doing justice to that debt.

Let me begin by thanking Chris Maloney, who assigned me to be a teaching assistant in a philosophy of science course during my third semester at Arizona, despite its being my last choice. Had I never been assigned that class, I never would have chosen this project.

Matt Zwolinski served as an early model for me of what a graduate student should be. He later served as a groomsman in my wedding. I am thankful for his example and for his continued friendship. I am also grateful for friendship and philosophical discussions with Chris Freiman. I wish that I had been more like Chris at his stage.

Numerous fellowships and grants made my love affair with philosophy possible. The Earhart Foundation provided me with funding for two years. The Institute for Humane Studies took a deep interest in my development, providing me with travel grants and fellowships over three years. They also connected me with brilliant people, exposed me to new ideas from the social sciences, as well as taught me some of the non-academic skills academics need to be academics. Nigel Ashford, Amanda Brand, and Nikki Sullivan especially deserve my thanks. My graduate career was deeply enriched by IHS.

Thanks to Mark Timmons and Michael Gill. Their challenges have made this a better work. They have helped to keep this work grounded and rigorous. They are fabulous teachers and philosophers.

David Schmitz has been more than a mentor for me. Philosophical writing is an art. Dave is a consummate artist. More importantly, he leads his life in a way that gives meaning to the concept of the *phronimos*. Reaching the stage at which I can call Dave a colleague in addition to a teacher will be one of my greatest honors.

Thanks most of all to Lauren, for her love, for starting this whole adventure with me and for being with me all the way, and most importantly, just for being who she is.

DEDICATION

To Lauren, the Best Ever

TABLE OF CONTENTS

LIST OF TABLES.....	7
ABSTRACT.....	8
1. THE CONTOURS OF PHILOSOPHICAL ETHICS.....	10
2. MORAL THEORY AND ITS DISCONTENTS.....	31
3. THEORY MULTIPLISM.....	61
4. MORAL THEORY AND THE ILLUMINATION OF TRUTH.....	95
5. THE MORAL COSMOS: ON BEHALF OF GENERALISM IN ETHICS.....	120
6. PHILOSOPHICAL DISSENSUS.....	150
7. TRUSTING OUR MORAL INTUITIONS.....	183
8. THE PRACTICAL VALUE OF MORAL THEORY.....	212
REFERENCES.....	239

LIST OF TABLES

TABLE 4.1, Varieties of Realism.....	102
--------------------------------------	-----

ABSTRACT

Anti-theorists claim that moral theories do not deliver all the goods we want and that consequently such theorizing is not a philosophically worthy pursuit. We suffer from certain misconceptions about the point and purpose of such theorizing and the theories it produces. In this essay, I treat moral theorizing as a genuinely theoretical enterprise that produces abstract knowledge about the general structure of morality.

Moral theories should be understood as tools—intellectual and practical tools with importantly different uses. Just as with hand tools where it is useful to have hammers for one sort of job and screwdrivers for another, it can be rational to accept multiple moral theories at the same time. The idea here is that all good theories illuminate some truths about morality, but are also misleading at times. A theory that is good at solving one moral problem may be bad at solving another; a theory that is illuminating in one place may be distorting in another.

Chapter one outlines the differences between moral theory, metaethics, moral metatheory, and morality itself. It argues that disagreement about moral theory need not reflect moral disagreement, and vice versa. Chapter two argues that even if moral theory turned out to be practically useless, it would still accomplish certain theoretical tasks. Chapters three and four explain how and why one might adopt different incompatible moral theories at the same time. Chapter five defends moral principles from various particularists and shows how the imperfections of moral principles mirror the imperfections of laws in other fields. Chapter six explains why philosophical inquiry is worthwhile despite the overwhelming disagreement displayed by philosophers. Chapter

seven shows that moral intuitions serve as a check on philosophical methodology just as much as methodology helps us verify our intuitions. It explains why a certain sort of psychology-based argument against deontological intuitions will not work. Finally, chapter eight explores the various ways in which moral theory is and is not practical. It concludes that the practical usefulness of theory is a matter of empirical contingency that philosophers have done little to investigate.

CHAPTER ONE

THE CONTOURS OF PHILOSOPHICAL ETHICS

1. MORALITY AND MORAL THEORY

“Hello. Welcome to the Pearly Gates,” said the grizzled old man.

Eric was dazzled. “The Pearly Gates? Does this mean I get into heaven?”

He frowned. “No, sorry about that. According to my records, you’ll be spending most of your time down...south, if you know what I mean.”

“Oh... But I tried to live a good life. I—”

“Yes, yes. You’re quite right. In fact, the Boss was rather fond of you. We could always count on you to do the right thing for the right reason and know what you were doing. You always had the right attitude.¹ We see very few like you, even here. But you see, there’s just one problem.”

“What was that?”

“You never articulated and subscribed to the right moral theory.”

“Huh? What’s a moral theory?”

“Oh, you’ll find out. We force the inmates to read them. Keeps them from having fun.”

This story illustrates something we already know. You do not need to have a philosophical moral theory in order to be a good person. Eric should get into heaven

¹ Cf. Hursthouse, *On Virtue Ethics* (Oxford: Oxford University Press, 1999), pp. 123-125.

because he was good, regardless of his knowledge of moral theory. In an alternate universe, one in which no moral *theory* had ever been produced, many people could still deserve to get into heaven. The fate of morality does not hinge on moral theory. Still, we also know that being a good person means living a moral life, and living a moral life entails understanding morality in some way. What then is a moral theory, if it is not necessary in order to be good?

Philosophers sometimes talk as if the point of moral theories was to find out what we need to do to be good. Perhaps they are right to talk this way. Yet it would be rather immodest—a philosopher's fantasy—to suggest that people would not have led good lives without philosophical ethics. There were virtuous people before the invention of moral philosophy, and there are virtuous people now who have never been exposed to Aristotle or Kant. Moreover, there are philosophers living today who have produced excellent, insightful moral theories but are not good people. What then is the connection between a good moral theory and being good?

Morality is not the same thing as moral theory, though a philosopher might carelessly speak as if they are the same. One might say, "I must do X because it is demanded by the theory of justice", but of course, that is not true. The reason one must do X is because it is demanded by justice *itself*, not a *theory* of justice. I fall to Earth because of gravity, not because of any *theory* of gravity.

In general, theorizing—the practice of producing theory—is an attempt to understand a subject matter. With different types of theories, people wish to accomplish different things. Even with the same theory, people wish to accomplish different things.

However, what unites all forms of theorizing is that they aim at understanding. Thus, if we ask what makes theorizing successful, a simple, first shot answer would be that it achieves its aims. Successful theorizing helps us understand. By extension, we could say that successful *theory* helps us understand. However, we should remember that *theorizing* and *theory* are two different things. It may turn out that theorizing helps us immensely to understand, even when individual theories do not. Theorizing is a process of reflection and articulation. Sometimes such reflection produces theories. It might turn out the process of reflection is far more fruitful than any theory (or all theories combined) could be.

We know that it is possible to be a good person without having a moral theory. We do not need a moral theory to tell us that this is true, and we can agree to this point regardless of other disagreements about moral theory or morality. We also know that a good person in some way understands what she is doing. If this is the case, then the understanding one obtains from a good theory is not the same understanding one has from being a good person. How can this be?

Luckily for moral theory, this split in types of understanding is ubiquitous. Many people consider Jimi Hendrix to have been an excellent blues-rock guitarist. However, Hendrix did not know music theory, how to read music, or have a theory of how to play guitar. Rather, Hendrix practiced to records, learned riffs from friends (who also lacked formal instruction in theory), and invented his style in order to pursue his artistic vision. In his own way, Hendrix understood guitar. He could not tell you why Mixolydian mode works for soloing over dominant seventh chords. However, he often used Mixolydian

mode when soloing over dominant seventh chords, so he was responsive to the music in the right way. It would be odd to say that he had a theory of guitar, but that this theory was internalized and implicit. He would have protested to the contrary. It is better just to say that he had a skill.

Otherwise, we have to posit theories everywhere we find knowledge. Children know how to play, but we do not want to say that they have implicit theories of how to play. But there are such things as theories of play. Many people know how to ride a bike, but they do not have theories of bike riding. The difference between theory and skill is the type of knowledge involved.

Hendrix did not know guitar theory, the theory of how to play guitar. Yet, there are books and courses on it. There are even university degrees in guitar performance, and a large percentage of these degrees are earned by learning guitar theory. It is possible for one to learn much about the guitar through these theories without actually becoming able to play. People can instruct others without being able or willing to play themselves, though that would be unusual. A good guitar teacher (or football coach) can continue to aid a student in performance even when the student's ability to perform far exceeds the teacher's, or even when the teacher never could perform at all.

The person with theory understands something that Hendrix does not, but Hendrix understands something that the person who only has theory does not. The person who has mastered both theory and practice understands even more. If a guitarist who knows theory talks to an equally skilled guitarist who does not, the two can discover that the first guitarist understands something the second does not. Musicians with theory can better

articulate what they are doing and can communicate with each other far easier. (A quick way to impress upon a guitarist the value of theory is to have him try to communicate musical ideas to a keyboard player.) They can understand the workings of instruments they do not play, and more quickly adapt to novel forms of music. (Towards the end of his life, Jimi Hendrix felt limited by his lack of music theory.) Knowing music theory also means having literacy (the ability to read and write music), something most guitarists lack. A guitarist with theory understands not merely how to perform, but the deep, underlying structure of music. With theory, one can articulate *why* certain musical ideas make sense. Music theory is not just a mode of discourse. It is a genuine theory. Moral theory is similar.

If a virtuous person with a worthwhile moral theory talks to a virtuous person without one, the two can discover a gap in their understanding. Suppose the Categorical Imperative shows the deep, underlying structure of ethics. The virtuous person need not be aware of this, though he can give reasons for what he does. He has something to learn from theory.

The typical person has many moral beliefs. Many of these beliefs are evaluations of particular objects, acts, and states of affairs. This person may also employ various principles in her deliberations. When asked, she can probably articulate a number of principles to which she subscribes: slavery is wrong, keeping promises is right, kindness is good, etc. She was taught various rules as a child and probably teaches these rules to her own children. She has a certain skill in making moral evaluations and acting morally. Much of the time, though, the average person acts on sentiment and automatic

dispositions, rather than on conscious principles. The difference between the virtuous person and her is mostly a degree of skill. The virtuous person excels at evaluating and acting morally. Not only does acting morally come easily to the virtuous person, she *enjoys* it. Strangely, she feels both *compelled* and free in acting morally.

I have claimed that such people do not have moral theories. What then must they acquire to have a moral theory? We can ask a similar question about any other field. My parents have beliefs about the way the economy works, but do they have economic theories, like I do? A five year old has beliefs about the physical world, but do these beliefs constitute a physical theory? To borrow an example from Ronald Giere: pre-historic peoples knew that objects fall to the earth when dropped. But it seems strange to claim they had physical theories and laws of physics. They lacked the concepts of "theory" and "physical law".² Similarly, the average person cares about morality, but she does not even know what a "moral theory" is. I do not wish to attribute acceptance of theories or laws to anyone who does not know what such things are, even if we might say that they act as if they believe in such theories or laws.

Concern for morality—and having moral beliefs—is not the same as having a moral theory. Consider the typical ethics professor who in fact does subscribe to a moral theory. Quite often, one can produce something that the professor will agree is a counterexample to her theory. Her knowledge of morality goes beyond her explicit theory. Otherwise, she would never agree that her theory has counterexamples, i.e., that

² Ronald M. Giere, *Science without Laws* (Chicago: University of Chicago Press, 1999), p. 2.

her theory misrepresents morality. One might contend that the theorist has two theories—an explicit and an implicit one—but that is unpalatable.

No clear boundary demarcates belief from theory. But we can elaborate some differences. Robert Louden suggests that one begins to have a moral theory when one offers reasons for one's views and fits these within a greater explanation of "how human beings should live".³ Theories are abstractions, "taking a larger view of things...not readily available to common sense."⁴ Theories are linguistic representations: a person can act on certain tacit beliefs, but these beliefs do not represent a "proposition", but at best just a set of functional behaviors that imply some sort of acceptance of a proposition. I would add that theories are explanations, not merely assertions. A moral theory does not merely try to elaborate the content of morality, but explains *why* the content is what it is. An average American may appear to be a divine command theorist because she will tell you that morality consists of God's commands, and God has commanded the Decalogue and the Golden Rule. However, that person begins to move from belief to moral theory when she can articulate why her God choose those rules and not others.⁵ I would add that the most fundamental moral theories offer in addition an account of moral agency. They try to explain how normativity and morality is possible for us, how there can be norms with moral force. Moral theorists might also try to explain what the

³ Robert B. Louden, *Morality and Moral Theory* (Oxford: Oxford University Press, 1992), pp. 140-141.

⁴ *Ibid.*

⁵ Thanks to David Schmidtz for suggesting this example.

purpose of morality is, what its essence is, what its cause is, and an elaboration of the nature of moral properties and moral reasoning. Perhaps people have some beliefs about these issues, but typically nothing quite like a theoretical framework.

A moral theory is not merely a set of principles. Suppose the average American is right about morality. An anthropologist studies this American and produces a list of the principles the American endorses, though the anthropologist has no opinion of the principles. This anthropologist does not seem to be a moral theorist in virtue of making his list. If we add that he endorses the principles, he still does not seem to be a moral theorist. Once he starts providing systematic justification for the principles, he comes closer. When his justification reaches a level of abstraction that responds to multiple human perspectives, not merely parochial concerns, he begins to cross the threshold from anthropologist to philosopher. Lastly, when he explains the nature of norms and how normativity is possible, he is clearly a moral theorist.

We might carve up the distinctions differently. Someone may wish to say that there are levels of theory. A person with particular judgments has a low-level theory. A person with mid-level principles (like the Decalogue) and some parochial justification for them has a mid-level theory. Immanuel Kant has a grand, high-level theory. If this sort of description proves more useful for some purposes, I have no qualms. Only a verbal difference exists between this view and the view I have advanced. However, if one chooses to speak this way, one should be cautious. If you had told Jimi Hendrix that he really did have a low-level guitar theory, I bet he would have been amused. The point

remains that he was missing something. Similarly, it would be amusing to tell Forrest Gump that he really does have a moral theory.

To turn the tables on myself, the philosophers we label "moral theorists" often merely seek to articulate lower moral principles. They sometimes do not provide deep justifications or try to explain how normativity is possible. It would also be amusing to tell them that they are not doing moral theory at that time.

2. MORAL VERSUS THEORETICAL DISAGREEMENT

If moral theory and morality are different things, then disagreements about moral theory are not the same thing as disagreements about morality. Moral theories are meant to explain correct moral judgments. (This is one of many success conditions of a moral theory.) We might agree on the moral judgments, but disagree on the best explanation of those judgments. If we agree on all the moral judgments, we agree on morality, regardless of what theory we subscribe to. The question of what theory best accounts for those judgments is a theoretical, not a moral, question.

Some may think this a non-starter. They might argue that if two theories concur with the same set of judgments, they are not distinct theories. The content of a moral theory just is the judgments implied by the theory. This objection reeks of logical positivism. Logical positivists held that the content of a physical theory just was the set of observations it implied. Accordingly, if two theories had the same observational implications, any differences between the two theories were merely verbal. I will not rehearse the reasons for rejecting this view here, since almost every philosopher rejects

this position on the nature of physical theories. Similarly, we should reject any sort of moral logical positivism. The content of a moral theory is not merely its implications for particular judgments. More importantly, neither in science nor in moral philosophy do theories logically entail (or are entailed by) the full set of observations or judgments they are meant to account for. Scientific theories and moral theories are not like geometry. They are not systems of axioms from which to perform deductions. The relationship between a moral theory and the judgments it explains is not direct entailment. Thus, we have room to accept that the theoretical content of moral theory is something besides a set of judgments.

Philosophers sometimes fail to notice the distinction between moral disagreement and theoretical disagreement. This is not surprising, since philosophers sometimes fail to notice the distinction between morality and moral theory. Two moral philosophers could, in principle, agree on every moral judgment they make in their entire lives, and yet produce two quite distinct moral theories. It would be surprising if, given that much moral agreement, the two did not have significant theoretical agreement, but it would also be surprising if they had full theoretical agreement. In the philosophy of science, the view that careful observation alone simply pushes a scientific theory onto a scientist is now derogatively called naïve Baconianism. (This view is unfairly named after Francis Bacon.⁶) Naïve Baconianism is the view that nature alone will tell us the best theory, just so long as we pay attention. The truth is that theory construction, in science, involves

⁶ Ian Hacking, *Representing and Intervening* (Cambridge: Cambridge University Press, 1983), pp. 168-169, 246.

choice and skill. Two scientists can agree on any set of observations, but the observations alone do not tell the scientists what theory to produce. Anyone familiar with the history of science knows that people will produce significantly different theories, even with the same data, assuming the data prompts them to produce theory at all. (Part of the problem is that recognizing data as data, and not merely background noise, takes skill, serendipity, and, sometimes, theory.) Moral philosophers ought to avoid being naïve Baconians about moral theory. Judgments by themselves do not announce the best theory.

The failure to notice the distinction between moral and theoretical disagreement partly accounts for why moral philosophers' debates are often heated. Moral philosophers theorize about morality, and morality matters, *really matters*, in a way that some other philosophical subject matters do not. It matters whether slavery is wrong or right. Being a philosopher, I have to admit it matters whether events are fact-like or object-like, or whether intentionality is the essence of consciousness, but those types of issues will never matter in the same way.⁷ The metaphysician or philosopher of mind may not be very interested in philosophizing about the wrongness of slavery, but they are all very interested in not being enslaved. We are all interested in morality, but moral philosophy does not matter as much as morality itself. Geophysics matter less than the planet Earth.

Like most good distinctions—distinctions worth making—the distinction between moral and theoretical disagreement is fuzzy, and there is overlap. Sometimes theoretical

⁷As I write this, *The Monist* calls for papers for issues devoted to such subjects.

disagreement does produce moral disagreement. Sometimes moral disagreement necessitates theoretical disagreement. For an instance of the latter, consider a person who judges every particular act of killing right and another who judges every particular act of killing wrong. These two will almost certainly not endorse the same moral theory. As for an instance of the former, the fascist Carl Schmitt and the libertarian Robert Nozick did not merely disagree on the correct way to organize their respective *theory* of the good society; they fundamentally disagreed about justice itself.⁸ It is impossible to endorse Schmitt's theory and not have profound *moral* disagreement with Nozick.

The only thing philosophers produce more of than theories is disagreement. About what do moral philosophers do most of their disagreeing? Certainly, moral philosophers argue over morality, but they spend far more of their time arguing over moral theories. Moral philosophers generally agree on how they ought to treat themselves, each other, their families, friends, neighbors, and strangers. Otherwise, they would not get along or survive in larger society. Yet, ethics conferences are convivial affairs. The fact that philosophers with deep theoretical disagreements can get along shows us that the theoretical aims of moral philosophy do not translate directly into practical effects. It also shows us that virtuous behavior does not translate directly into a moral theory.

⁸ Carl Schmitt, *The Concept of the Political*, trans. George Schwab (Chicago: University of Chicago Press, 1996 [1932]); Robert Nozick, *Anarchy, State, and Utopia* (New York: Basic Books, 1974).

Most major philosophers in the history of ethics were not trying to replace commonsense morality wholesale. They did not want to scrap folk morality, start over, and then tell people what to do. Kant, Sidgwick, and Aristotle began with the assumption that the average person is morally decent, and that much of commonsense morality is more or less correct. Immanuel Kant even believed that the average layperson was morally better than the typical intellectual.⁹ This is not to say that such philosophers were obscenely conservative. Rather, they wanted to understand morality, and had to start somewhere. They realized one can start only one place, and that is where one already is. (Some starting points make getting to one's destination hard, and some make it impossible, but one cannot even attempt to get to a destination starting from nowhere.) They believed rational reflection about moral practice could lead to insight, and such insight might call for revision, even major revision, of common sense. Some of the suggested revisions (e.g., women's equality, the abolition of slavery) were among humanity's greatest achievements when followed. However, they also effectively recognized that one could not begin such reflection on the presupposition that one knows absolutely nothing about morality. Imagine asking someone to produce a theory of guitar beginning with the presumption that she knows absolutely nothing about guitar.

On the morality of daily life, moral philosophers agree with each other, and with laypeople, far more than they disagree. There is no obvious connection between the moral theory or theory of justice they endorse and their degree of virtue. (Of the most virtuous philosophers I have known, one was a classical liberal; another was a socialist.)

⁹ Robert B. Louden, *Morality and Moral Theory*, pp. 120-124, especially p. 123.

Accordingly, disputes in moral theory are not the same as disputes about moral judgments. Such theoretical disagreement does not entail different moral behavior. That is not the point of moral theorizing.

3. MORAL THEORIZING AS A SUBJECT OF PHILOSOPHICAL INVESTIGATION

Norms constitute much of morality, but not all norms are moral.¹⁰ We use the term "morality" in at least two different ways. We can speak of the morals accepted by a culture, as when we say "Roman morality" or "Chinese morality". However, without contradiction, we can say that cultures are immoral, that they do not follow morality. For example, I can speak of "Stalinist morality", meaning that Stalinists subscribed to and followed certain norms. Yet saying "Stalinist morality *is immoral*", "Stalinists lacked morals", and "Stalinist ethics is a rejection of morality" are all perfectly coherent (and true) sentences.¹¹ So, we have at least two senses of morality. On one hand, the norms (or code of conduct) accepted by a person or group constitutes his or their morality.¹² On the other hand, sometimes when we talk about morality, we mean not just any set of accepted norms, but the norms that *ought to be accepted*. When we describe Stalinists as

¹⁰ Cf. the discussion in Robert Nozick, *Invariances* (Cambridge, MA: Belknap, 2001), p. 238.

¹¹ Cf. Bernard Gert, *Morality: Its Nature and Justification* (New York: Oxford University Press, 1998), p. 10.

¹² Of course, not all of the norms I accept are moral norms in this sense. I accept certain norms of guitar playing, but these have nothing to do with morality.

immoral, we do not mean to suggest that they failed to follow norms. What we mean is that they followed the wrong norms. A culture's morality is one thing; morality simpliciter is something else.

With that distinction in mind, consider how widely morality (in both senses) is studied. At the typical university, one finds a wide variety of theorists studying the phenomenon. Anthropologists collect data on the norms accepted in different cultures. Economists describe how unintentional large-scale order emerges from the actions of individuals as they seek to promote their ends. Psychologists and neuroscientists study how the brain functions during decision-making. Rhetoricians study how moral terms are used in the effort to persuade.

Philosophers, on the other hand, investigate morality in a very different way.

A historian can tell you how people came to believe, over time, that slavery is wrong. A moral philosopher, however, asks if slavery is in fact wrong, and what makes it so. A sociologist can report people's moral attitudes to you. A moral philosopher, however, seeks to determine what, if anything, can *justify* moral attitudes. Social scientists describe norm-following behaviors. That's their job. Moral philosophers, on the other hand, ask about morality itself. That's their job. Social scientists do not generally distinguish between morality in the sense of accepted norms and morality in the sense of the norms that ought to be accepted. For a host of social scientific questions, it simply does not matter whether the norms accepted by a culture actually ought to be accepted. Many social scientists would think the question "What norms ought to be accepted?" is nonsensical. (Whether the question is nonsensical, however, is not

something that the social scientist has any expertise in determining. It's a philosopher's question.) For a moral philosopher, on the other hand, the question of which norms ought to be accepted is central.¹³

Moral philosophers produce *moral theories*. Among other things, moral theories are systematic accounts of what morality is, what actions are right or wrong, what character traits are virtuous or vicious, what things are good or bad. The moral theorist typically tries to determine the structure of moral norms and to discover underlying, explanatory principles. Just as the physicist looks for principles governing matter, the moral philosopher looks for principles governing morals.

Of course, not all philosophers study morality. Philosophical theories are not the only type of theory. There are also scientific theories, mathematical theories, literary theories, music theories, theories of how to play various sports, theories of leadership and management, conspiracy theories, theological theories, and so on.

A physicist studies the underlying nature of mass and energy. Chemists study the properties of certain forms of matter and how matter combines. Biologists study life. All of these scientists produce scientific theories. Scientific theories are (to put it crudely) systematic, unifying, predictive accounts of various parts of the world.

Scientific theories in their own right are interesting objects of philosophical study. While a physicist asks about the nature of matter and the laws regulating it, the philosopher of science asks about the theories the physicist produces. The philosopher of

¹³ Cf. Henry Sidgwick, *The Methods of Ethics* (Indianapolis: Hackett Press, 1981 [1907]), p. 2.

science does not produce scientific theories—she theorizes *about* scientific theories. The good philosopher of science does not confine herself to theories; she philosophizes about experimentation, engineering, and other aspects of scientific practice.¹⁴ She asks questions such as: What is the nature of scientific explanation? What makes one scientific theory better than another? What are the roles of experience, experiments, and speculation in scientific theorizing? Are scientific theories meant to be true descriptions of the universe, useful fictions, or something else? How do we know?

These questions cannot be answered without a firm grasp of science. The successful philosopher of science does not need to be a scientist herself, but she needs to know much about science before she can start asking about scientific theories in their own right.

As of yet, there is no equivalent of the philosophy of science for moral theories. There has not been a parallel, organized field within philosophy that takes *moral theories*—the accounts of morality produced by philosophers—as objects of philosophical investigation. Moral theorists theorize about morality, but they do not systematically theorize about moral theory itself.

There is a sub-branch in philosophy known as metaethics. Metaethics asks about the *semantics, epistemology, and metaphysics* of *moral judgments and terms*.¹⁵ When I apologize to someone, I might say something like, "It was wrong of me not to call you."

¹⁴ E.g., Ian Hacking, *Representing and Intervening*, chapters 9-16, especially chapter 9.

¹⁵ Mark Timmons, *Moral Theory* (Lanham, MD: Rowman and Littlefield, 2002), pp. 17-20.

The statement "It was wrong of me not to call you" is a moral judgment. It evaluates an action of mine. What is the meaning of the judgment "It was wrong of me not to call you"? Could that sentence be true or false, or does it merely express an attitude? Is moral language descriptive or fact stating, or purely evaluative? The answer to those questions is a matter of *moral semantics*. How do we know whether "It was wrong of me not to call you" is a good judgment or not? Is there such a thing as moral knowledge? The answer to that question is a matter of *moral epistemology*. Lastly, how does the world have to be for such judgments to be correctly made? When I say that something is wrong, does that commit me to believing that there is a property of *wrongness* out there in the world, much like there is a property of *hardness*? Is morality in some way real? Questions like this are a matter of *moral metaphysics*.

Moral theory tries to answer *moral* questions, while metaethics tries to answer *nonmoral* questions about morality.¹⁶ Good moral theory helps us understand those moral judgments we ought to endorse, those character traits we ought to develop, or those goals we ought to achieve. Metaethics helps us understand the metaphysical nature of morality. This distinction is difficult to draw, but an analogy will help. The philosophy of mind seeks to uncover the metaphysical nature of consciousness. E.g., it asks if consciousness is a physical or nonphysical phenomenon. Is it a function of physical processes, identifiable with particular physical states, or something else altogether? In some way, philosophy of mind is telling us about the nature of consciousness. Its questions and answers are different from, but compatible with, the findings of psychology

¹⁶ Timmons, *Moral Theory*, p. 19.

and cognitive science. Psychology and cognitive science ask how the mind works, but not what its metaphysical nature is. Similarly, moral theory asks how morality works, but metaethics asks what its metaphysical nature is. Of course, good psychology and good philosophy of mind will in some ways overlap, and will proceed while paying attention to each other. Similarly, good moral theory and good metaethics will in some ways overlap, and will proceed while paying attention to each other.¹⁷ The relationship between moral theory, morality, and metaethics is complex.

An equivalent of philosophy of science, that takes moral theorizing as its subject, would be *moral metatheory*. Generally, moral metatheory is not about the semantics, ontology, or epistemology of moral *judgments* and *terms*. Rather, it concerns the semantics, ontology, epistemology, and *methods* of moral *theories* themselves. The questions of moral metatheory include: What is the nature of explanation in a moral theory? What makes one moral theory better than another? What exactly does the moral theorist know? What are good philosophical methods? How, if ever, can we confirm moral theories? What are the roles of experience, intuition, and judgment in moral theorizing? Are moral theories meant to be true descriptions of morality, or merely useful fictions? Does one need a moral theory in order to be a good person? What can we learn about moral theories by investigating the history of moral theorizing? Are moral theories even possible? How do we know?

¹⁷ For an argument to this effect, see Stephen Darwall, "How Should Ethics Relate to (the Rest of) Philosophy?: Moore's Legacy", *The Southern Journal of Philosophy* (2003), vol. XLI, pp. 1-29.

One must keep in mind four important distinctions that we have already made, distinctions that are glossed over even by professional philosophers. Morality is not the same thing as a moral theory. Rather, morality is the subject matter studied in a moral theory. Moral theories are not the same thing as metaethical theories. Lastly, metaethical theories are not the same as moral metatheories. Metaethical theories are about the moral judgments and terms found inside the practice of morality, but moral metatheories are about moral theories. In practice, the distinctions are not so neat. There is significant overlap between morality, moral theory, metaethics, and moral metatheory. For reference, here is a chart:

Type of Theory	Subject Matter
Scientific	Nature
Philosophy of science	Scientific theory and practice
Moral	Morality
Metaethical	Moral judgments and terms
Moral Metatheory	Moral theory

Moral metatheory is largely a new field. Many of the issues I will explore are similar to issues in the philosophy of science, but as scientific theories and moral theories are quite different, not only in subject matter but in form, one cannot simply adopt one's favored philosophy of science as one's moral metatheory. Some of the issues I will examine have been discussed at length. There have been explicit debates about how moral theory links with moral practice, and whether moral theory is a worthwhile endeavor. In arguing over theories, philosophers will appeal to and discuss the standards

of moral theory. For example, Brad Hooker, in his *Ideal Code, Real World* (a book of moral theory), advances desiderata that moral theories are meant to satisfy.¹⁸ When I say that moral metatheory is largely a new field, I do not mean to suggest that the ground we will explore has never been trodden upon before. Others have briefly crossed this territory in search of different destinations. However, this territory is our destination. We are here to stay.

Moral metatheory is a challenging field. Just as the philosopher of science must have a firm grasp of science, so the moral metatheorist must have a firm grasp of moral theory. We are philosophizing about philosophy. In order to avoid making obvious mistakes as made in a parallel field, the moral metatheorist must be strongly acquainted with the philosophy of science, and this field has its own challenges. Successful moral metatheory requires finding connections (and differences) over a giant terrain. I offer here my best attempt.

It seems relatively unproblematic why someone would want to theorize about morality. Why someone would want to theorize about moral theory is less obvious. Here I will offer only some answers. For me, personally, it is because I find the subject matter inspiring and enthralling. It is simply something I must do. Less autobiographically, this is an attempt to ask just what it is moral philosophers, and philosophers more generally, are doing. Understanding their activity is worthwhile.

¹⁸ Brad Hooker, *Ideal Code, Real World* (Oxford: Oxford University Press, 2000), p. 4.

CHAPTER TWO

MORAL THEORY AND ITS DISCONTENTS

1. ANTI-THEORY, REDUX

Richard Posner, one of the latest critics of moral theory, claims moral theorizing has two goals. First, he thinks the main “ambition” of theorists such as Thomas Scanlon or Frances Kamm is “to change people’s moral beliefs to the end of changing their behavior...”¹ According to Posner, the principal purpose of moral theory is moral persuasion. He claims the second major purpose of moral theory is to prescribe particular actions. Theorists aim to tell us what to do in particular situations. Theorists seek to discover principles or concepts from which they can “*deduce* the answers to contemporary moral questions...”² Moral theorists are preachers, endeavoring to teach others right from wrong.³

However, Posner claims that moral theory does not achieve either of these goals, and is thus worthless. He says, “There is nothing to academic moralism”.⁴ Moral theory is plagued by abstraction; it is “too lofty ever to touch the ground and resolve a moral

¹ Richard A. Posner, *The Problematics of Moral and Legal Theory* (Cambridge, MA: Harvard University Press, 1999), p. 38.

² Posner, *The Problematics of Moral and Legal Theory*, p. 48.

³ Posner, *The Problematics of Moral and Legal Theory*, p. 15.

⁴ Posner, *The Problematics of Moral and Legal Theory*, p.12.

issue".⁵ Because of this abstraction, it cannot prescribe particular duties. Its principles are devoid of substantive content. Moral theory is futile; it "has no prospect of improving human behavior" because its arguments are feeble.⁶ Arguments do not persuade people; "emotional appeals that bypass our rational calculating faculty" do.⁷ In fact, Posner thinks moral theory is worse than useless. It becomes pernicious when employed for legal decision-making.

The overall structure of Posner's argument is this:

1. Moral theory does not convince people to act morally better.
2. Moral theory does not provide a useful way of ascertaining what one ought to do in particular situations.
3. Therefore, moral theory is worthless.

Many theorists do not take Posner seriously.⁸ Partly, this is Posner's fault. His writing borders on diatribe and he superficially assays many complicated issues (such as motivational internalism⁹). Still, most moral philosophers recognize that there is something disturbingly correct about premises 1 and 2.

⁵ Posner, *The Problematics of Moral and Legal Theory*, p. 19.

⁶ Posner, *The Problematics of Moral and Legal Theory*, p. 7.

⁷ Posner, *The Problematics of Moral and Legal Theory*, p. 42.

⁸ One notable exception, of course, is Steven L. Ross, "Taking Posner Seriously", *The Philosophical Forum* XXXII:1 (2001), pp. 1-23.

⁹ Posner, *The Problematics of Moral and Legal Theory*, p. 39.

Moral theorists are academic experts. They had better be, since universities pay them to write and teach. But what exactly do they know? They do not seem to be expert moral performers the way Andre Agassi is an expert tennis player. They do not seem to be expert moral coaches the way Bill Belichick is an expert football coach. In improving people's behavior, Dr. Phil probably gets better results than Dr. Kant. Moral theorists do not seem to be expert critics, either. Roger Ebert cannot make good movies, but he can expertly distinguish good movies from the bad. In contrast, moral theorists do not seem to be experts in distinguishing particular good actions from the bad.

Posner exaggerates his case, but he is largely correct about 1 and 2. Yet, he is mistaken in thinking this shows moral theory is worthless. Moral theory is not typically necessary to be a good person, nor is it typically helpful. But that is not theory's point. The goal of moral theory is to produce a certain kind of knowledge. Theory answers a number of related questions about morality. Moral theory is about a normative domain, but its goal is theoretical. It tries to provide an explanation of that domain. The theorist's knowledge is different from that of the expert moral performer, coach, or critic. The point of this essay is to show that even if moral theory could not convince people to act better and could not provide a formula for right action, it still offers a valuable form of knowledge.

2. WHAT IS MORAL THEORY FOR?

Posner thinks the point of moral theory is to tell us what to do and convince us to do it. I think moral theory aims at something else. To illustrate this, I will turn to some

thought experiments showing that even if we always knew what to do and were motivated to do it, moral theory might provide additional knowledge.

Suppose God hands you a magic book, the Book of Judgments. The Book of Judgments lists all the correct moral evaluations of every object of assessment. One need only point to an object, action, person, etc., and the book will produce the correct judgment about that object. Also, because knowing whether something needs to be judged requires good judgment, the book not only signals judgments on request, but also signals when a judgment is called for.

If you had the Book of Judgments, would there be any point to having a moral theory? Posner might think that at best, moral theory is shorthand—an easy memorable formula that produces the same judgments as the Book. However, suppose God simultaneously boosts your IQ, allowing you to memorize the entire Book of Judgments. You would then not even need a convenient summary of the book. You certainly would not need some grand principle from which to derive particular judgments. Would this mean that you would have no need of moral theory?

Moral theory is a poor surrogate for the Book of Judgments, but it is not meant to be a surrogate. Let me illustrate this with a parallel thought experiment. Suppose God also implants in your mind a fully accurate, completely detailed Map of the Universe, which allows to you to visualize all of space-time simultaneously. In this case, you would not need a scientific theory to make future predictions or to determine if various entities exist. You could see what will happen tomorrow and you could see all the fermions and bosons. However, if God then tells you that the universe is best

characterized by Einstein's theory of general relativity, you have learned something new, something that was not on the Map. The Map of the Universe is a description, but not an explanation. The Map tells us what there is; general relativity tells us how it fits together.

The Book of Judgments is a map of the moral universe, but one still learns something extra, something not on the map, in acquiring a moral theory. Even if one knows all the judgments, one learns something new in discovering that these judgments fit together via the Categorical Imperative. The Book does not tell you whether morality is orderly or chaotic, characterizable in general terms or radically particular. It does not tell you how moral truths are known or what the nature of these truths is. The Book tells you the substance of morality, but does not explain its essence. It tells you what is right and wrong, but not what rightness and wrongness ultimately consist in. It does not tell you what makes a norm a moral norm, as opposed to legal or aesthetic, nor does it tell you how normativity is possible. The goal of moral theory is to explain these issues, the questions not answered by the Book of Judgments.

Thomas Scanlon characterizes the attempt to answer these sorts of questions as "Philosophical Enquiry into the nature of morality".¹⁰ Philosophy Enquiry seeks to answer questions such as:

When we address our minds to a question of right and wrong, what kind of question are we considering? What kind of fact or truth, if any, could the answer represent? How can we discover truths of this kind? Why do they

¹⁰ T. M. Scanlon, "The Aims and Authority of Moral Theory", *Oxford Journal of Legal Studies* 12: 5 (1992), pp. 1-23, here p. 4.

have such importance to us? That is, what kind of reason for action do they provide and how could these reasons have the kind of priority over other considerations which morality is commonly thought to have?¹¹

According to Scanlon, and I agree, these are the questions Kant and Aristotle sought to answer in the *Groundwork* and the *Nicomachean Ethics* respectively. Consider Scanlon's own contractualist principle, in its latest version, that an act is wrong just in case people seeking principles for the general regulation of conduct could reasonably reject any principle permitting that act.¹² Scanlon says this is not thickly substantive principle telling us what to do. It is not something from which to derive particular duties. Rather, it is a "philosophical account of the nature of moral wrongness".¹³ I do not discuss Scanlon here because he is an unusual counterexample to Posner. Posner himself identifies Scanlon an academic moralist seeking to prescribe behavior.¹⁴ Yet, that is not how Scanlon understands himself.

In fact, many of the theorists Posner mentions explicitly deny they are seeking to change behavior or prescribe policies. For example, Robert Nozick says that his book *Anarchy, State, and Utopia* "is not a political track; it is a

¹¹ Scanlon, "The Aims and Authority of Moral Theory", p. 5.

¹² T. M. Scanlon, *What We Owe to Each Other* (Cambridge, MA: Harvard University Press, 1998), p. 4.

¹³ Scanlon, "The Aims and Authority of Moral Theory", p. 5.

¹⁴ Posner, *The Problematics of Moral and Legal Theory*, p. 5.

philosophical exploration.”¹⁵ As Jeremy Waldron points out, Karl Marx and G. W. F. Hegel denied their theories were for recommending policies.¹⁶ The value of political philosophy, though it is more “applied” than general moral theory, lies not so much in its conclusions but in its arguments, because these show how different moral ideas intersect and interact.¹⁷

Posner conflates the question what is right action with the question which actions are right. Of course, the questions are not unrelated. The answer to the first question cannot be morally neutral, as Scanlon says, but it is not meant to fix the full content of morality either. Rather, the point of answering the first question is to explain something deep and interesting about morality. In fact, moral philosophy operates at three or more levels. The question of what is right action is different from the question of whether abortion is presumptively right or permissible. There is a third question, still: Is it right for Joan to have an abortion now? In moving from the first question to the third, we are moving from moral theory to applied ethics, from philosophical explanations about the nature of morality to attempts to evaluations of particular actions. Posner complains that philosophers’ principles are vague, abstract, and empty, because these principles cannot successfully sort all particular right actions from particular wrong actions.

¹⁵ Robert Nozick, *Anarchy, State, and Utopia* (New York: Basic Books, 1974), p. xii.

¹⁶ Jeremy Waldron, “What Plato Would Allow”, *NOMOS XXXVII* (1995), pp. 138-180, here p. 160.

¹⁷ Waldron, “What Plato Would Allow”, p. 166.

But that is not the point of theory. He is right to say that moral theorists are not particularly good at answering low-level questions such as whether Joan should have an abortion today. But, again, that is not what moral theory is for.

This is just a sketch of the content of moral theory. I have more to say below. For now, let us turn to the issue of whether the point of theory is to convince people to be moral.

Posner portrays moral philosophers as rhetoricians manqué. They seek to persuade people to act morally, but they fail to realize this aim. They are ineffective preachers. Long, hairsplitting arguments do not have the motivational force of emotional pleas and pictures of starving children.

There is something so obvious about the philosophy's lack of motivational force that it is strange to accuse philosophers of trying to motivate others. Philosophers are smart people; they generally know how to take effective means to their ends. Philosophers know that their moral theories will be published by academic presses in prints of only a thousand copies. A typical philosopher would be delighted to learn that two hundred people read her work. Is it at all plausible to think most philosophers are trying to induce moral behavior among the masses?

Posner misunderstands why philosophers try to answer the question why be moral. They try to answer the question because it is philosophically interesting—so interesting, in fact, that many articles are devoted not to answering the question but to understanding what would count as an answer. (For

example, Scanlon argues that “Because it is right” and “Because it is in your self-interest” are both bad answers. Somehow a good answer has to be in-between those two.¹⁸) One of the first examples of the attempt to answer the question is Socrates in the *Republic*. Note that Socrates, at least, was manifestly *not* trying to persuade anyone to be moral. Socrates, Glaucon, and Thrasymachus all understand that Glaucon will continue to act morally and Thrasymachus will continue to act immorally regardless of whether Socrates makes his case. Socrates is not looking to change anyone’s behavior; he is looking for an answer to a question. (Presumably, Plato, in writing the *Republic*, is doing the same.)

Posner seems perplexed that philosophers could be doing anything other than trying to influence behavior. For instance, he claims that philosophers aim to change others’ moral beliefs with the goal of changing their behavior, otherwise there would be no reason to try to change their beliefs.¹⁹ However, it is not clear that philosophers are even trying to change others’ beliefs. Rather, they are trying to produce good arguments for interesting positions. They know better than to think they will win many converts. Posner says that if philosophers are not trying to alter behavior, they must simply be “dilettantes or careerists”.²⁰ Posner is mistaken about philosophers’ aims. They are primarily concerned with answering

¹⁸ Scanlon, “The Aim and Authority of Moral Theory”, p. 6. See also Scanlon, *What We Owe to Each Other*, chapter 5.

¹⁹ Posner, *The Problematics of Moral and Legal Theory*, p. 38.

²⁰ Posner, *The Problematics of Moral and Legal Theory*, p. 40.

certain fundamental questions about morality. Posner's mistake stems from his pragmatism. He is such a thoroughgoing pragmatist that he does not separate the issue of truth from usefulness.

Though I have many pragmatist tendencies myself, it seems to me that there could be useless knowledge about morality. Consider a character that I will call Mr. Virtue. Mr. Virtue is a perfectly virtuous person, who always does the right thing for the right reason and feels the right way about it. Mr. Virtue could learn something from moral theory even though that theory would have no effect on his behavior.

Mr. Virtue, like the average person, has moral beliefs but lacks a moral theory. (Posner himself grants that moral beliefs do not constitute a theory.²¹) Moral theories are systematic accounts of the nature of morality. The moral theorist attempts to determine how norms fit together and to discover underlying, explanatory principles. The moral theorist tries to explain normativity itself and why moral norms have their special force. Mr. Virtue can be virtuous without these explanations.

Mr. Virtue may know various low-level principles and be able to give reasons for what he does. His dispositions, attitudes, responses, and particular judgments constitute a kind of understanding of morality—the most practical kind—but not a moral theory. It remains an open question whether the underlying structure of Mr. Virtue's morals is Kantian, utilitarian, contractualist, or whatnot,

²¹ Posner, *The Problematics of Moral and Legal Theory*, p. 16.

if it has any structure at all. The moral theorist who discovers this structure is not likely to be Mr. Virtue. She will thus lack some understanding of morality. Yet, she will also understand something Mr. Virtue does not.

In parallel, Jimi Hendrix was an excellent blues guitarist, but he did not know the theory of blues guitar. This does not mean he did not know what he was doing. Also, many do know guitar theory, but are not expert performers like Hendrix. The person with theory understands something that Hendrix does not, but Hendrix understands something that the person who only has theory does not. The person who has mastered both theory and practice understands even more. (Hendrix agreed; he regretted not knowing music theory.) A guitarist with theory understands not merely how to perform, but the deep, underlying structure of music. With theory, one can articulate, at a fundamental level, *why* certain musical ideas make sense and what musicality is all about. Moral theory is similar.

3. THE ELEMENTS OF MORAL THEORY

Above, I gave a brief overview of the aims of moral theory. I began to argue that even if Posner is right that moral theory has no practical value, it still has theoretical value. To complete this argument, I will first provide a fuller account of the contents of moral theory. In the next section, I will provide a more detailed account of the types of explanations moral theory provides.

There is no clear demarcation between having a theory and not. There is reasonable disagreement about what is in a theory and what is ancillary.²² However, I have asserted that the typical good person lacks a theory. To avoid trivializing this assertion, I need a broad conception of theory. The terms "moral theory" and "ethics" are used in broader and narrower senses. I use a broad sense below, but I am not the only one to do so.²³ Since I am evaluating the expertise of moral theorists, I want to analyze what we expect moral theorists to know. When departments advertise for a specialist in ethics in the American Philosophical Association's *Jobs for Philosophers*, they work with this broad conception.

There is another reason to use a broad notion of "moral theory": to parallel the use of the term in other fields. In science, a law is (roughly) a generalization that describes what occurs under certain conditions. For instance, the law of universal gravitation describes the relationship between force, mass, and distance. Scientific theories, however, explain the laws and/or embed them within a deeper unifying framework. Following this parlance, moral principles are the analog of scientific laws, while moral theories are the analog of scientific theories. This means distinguishing between theory and principle, and thus working with a broad conception of theory as I favor.

²² Cf. Robert B. Louden, *Morality and Moral Theory* (Oxford: Oxford University Press, 1992), pp. 140-141.

²³ E.g., see Norman Daniels, "Wide Reflective Equilibrium and Theory Acceptance in Ethics", *The Journal of Philosophy* 76:5 (1979), pp. 256-282.

First, a theory contains a set of principles, laws, and/or functional explanations. This is the core of any theory. Typically, these principles are presented as *criteria of right* or *criteria of good*, i.e., they are to specify, at a fundamental level, what the right and the good are. Some philosophers of science write as if a scientific theory simply is a set of laws. Similarly, some philosophers identify moral theories simply as sets of moral principles. These are narrow conceptions of theory.

A theory asks what right action is and what good things are. Asking what right action is is not quite the same as providing a set of rules for living one's life. Epistemologists provide accounts of epistemic justification; they explain what justified belief is. However, no one expects that reading an epistemology book would allow him to *infer* which beliefs are justified. Having an explanation of justified belief is not enough to justify all of my beliefs or to derive all justifiable beliefs. Similarly, in economics, one of Posner's favorite fields, the dominant theory of prices is marginalism. Marginalism says that on a normally functioning market, the price of a unit is its value (its utility) to the marginal buyer. This theory is well confirmed, but it does not allow us to predict particular prices. (If anything, it explains why we cannot predict prices.) Moral theory's explanations follow this trend. Kant's explanatory analysis of right action is the Categorical Imperative, but we cannot infer from this explanation which actions are right in every circumstance.²⁴

²⁴ Aside from a few illustrations in the *Groundwork*, Kant does not make many derivations from the Categorical Imperative. Rather, in the *Metaphysics of Morals*, he lists many basic duties, which are shown to fit into the model of the Categorical

Second, a moral theory provides an account of the nature and scope of moral judgments and norms. In particular, a moral theory explains what makes a moral judgment a *moral* judgment, rather, than say, an aesthetic judgment or a mere expression of preference. Moral theories typically explain the difference between moral norms and other types of norms: legal, prudential, religious, norms of etiquette, norms of musical or sports performance, and so on. Note that providing an account of the nature and scope of judgments and norms is not the same as providing a set of explanatory principles. This distinction between moral and non-moral norms can be made before offering any principles.²⁵ For instance, one might distinguish moral versus non-moral norms by saying that moral norms are universal, overriding, and public.²⁶ This claim is compatible with almost any explanatory moral principles. Also, one can differentiate moral and non-moral normative judgments while denying that there are any moral principles. Moral

Imperative, but not *derive* from it. Kant's list of duties shapes the Categorical Imperative as much as the Categorical Imperative shapes the duties.

²⁵ E.g., in his ethics textbook, Louis Pojman does an excellent job distinguishing morality from religion, law, and etiquette. However, this distinction is established well before Pojman describes any moral principles. See his *Ethics: Discovering Right and Wrong* (Belmont, CA: Wadsworth, 1999), pp. 3-9.

²⁶ I take this list from Pojman, *ibid.*

particularists deny that there are moral principles, but they do distinguish between morality and other normative domains.²⁷

The third component of theory is closely related to the first two. A theory typically provides a formal definition of morality. E.g., following but modifying a definition provided by Bernard Gert, I might claim that morality is a "public system [of norms] applying to all rational persons".²⁸ This is a formal definition because on its face it does not entail any substantive principles. However, as Scanlon has said, for this formal definition to count as a definition of morality, rather than something else, it must bear the right relationship to what we characteristically understand as morality.²⁹ Thus, a formal definition has to be connected to substantive moral ideas.

Providing a formal definition is not the same as providing explanatory principles of right action, because a particularist could accept this definition of morality but continue to deny the existence of principles. Providing a formal definition is not the same as distinguishing moral judgments and norms from non-moral judgments and norms. A theorist might claim legitimately that while this distinction can be made, morality does not admit of a definition. (This theorist would thus deny that definitions

²⁷ See Jonathan Dancy, *Morals without Principles* (New York: Oxford University Press, 2004), pp. 3, 43-45.

²⁸ Bernard Gert, *Morality: Its Nature and Justification* (New York: Oxford University Press, 1998), p. 13.

²⁹ Scanlon, "The Aims and Authority of Moral Theory", p. 6.

are a proper part of moral theory, though she might say that many moral theories mistakenly offer such definitions.)

Fourth, a theory provides an account of the nature of moral agency and moral rationality. This typically includes an account of relevant moral psychology, such as how a moral agent "perceives" moral distinctions. Bernard Gert, e.g., argues that a moral theory must elucidate what makes a moral agent an agent.³⁰ Such an account should be psychologically and anthropologically informed. This account should explain which types of beings can be rational and which cannot. It should explain distinctions between normative categories such as rationally required, rationally allowed, and rationally prohibited.³¹ Typically, a theory will analyze related concepts, such as impartiality.

A fifth, related element proper to moral theory is an account of reasons. While a virtuous person can easily relate what substantive reasons one has in a given scenario, a theorist has a different task: to explain what reasons are. How can there be such things as reasons, and how do reasons connect with rationality?

An account of reasons normally attempts answer what Christine Korsgaard calls "the normative question".³² A common task in the history of philosophy has been explaining why moral principles have their characteristic force. How is normativity possible at all? Why must one do what one must? These are perhaps the most difficult

³⁰ See Bernard Gert, *Morality*, pp. 3-4.

³¹ Gert, *Morality*, chapter 2.

³² Christine Korsgaard, *The Sources of Normativity* (Cambridge: Cambridge University Press, 1996), pp. 8-18.

questions a philosopher faces. Perhaps Korsgaard's answer to these questions is correct. Mr. Virtue need not know her answer; the question need not have even occurred to him.

Sixth, a theory must explain not merely moral agency, but moral patiency. To whom and what are moral duties owed? Again, a virtuous person can answer that question in his characteristic way: "You owe your brother gratitude because he helped you." However, for a theorist, the question is more formal. What is it to be the type of thing that can be owed duties in the first place? The answer to this question is not quite the same as providing a substantive list of particular beings to whom one owes duties. For instance, a theory might claim that what makes something a moral patient is that things can go better or worse for it. (This explanation may be established and defended without providing a substantive account of better or worse ways things can go.) This explanation is still philosophical; it is not the same as giving a substantive list of moral patients.

Seventh, theories analyze morally significant concepts. For example, a theory might provide an analysis of "liberty", "harm", or "desert". Analyzing these is not the same as advocating a principle. A theorist might successfully analyze "desert" but be neutral as to what to do about desert.

Eighth, theories contain explanatory *structures*. We might even say that theories *impose* structures. In chemistry, the periodic table provides structure to the theory of the elements. The structure is a map of sorts, but it is also explanatory, as an element's position on the map tracks its chemical properties. (Alternatively, we might say that these properties determine why the periodic table has one shape rather than another.)

Moral theories not only categorize moral phenomena, they often fit them within an explanatory structure. E.g., in deontic logic, it is common to construct a "deontic square of opposition", which shows the relationships holding between categories such as "forbidden to A", "permitted to A", "permitted not to A", and "required to A". The point of having these categories in a square structure is to show that some are incompatible with one another, some entail others, and so on. Different theories have different structures. E.g., Paul McNamara argues that the deontic square is incomplete; deontic categories are best related via a tetradecagon.³³

The ninth component is closely related. A theory is not merely a set of principles, but normally is a set of theoretical categories and distinctions. E.g., Kant's moral theory is not merely the Categorical Imperative, it is an analytical framework that includes such concepts (theoretical terms) as pure practical reason, synthetic versus analytic, categorical versus hypothetical, maxim, apodictic versus assertoric versus problematic, autonomous versus heteronomous, and so on. Theories do not merely contain different theoretical terms; they claim that these terms refer to something.

In an often misunderstood chapter, Thomas Kuhn claims that scientists who accept different theories inhabit different worlds.³⁴ One thing Kuhn means is that

³³ Paul McNamara, "Must I Do What I Ought? (Or Will the Least I Can Do Do?)" in José Carmo and Mark Brown (eds.), *Deontic Logic, Agency, and Normative Systems* (Berlin: Springer Verlag, 1996), pp. 154-173.

³⁴ Thomas S. Kuhn, *The Structure of Scientific Revolutions, 3rd Edition* (Chicago: University of Chicago Press, 1996), pp. 111-135.

scientists' different background expectations cause them to see the world differently. More importantly, accepting different theories means positing different entities and categorizing the world differently. In effect, to accept a theory is not merely to accept laws, but to accept an ontology.

Following Deidre McCloskey's philosophy of economics, we may wish to say that the tenth and eleventh parts of a theory are a *stories* and *metaphors*. (Metaphors and models are more or less the same.) An example of an economics story is, "Once upon a time we were poor, then capitalism flourished, and now, as a result, we are rich."³⁵ An example of an economics metaphor is to say that some market is a demand and a supply curve.³⁶ McCloskey argues, providing numerous examples, that theories of all sorts comprise stories and metaphors. Moral theory is no different. On the broad conception of theory I am using, Rawls' political theory is not merely his two principles of justice. The original position is part of the theory, not just a device for illustrating and defending principles. The original position tells a story: First we had a problem of determining fair terms of cooperation to govern a society marked by both identity and conflict of interest. Then we imagined some bargainers unanimously agreeing under fair circumstances to principles to govern the people they represent. Then the problem was solved. Kant's moral theory uses metaphors: Wrongness is treating people like tools.

³⁵ Donald N. McCloskey, *If You're So Smart* (Chicago: University of Chicago Press, 1990), p. 1.

³⁶ *Ibid.*

Next, we may add that a full theory of ethics will include *metaethics*. Though it is helpful to speak of the differences between normative ethics and metaethics, in practice, most moral theories have both normative and metaethical components.³⁷ Metaethics attempts to answer many questions: Are good moral judgments true propositions, expressions of attitudes, or something else? Is moral knowledge possible? What do moral terms mean? Are there moral facts?

Finally, as a thirteenth component, we may add that a moral theory normally includes an articulation and defense of substantive moral principles and judgments about particular actions. E.g., Kant does not merely argue that right actions are those that could become law in a hypothetical kingdom of ends. He argues that beneficence and self-respect are virtuous, while pedantry and suicide are vicious. Just as theories normally account for what reasons are, they also often articulate some of the reasons the typical person actually has.

A theory can be partially disconnected from its substantial judgments. Neo-Aristotelians accept variations on Aristotle's theory, but they do not accept the existence of natural slaves. Neo-Kantians do not regard masturbation as degrading. Modern utilitarians may view England's colonization of India differently from Mill. Substantive

³⁷ See Stephen Darwall, "How Should Ethics Relate to (the Rest of) Philosophy?: Moore's Legacy", *The Southern Journal of Philosophy* XLI (2003), pp. 1-29. Darwall argues that normative ethics should include metaethical and other philosophical issues. I would add that Darwall's argument does not merely show that moral theorists should attend to metaethics, but explains why most in fact do.

judgments are quite ancillary to a theory. If Posner is correct, moral theorists are not superior at making such substantive judgments using their theories. But this does not show theory is worthless, for as we have seen, it has much more work left over to do.

I would not claim this list is complete. Nor would I claim that each component is necessary for something to be called moral theory. Moral theories are not all one thing. Still, a common cluster of issues, such as those listed above unites moral theories.

When I claim that the average moral agent and the virtuous agent do not have moral theories, I mean that they normally will be missing most of the elements above in their worldviews. Many of the questions will never have occurred to them, or only occurred to them in superficial ways. Still, there is no sharp boundary between having a moral theory and merely having moral beliefs.

4. MORAL EXPLANATION

As we have seen, moral theory has many components and is an attempt to answer many different questions. Posner complains that philosopher's arguments in support of their answers are not persuasive and concludes that philosophy is therefore useless. However, philosophical argumentation is not principally aimed at persuasion. Convincing others of a position typically requires not a sound argument, but good rhetoric. Posner himself is probably aware of this firsthand. Take economics, a subject dear to him. Economic conclusions are counterintuitive. E.g., the economics of rent control are well understood. The overwhelming majority of economists believe that rent control is a bad policy, which hurts the poor by reducing the supply and quality of housing. Yet, it is

difficult to convince laypersons (or non-economist intellectuals) using economic reasoning that rent control is bad policy. This does not mean the arguments or the conclusions are bad. Proof need not persuade.

Proof and persuasion are different. There is overwhelming evidence in favor of the theory of evolution. At present, any person committed to seeking the best theory in biology should accept evolution. However, many people are not committed to seeking the best theory, and thus they ignore or deny the overwhelming evidence with ease. The arguments for evolution are not as convincing as the arguments for creationism, not because the arguments for creationism are sound, but because they indulge our irrationality.

Following Nozick, we should say that philosophy seeks explanations. The purpose of philosophy is to render things “coherent and better understood.”³⁸ What sorts of explanations does moral theory provide? I have already sketched an account of the typical components of theory and the questions it tries to answer. Any account of explanations will be contentious. For better or worse, my account invokes Aristotle, not because Aristotle's view is uncontroversial, but because it is familiar and general. At present, there is no uncontroversial account of explanation in any field. Probably, whatever winds up being the best account of explanations in general will also supply the framework for the best account of moral explanations.

³⁸ Nozick, *Philosophical Explanations*, p. 8.

Aristotle posited four types of explanations (*aitia*).³⁹ One can explain something by describing 1) the matter out of which it is constituted, 2) the cause of the thing, 3) the purpose of the thing, and 4) the form, structure, or essence of the thing. These are called material, efficient, final, and formal explanations, respectively. With the Book of Judgments, one knows much about morality, but one does not know its underlying matter, its pattern or essence, what brought morality into being, or what purpose morality serves, if any. It takes some theory to discover that. Certain theoretical tasks remain.

A metaethical theory of moral properties can be called a material explanation of morality. When we ask about the matter of morality, we might be inquiring about its metaphysical nature. When we ask what makes a judgment such as "Slavery is wrong" true (if it is the type of thing that can be true), we might refer to a property of wrongness. An account of moral metaphysics is a material explanation of morality. We do not expect the virtuous agent to have such metaethical knowledge.

As for efficient explanations, asking in virtue of what does the good moral agent *know* moral truths might involve some sort of causal story describing the genesis of moral judgments. We can also ask for literal causal explanations. The crude divine command theorist holds that the efficient explanation for morality is that some god has made commands. A crude contractarian holds that the "cause" of morality is an agreement between agents. Hume thinks sympathy gives rise to much of morality. If efficient explanations of morality exist, we need not expect the virtuous person to know this.

³⁹ Aristotle, *Physics* (Cambridge, MA: Harvard University Press, 1986), II.3.

These explanations are not necessary to be virtuous. Such explanation is the realm of the moral theorist.

A moral theory might also be able to provide a final explanation of morality. Perhaps morality serves a purpose; morality is *for* something. Our crude divine command theorist might claim that morality exists in order to please God. Robert Nozick argues that morality has the evolutionary purpose of facilitating cooperation.⁴⁰ David Hume argues that justice exists to solve the problems of moderate scarcity and limited altruism. Friedrich Nietzsche averred that morality is developed by the weak to suppress the strong.⁴¹ Some theorists have held that moral language exists to prescribe and modify behavior.⁴²

However, the most important function of moral theory is the *formal* explanation of morality. In looking for a formal explanation, we seek out morality's underlying structure. The formal explanation of morality is morality's *essence*. An account of the essence of morality is an account of morality's core nature, of that in virtue of which a norm is moral. An essence is not an exhaustive list of the features of a moral action. To be essential is to be fundamental. A thing's essence describes not merely (nearly) necessary characteristics, but those characteristics upon which other important characteristics depend.

⁴⁰ Robert Nozick, *Invariances* (Cambridge, MA: Harvard University Press, 2002), chapter 5.

⁴¹ Friedrich Nietzsche, *On the Genealogy of Morals* (New York: Vintage, 1989 [1887]).

⁴² R. M. Hare, *The Language of Morals* (Oxford: Oxford University Press, 1952).

Kant argues that the Categorical Imperative is the essence of right action. Mill maintains that it is the Principle of Utility. Hume holds that the virtues are mental characteristics that serve those ends with which we have a sympathetic interest.⁴³ And so on. One or more of these philosophers may have provided a compelling and illuminating formal explanation of certain aspects of morality. However, Mr. Virtue's virtue alone does not reveal such explanations. If he wants to know these truths, he needs a moral theory.

Posner complains that academic moralists are not unusually skilled at making particular moral judgments such as “that act was wrong” or “this thing is good”. They are not expert performers or critics. In the Aristotelian framework we are entertaining, this is because moral theory uncovers the form of morality but not its substance. A perfectly virtuous person knows the substance of ethics, but not its formal essence.

High-level moral principles, like the Categorical Imperative, are attempts at explaining morality's essence. Such principles are ethical generalizations. Yet, not just any generalization is explanatory. Take these examples:

1. (As far as we know,) all morally right actions have been committed on Earth by humans in the past 400,000 years.
2. It is physically impossible to commit a morally wrong action in the center of the sun.
3. All right actions are right.

⁴³ Geoff Sayre-McCord, "Hume and the Bauhaus Theory of Ethics", *Midwest Studies in Philosophy* 20 (1996): 280-298, here p. 285.

These generalizations are true, but they explain nothing. Theorizing involves finding what right actions, virtuous character traits, and good things have in common, but as the examples show, not every such finding is explanatory. For a generalization to explain, it is not enough for it to be true, or even necessarily true. An explanatory generalization cannot just link some moral status *S* to some feature *F*; it must be that case that *S because F*. Some generalizations are imperfect, but are more explanatory than these three perfect generalizations. Kant argues that what wrong actions have in common is failing to respect people as autonomous agents. This generalization explains much, but it often seems silent (or incorrect) about moral obligations concerning nature, art, and the like.

This is not nearly a full treatment of the nature of explanation or the expertise of moral philosophers. Rather, it is an overview and somewhat formal description of their expertise. The substance of their expertise is shown ostensively: look at the content of actual moral theories. However, this is enough to show that moral theorists are genuine experts.

Most philosophical questions do not occur to most people. Moral theorists are not expert moral performers, coaches, or critics. However, in some respects, they are expert inventors. They invent problems and then solve them, often simultaneously. To say they invent the problems is not to say the problems are not real. Charles Goodyear invented vulcanized rubber, and thus vulcanized rubber became real. To invent something is to make it real. Posner complains that moral theorists concern themselves with abstract problems that lack political significance, e.g., with trying to defend

democracy though it is not under any real attack.⁴⁴ However, philosophers qua philosophers are not trying to save democracy in the real world. Rather, they are trying to see if it survives scrutiny. Philosophers model themselves on Socrates. If democracy is not being attacked, it is philosophers' job to attack it for the sake of seeing if it survives.

Some would claim that moral philosophers do not even have the expertise I describe. Unless a moral theorist is fully virtuous, her moral perception will be distorted, which will in turn infect her theorizing. There is wisdom in this objection, but it goes too far. Theorists attempt to answer special sorts of questions about morality. They need the skills to construct moral cosmologies, to provide deep explanations about certain moral phenomena. Physicists can do excellent work without having perfect empirical perception. Why would moral theory require its theorists to have perfect moral perception *in practice*? Philosophers need well-developed imagination and the ability to make accurate moral judgments when assessing carefully described, controlled thought experiments. They need to be sensitive to generalities and construct accounts of these generalities. This is compatible with the philosopher being of average character. Moreover, one reason moral philosophers can produce excellent theory despite having imperfect character is simply that in the cool hour of theorizing, vice offers fewer temptations.

Moral theory is a type of expertise about morality. Peter Singer suggests how this expertise is possible. First, producing theory requires one to be able to collect and assess

⁴⁴ Posner, *The Problematics of Moral and Legal Theory*, p. 59.

evidence. Moral theorists have more time to devote to these tasks. They are full-time thinkers. Second, philosophers are trained to recognize, formulate, and evaluate moral arguments. They are less likely to fall prey to fallacies and bad inferences. Philosophers are comfortable questioning assumptions and challenging norms. Third, philosophers are more accustomed to considering new possibilities and making important conceptual distinctions. They are thus less likely to avoid conflating disparate ideas.⁴⁵ To this list, we may add that on average moral philosophers are simply smarter than the average person. Producing theory is an intellectual exercise, and requires the intellectual virtues: insight, creativity, intellectual honesty, perseverance, caution, modesty, etc. Philosophers' training develops these intellectual virtues.

5. WHY THEORIZE?

Morality is a practical affair, concerning how we are to live. Posner objected that moral theory cannot convince us to act morally, nor can it aid us in determining what to do in each instance. If we grant Posner's objections, then theorizing about morality might be utterly impractical. (Whether we really should grant Posner's objections is another issue.) However, this does not entail that theorizing is worthless. Successful moral theories provide significant knowledge about morality.

If Posner's two objections hold, moral theory might be practically useless. However, useless knowledge is not worthless knowledge. Knowledge might be valuable

⁴⁵ Peter Singer, "Moral Experts", in Peter Singer, *Writings on an Ethical Life* (New York: Harper Collins, 2000), pp. 5-6.

for its own sake. Moral theory might properly be the realm of the curious and probing, of those who wonder about fundamentals. It may not be necessary for the hardheaded practical person.

Posner would demand a pragmatic justification of moral theory. It is this. Every normal adult has philosophical beliefs. I would not claim that they have *theories*, but they do possess less robust metaphysical, epistemological, and moral views. Some such people may never have reflected on their metaphysical beliefs, and they may have difficulty articulating the beliefs, but they still have them. Everyone has implicit views about the fundamental nature of reality, of the possibility and nature of knowledge, of issues of right and wrong. Human beings cannot help having such beliefs.

Normal human cognition results in everyone having some sort of philosophy. This fact gives us an answer to the first question, why philosophize? We do not have a choice in whether we will have a philosophy or not. But we can choose to philosophize. We can choose to reflect, to do our best to integrate our ideas into a coherent system, to weigh our philosophy against the evidence, and to formulate reasonable abstract principles. Or, we can choose to have our philosophies in an unphilosophical manner, without thinking.

The anti-philosophy skeptic advises us to forget "abstruse theoretical issues" and "focus on [our] practical needs".⁴⁶ Nicholas Rescher responds to these objections by arguing that philosophy is a form of "intellectual accommodation" into the world. A

⁴⁶ Nicholas Rescher, *The Strife of Systems* (Pittsburgh: University of Pittsburgh Press, 1985), p. 249.

good moral theory—even if it does not convince us to be moral or tell us what to do—might help us feel a little more at home in morality. Just as we need to physical comfort and security, we need intellectual comfort and security. The need for a good philosophy is like the need for a well-built house.

In closing, I wish to show that the view presented here passes a type of self-inspection. If what I have said about philosophical explanations is true, then such explanations will not allow one to derive all particular bits of knowledge. Applying this theory to itself, it would mean that my articulation and justification of moral theoretic knowledge, even if correct, would not suffice to derive the content of a moral theory. Just as reading a moral theory does not make one an expert moral performer, reading a theory of theoretical expertise does not make one an expert theoretician. My explanation of expertise fits its own model.

CHAPTER THREE

THEORY MULTIPLISM

1. A TOOL BELT FULL OF MORAL THEORIES

What is the nail for which moral theory is the hammer? This simple question assumes too much. We have not just nails to be hammered, but boards to cut, screws to tighten, and so on. We have no truly all-purpose tools. Likewise, I argue, we have no all-purpose moral theories. Some moral problems are nails; some are screws. Maybe Kantianism is a hammer and utilitarianism is a screwdriver. A well-equipped person will have, and sooner or later will need, both.

No one moral theory adequately addresses the full range of moral problems. Our best virtue, Kantian, utilitarian, contractualist, and pluralist theories are each well suited for explaining certain moral issues.¹ Yet, the best theories from each tradition are widely regarded as having defects, though we may disagree over the particular defects. This essay describes how to make the most of defective moral theories. Defective theories, defective at least when treated as “all-purpose” theories, are all we have, and probably all we’ll ever have. As with other tools, the rational response is to use each theory where it is most effective.

¹ Moral theories are normative explanations. They outline the general structure of ethics. They describe the essence of morality rather than its contours. Moral theories are less about telling us what things are good than about tell us why they are good.

I advocate *theory multiplism*. Theory multiplism holds that one can be justified in simultaneously accepting multiple incompatible theories about the same domain. Theory multiplism is not a moral theory. It's a methodological stance towards theories. We should not assume theory singularism, the doctrine that one should endorse exactly one theory. If no one theory is adequate for all problems, we need not commit ourselves to one theory.

Swiss army knives try to be all things; their virtue is versatility. However, they are not particularly good scissors, saws, screwdrivers, or corkscrews. If one has space for separate better tools, there's little reason to limit oneself to a Swiss army knife. Versatility is not the only virtue of tools or theories. One might worry that theory multiplism suggests having a different theory for each problem. This is not so. Robustness and generality are virtues of a theory, though not the only virtues. There is a point to having many tools rather than just a Swiss army knife. Eventually, though, adding more tools just adds clutter. The point of multiplism is to extract the maximum explanatory utility from theories. *Ceteris paribus*, adding theories, like adding tools, has diminishing marginal returns.

In the first few sections of this essay, I motivate moral theory multiplism by comparing it to *de facto* (and sometimes explicit) multiplism in the social and physical sciences. The remaining sections describe how multiplism would function (and sometimes has functioned) in moral theory.

2. THEORY MULTIPLISM BETWEEN THE SOCIAL SCIENCES

The best theories in political science, economics, sociology, history, and anthropology are not fully compatible. Each discipline has a core of problems where its theories function comparatively well. Economics is better at explaining markets while sociology is better for explaining relationships. To varying degrees, each discipline attempts to explain all human behavior. Economics is most notorious for this.² Economists not only try to explain markets, but also political behavior (e.g., via public choice economics), family interactions, the development of art,³ the behavior of sumo wrestlers,⁴ and more. Sociologists and anthropologists also tend to be as ambitious. E.g., the sociology of economics tries to explain market behavior through sociological theory.⁵

² E.g., Nobel Laureate Gary Becker argues that economics can explain all human activity in his *The Economic Approach to Human Behavior* (Chicago: University of Chicago Press, 1978). For more recent work, see Mariano Tommasi and Kathryn Ierulli (eds.), *The New Economics of Human Behavior* (New York: Cambridge University Press, 1995).

³ E.g., Tyler Cowen, *Good and Plenty: The Creative Successes of American Arts Funding* (Princeton: Princeton University Press, 2006); Tyler Cowen, *Creative Destruction* (Princeton: Princeton University Press, 2002); Tyler Cowen, *In Praise of Commercial Culture* (Cambridge, MA: Harvard University Press, 2000).

⁴ Steven D. Levitt and Mark Duggan, “Winning Isn’t Everything: Corruption in Sumo Wrestling”, *American Economic Review* 92:4 (2002), pp. 1594-1605.

⁵ E.g., see Jens Beckert, *Beyond the Market* (Princeton: Princeton University Press, 2002).

Max Weber tried to explain the productivity of capitalism by reference to a Protestant work ethic. Sociologists try to explain the development of philosophy, but their analyses contradict explanations offered by historians of ideas.⁶

Social scientists often complain that no factors are fully exogenous. They argue that X explains Y, but in some places, Y explains X instead. In a sense, social scientists respond to this problem by specializing. Anthropologists focus on culture as an explanation, sociologists on networks and interactions, economists on incentives and information, and psychologists on the relationship between mental states.

No field is fully successful on its own. At various points, the dominant theories in each discipline slip into epicycles and distortions. The further a theory extends beyond its core problems, the more it tends to rely on epicycles, ad hoc theses, and tautologies. E.g., economics models human beings as self-interested utility maximizers. When explaining market or even politician and bureaucrat behavior, this model works well. However, the economic model of human behavior is empty for explaining why mothers risk their lives to save their children. (“Saving children is part of their utility function.”)

Other disciplines suffer from similar defects. The best theories from each discipline are imperfect. Plus, the theories do not always cohere. Economics and sociology often explain the same phenomenon in contradictory ways. What should one do about such contradictions? One could eliminate the contradiction by ignoring all but one field. However, to understand humanity, one needs all of the social sciences.

⁶ Randall Collins, *The Sociology of Philosophies* (Cambridge, MA: Harvard University Press, 1998).

Economists and sociologists have different perspectives. One's own perspective is less distorted when one can see through both of their perspectives. Our social sciences are imperfect. To get the most from them, one must be a theory multiplist.

There often are major successes in interdisciplinary work or in attempts to combine two disciplines into a single new discipline. Many such attempts fail, producing explanatory mush. The continued specialization and division of the social sciences is no accident. Multiplism about disciplines typically is preferable to combining the disciplines.

3. MULTIPLISM WITHIN SINGLE DISCIPLINES

Multiplism also occurs within particular disciplines, not just between them. Whether multiplism will be permanent is contested. I will examine the cases of microeconomics versus macroeconomics and general relativity versus quantum mechanics.

Neoclassical microeconomics studies the behavior of individual economic actors, such as firms, consumers, or (in public choice theory) voters and politicians. It focuses on explaining the prices and quantities of specific goods and services. In contrast, macroeconomics looks at the economy as a whole. It examines aggregate and average figures, e.g., average prices, and employment levels.

Microeconomic and macroeconomic explanations often appear to contradict one another. Why are goods offered at one average price rather than another? The microeconomist claims that the average price is fully explained by individual buyer and

seller behavior in response to scarcity and diffuse knowledge. The macroeconomist explains the same price by producing an equation that correlates aggregate input prices and aggregate demand to the average price of the good. The two describe the same phenomenon, but they offer different explanations. The compatibility of the explanations is often unclear. Some economists think the two can be reconciled and thus use both types of theory.⁷ Some deny reconciliation is possible and only use one theory. However, the majority holds that both fields should be studied, despite their tensions, because both provide useful models and explanations. The default methodological approach is that when the best micro and macro theories conflict over a particular issue, the economist must choose the theory she believes is more successful there.

In physics, the two best theories are general relativity and quantum mechanics. However, the two are currently incompatible. They each describe the universe in radically different ways. In practice, physicists use both theories, but for different purposes. General relativity has implications for the behavior of sub-atomic particles, but quantum mechanics provides more successful models for that class of problems. Quantum mechanics has implications for the behavior of planets, but general relativity is more successful there. Many scientists have tried to unify the two within a single theory. So far, no one is regarded as successful. The de facto response is multiplism: physicists use each theory where it is strongest. In a recent interview, prominent theoretical

⁷ A classic reconciliation can be found in Roy E. Weintraub, *Microfoundations: The Compatibility of Microeconomics and Macroeconomics* (New York: Cambridge University Press, 1979).

physicist Stephen Hawking endorsed just such multiplism. Hawking once championed the view that science should aim for (and will find) fundamental unity and simplicity. He now believes that at best we can produce a patchwork of theories and models. These theories contradict each other. One theory Hawking endorses describes the universe as nine-dimensional everywhere; another says it is ten-dimensional everywhere. Our theories are essentially incomplete and imperfect, with limited applicability.⁸

In moral theory, as in physics and economics, we have produced a patchwork of theories and models. Moral theories may have implications for all moral phenomena, but their explanations are misleading or incorrect at various points. Using multiple theories is acceptable in the social and physical sciences, because it's the best we can do. To be fair, it should be acceptable in moral theory.

4. MULTIPLISM IN COMPUTER SCIENCE

In certain respects, a moral theory is to morality as programming languages are to machine code.

Computers store data in binary, i.e., in millions of switches (bits) that are either on or off. They recognize and read machine language, a set of functions written in binary telling them how to operate. Each computer type has a machine language unique to its particular architecture. Machine languages are implemented inside every computer. One

⁸ See Leonard Mlodinow, "Prometheus Bound", *Discover* 28:10 (October, 2005), 64-65.

See also Stephen Hawking, "Gödel and the End of Physics", Lecture: Dirac Centennial Celebration, Cambridge: University of Cambridge, 2002.

level up from machine language is assembly language. Assembly languages are like shorthand for machine languages. They are complicated, consisting of numerous functions and terms, but they are more accessible to human understanding, easier to remember, and so on.

At the highest level are programming languages, the languages in which programmers normally write. A compiling program translates the programming language into machine language. The translation from programming language to machine code is never perfect. Though the compiler might not find any bugs, in actual operation, the code created using the programming language will occasionally malfunction.

Different programming languages have different strengths and weaknesses. They are created for different purposes. Some languages are good for making video games but bad for running mainframes. Some are more powerful than others. Some are more flexible. Programming languages even have different ontologies. Some, like Python, have “objects”; some, like Ada, do not. Some programming languages make a rigid distinction between different types of data (e.g., integers and real numbers); others do not.

Multiple programming languages can imperfectly represent the same binary code. E.g., a skilled programmer could read the binary code in your computer and reverse engineer it into any number of programming languages: Fortran, Visual Basic, C++, Lisp. If two programmers were asked to do such reverse engineering, they would probably produce different codes written in different programming languages. Moreover, even if

the engineers were asked to use the *same language*, they would still produce different codes in that language. Programming languages can represent machine languages in potentially infinite ways. No such representation is perfect. Some are better than others. But no programming language dominates all others in all respects. There is no one best programming language.

Programming languages require discipline. If one does not use them as intended by their designer, they malfunction.

Here is the analogy. We might think of morality as being like a machine language. Morality is the set of functions, dispositions, capacities to discriminate and choose, etc., implemented in the mind (the hardware) of a virtuous person. This virtuous person can articulate some principles, rules of thumb, moral analogies, considered judgments, etc. These articulated principles are much like an assembly language. A moral theorist, like any other person, has her machine code (morality) imbedded in her, and has access to her assembly language. But her task is to reverse engineer the code into a programming language. In actual computer science practice, the main use for programming languages is to produce new programs in machine code. Yet, a programming language does more than that. It also *makes sense* of machine code. A programming language provides an understandable, explanatory linguistic structure for machine code. A programming language explains how a computer functions. It maps the otherwise bewildering code of ones and zeroes. Similarly, moral theories provide an understandable, explanatory linguistic structure for morality. They imperfectly map the bewildering complexity of the moral.

Different theories have different strengths and purposes, just like programming languages. COBOL is a bad language for producing video games but excellent for mainframes. A contractarian theory might provide an excellent account of why one should be moral, but a utilitarian theory might provide a better account of need. Most people do not know, and have no need for, a programming language. However, computer scientists need programming languages both for programming and for understanding computers. For most computer scientists, the optimal number of programming languages is not one. It is not infinity either. It is best to have multiple programming languages at one's disposal. Which languages one should use will vary between different programmers, though there will natural winners and losers. (Most programmers would not use Ada anymore, and most theorists would not pick ethical egoism.) Similarly, the optimal number of moral theories, for most moral theorists, need not be one.

In reverse engineering, machine code underdetermines which programming language one uses to represent it. Similarly, morality underdetermines which moral theory one uses to represent it. Even if there are moral facts, there is no unique truth about how a theory must represent those facts.⁹ (It is possible to make mistakes, though, by producing a misleading theory.) One theorist represents morality via a good version of deontology, another by a good version of utilitarianism, while a third uses both. One programmer, faced with reverse engineering, represents machine code via Perl, another

⁹ Cf. David Schmidtz, "When Justice Matters", *Ethics* (forthcoming).

via Java, and a third uses both. The third person, in each case, is a theory multiplist. In computer science, unlike moral theory, multiplism is normal.

A programming language malfunctions when not used as intended. A poorly designed language always malfunctions, i.e., no possible compiler could translate the programming language into machine code. Similarly, a moral theory malfunctions when not used as intended. In philosophical disputes, philosophers deliberately induce malfunctions in rival theories. They search for counterexamples and find them. Programming languages, when translated into machine code, always have mistakes in operation—so do moral theories.

In principle, one could dispense with programming languages and understand computers simply via machine code. Similarly, one could dispense with moral theories. Yet, one should not do either.

Machine languages are unique to each computer. In some respects, a person's moral psychology is unique to him. Programming languages, however, are portable; they are universal and can be applied to multiple computers via compilers. Similarly, moral theories are universal and apply to multiple people.

Programming languages recognize different classes of objects, make different distinctions, and function differently, even though they represent the same code. In some respects, C++ and Perl contradict each other, but in other respects, they say the same thing. Similarly, moral theories recognize different classes of objects, make different distinctions, and function differently, even when they represent the same morality. In

some respects, neo-Kantianism and rule consequentialism contradict each other, but in other respects, they say the same thing.

A set of instructions written in C++ cannot directly translate into computer behavior. Implementing code requires a compiler. In addition, one needs good hardware, or even great code will not run. Analogously, a set of substantive principles written in the language of deontology cannot directly translate into moral behavior. Moral education and good judgment are needed intermediaries. They are the compilers of moral theory. Plus, one needs good hardware. Morality does not run on a sociopath any more than Windows XP runs on an IBM 8086.

5. DIFFERENT THEORIES CAN SOLVE DIFFERENT MORAL PROBLEMS

One model of applied ethics is this. A Kantian inquiring about duties to preserve the environment asks what the Categorical Imperative implies, whereas a utilitarian asks what the Principle of Utility implies. Many times, this is what theorists do.

Other times, a theorist recognizes that her favored theory cannot adequately deal with a question. In such cases, many theorists turn to different theories. Thomas E. Hill, Jr. is a Kantian moral theorist, but for him virtue ethics provides a better explanation for why one should respect the environment.¹⁰ Perhaps (with epicycles) everything Hill says can be integrated into Kantian theory, but instructively, Hill does not attempt this

¹⁰ Thomas E. Hill, Jr., "Ideals of Human Excellence and Preserving Natural Environments", *Environmental Ethics* 5 (1983): 211-224.

integration. Bryan Norton realizes that neither rights-based theories nor cost-benefit analysis utilitarianism can adequately explain why a child ought not strip-mine sand dollars from a beach.¹¹ He needed a different type of explanation. Rights-based theories and utilitarianism have implications for the sand dollar case, but what they have to say is not illuminating. The theory that best explains the sand dollar case might have limited application elsewhere. Norton favors a deontological theory overall, but as a pragmatist, also recognizes that such theories are not particularly persuasive for influencing policy-makers. Lynn Paine's "Managing for Organizational Integrity" is theoretical hodgepodge, using consequentialist analyses for some issues, virtue ethics for others, and straight casuistry elsewhere.¹²

Another robust example comes from Rawls. In *Political Liberalism*, Rawls describes his theory of justice as an attempt to articulate an overlapping consensus implicit in modern democratic societies. On Rawls' view, individual citizens may have comprehensive moral worldviews. Some comprehensive views are religious; some are moral theories. A comprehensive moral theory is general and universal, applying to all

¹¹ Bryan G. Norton, *Toward Unity among Environmentalists* (Oxford: Oxford University Press, 1991), pp. 3-13. A rights-based theory says the child can't take any sand dollars. Cost-benefit analysis licenses her to take anything within the sustainable yield. Norton knows extra-theoretically that the girl may take a few sand dollars. He needs a theory to explain why.

¹² Lynn S. Paine, "Managing for Organizational Integrity", *Harvard Business Review* 72:2 (1994), 106-117.

moral aspects of a human life.¹³ In contrast, Rawls' theory of justice is meant not to be a comprehensive doctrine. It is meant to solve one and only one problem—justice in the basic structure of society.

On Rawls' view, citizens of just societies will recognize that reasonable disagreement and moral pluralism are permanent features of modern democracy. Since citizens value reciprocity and respect one another, they effectively adopt two moralities. For issues of justice, they employ a “political liberalism” (Rawls' two principles of justice). In their personal lives, they employ their comprehensive doctrines. Their comprehensive doctrines have implications for the justice of a basic structure, and sometimes these implications contradict the principles of political liberalism. (For instance, one's comprehensive doctrine might imply that abortion should be illegal, whereas the principles of political liberalism do not.)

In effect, Rawls recommends that citizens be theory multiplists, accepting both their comprehensive moral doctrines and a freestanding political theory, though the two can conflict. One accepts one's comprehensive doctrine for the purpose of living a good life and knowing the truth about morality. One accepts the freestanding principles of justice for the purpose of allowing people to live together prosperously despite disagreement. Generally, one keeps one's comprehensive doctrine isolated from the public arena.

As a committed eudaimonist asked to write about issue X, I may struggle to wring a good explanation from my theory. This may allow me to ignore the imperfections of

¹³ John Rawls, *Political Liberalism* (New York: Columbia University Press, 1996), p. 13.

my theory, but just as often it results in ad hoc theorizing. However, suppose a student comes to me for personal advice on X. That is not the time to be ideological; it's a time to be honest. I recognize that to help the student, I should explain the problem in utilitarian terms, even though I wish eudaimonism provided a better explanation.

Theories have different strengths and weaknesses.

Why theories have different strengths and weaknesses is an interesting question. Here is a first pass. Many research traditions exist in moral philosophy. Each tradition has multiple theories. E.g., the research tradition of Kantianism includes the distinct theories of Onora O'Neill, Thomas Hill, Christine Korsgaard, and Kant himself. Traditions as types have strengths and weaknesses, as do token theories. Providing an exhaustive list, explanation, and defense of these strengths and weaknesses would be exhausting. For illustrative purposes, I will quickly evaluate each tradition.

Contractualist theories are good for explaining the congruence between morality and self-interest and for illustrating how morality solves coordination problems.

Consequentialist theories are best at showing ways to mediate between commensurate goods and at dealing with cases where numbers matter. *Kantian* theories are good at explaining intuitions about the special dignity of agency and duty, the nature of practical reason from the first person perspective, the universal aspects of morality, and the public nature of reasons. *Pluralist* theories provide the simplest account of moral conflicts, for the apparent fragmentation of value, for the specialization of moral domains, and for deliberation and moral phenomenology. *Eudaimonist virtue ethical* theories are best for explaining the role of virtue in a good life, for describing what makes life meaningful, for

descriptions of virtue from a third person perspective, and for understanding moral education and development. *Feminist* moral theories are best for accounting for situated perspectives, for tempering universal principles by elaborating the significance of partial ties, and for highlighting the importance of the personal. This list is not complete. Your list may be different. Yet, most of us agree that no research tradition trumps all others for these important issues.

Perhaps unavoidably, theories emphasize different features of morality. Contractualist theories emphasize *reciprocity*. Deontological theories emphasize *respect*. Consequentialist theories emphasize *beneficence*. Virtue theories emphasize *wisdom*. Pluralist theories emphasize *moral complexity*. Divine command theories emphasize *reverence*. Ethical egoistic theories emphasize *prudence*. And so on. Theorists cannot help emphasizing certain aspects of morality over others, much like storytellers must emphasize some details.

Theories have different starting points. Aristotle begins the *Nicomachean Ethics* by asking if all of our actions have a unitary aim. Mill starts *Utilitarianism* by asking what the highest good is. Contractarian theories often begin questions about the rationality of morality. Deontological and Rossian theories tend to begin with our intuitive ideas of duty. Each of these theories eventually turns to the issues other theories began with. However, a theory constructed by first focusing on one issue may be inadequate when applied to another issue.

Many philosophers agree that theories are tools—maps, in particular.¹⁴ Different maps have different purposes and thus represent the same terrain differently. (Compare geographical maps to street maps.) Some maps purposely distort and exaggerate certain features. If a friend asks me to draw a map from his house to mine, I'll omit most streets and draw my house as a giant box. Taken literally, my map says my house is the size of city block. The hand-drawn map looks little like a Rand-McNally street map of the same area. Neither map looks like the actual city, except in abstraction. If theories are maps, then different problems may call for different maps. For instance, David Schmitz suggests, "a theory that maps a public official's duties may be quite different from a theory that maps a parent's duties."¹⁵ If so, then we will need two different theories. The two theories may conflict, as different maps often do.

6. DIFFERENT THEORIES CAN SERVE DIFFERENT THEORETICAL FUNCTIONS

Moral theories are not all one thing. Moral theorists do not ask all the same questions. They begin and end in different places, and they theorize for different

¹⁴ For a recent use of the theories as maps metaphor applied to scientific theorizing, see Ronald N. Giere, *Science without Laws* (Chicago: University of Chicago Press, 1999), pp. 25-26, 81-82, 214-215, et passim. For a recent use of the theories as maps metaphor applied to moral theorizing, see David Schmitz, *Elements of Justice* (New York: Cambridge University Press, 2006), pp. 21-28, 227.

¹⁵ Schmitz, *Elements of Justice*, p. 23. I owe much of this thinking about theories as maps to Schmitz.

purposes. Individual moral theorists do not ask the same questions each time they theorize. They adopt multiple perspectives, and theorize for different reasons.

However, there is no finite, known, unchanging set of problems for which moral theory is a solution. Since Plato, moral philosophy has possessed a common core of problems, but there is rich variation along the periphery. The philosophical problems of tomorrow may be unlike today's. Accordingly, our current tools might obsolesce.

Moral theorists theorize for different reasons. Accordingly, moral theories have different functions. Below, I provide a (non-exhaustive) list of such functions. I present them as ideal types, though perhaps no real theory belongs merely to one category. (Similarly, in economics, "entrepreneur", "laborer", "capitalist", "landowner" and other functional categories are ideal types, but any real person is a token of two or more.)

Speculative theories explore the limits of moral thought and understanding. They probe into darkness with hope of discovery. They expand our moral imagination and transform our moral experiences by providing strange but illuminating perspectives. Much of Robert Nozick's work is avowedly speculative, e.g., chapter 10 of *Anarchy, State, and Utopia*, or nearly all of *The Examined Life*.¹⁶ Honestly declaring his limitations, Nozick investigated areas most analytic philosophers avoid (e.g., why there is something rather than nothing).

Calculative theories are meant to assist in making moral decisions. A more ambitious (or injudicious) theorist might attempt to construct a decision algorithm that specifies the particular actions required in any situation. Many mid-twentieth century act

¹⁶ Robert Nozick, *The Examined Life* (New York: Simon and Schuster, 1989).

utilitarian theories fall in this category. Less ambitious, but more feasible, are attempts to specify a procedure for deriving a few principles for a limited context. John Rawls' "Outline of Decision Procedure for Ethics" is an example.¹⁷ Least ambitious, but more fruitful, are theories that elaborate deliberative strategies, portraying how good people (competent moral judges) have made past decisions, which a decent person could emulate with some success.¹⁸

Modeling theories impart understanding of complex moral issues through simplified idealizations that highlight morally significant features. Models can be useful when differences between reality and the idealization do not matter. (Importantly, the assumptions of theoretical models rarely hold.) Models are common in economics, physics, chemistry, etc. The same thing can be modeled differently for different problems. For one question, an economist might model General Motors as a firm in perfect competition, for another question as a firm in oligopolistic competition.¹⁹ Modeling is perhaps omnipresent in moral theory. Kant's three versions of the Categorical Imperative provide models. The first formulation models wrong action as

¹⁷ John Rawls, "Outline of a Decision Procedure in Ethics", *Philosophical Review* 60:2 (1951), 177-197.

¹⁸ See Robert Louden, *Morality and Moral Theory* (Oxford: Oxford University Press, 1992), pp. 132-135.

¹⁹ See Milton Friedman, "The Methodology of Positive Economics", in Daniel M. Hausman (ed.), *Philosophy of Economics* (Cambridge: Cambridge University Press, 1984), pp. 210-244, here p. 234.

contradictory or unwillable universal laws, the second formulation models wrong action as treating people like hammers, while the third formulation models right actions as conforming to laws of an ideal kingdom.²⁰ We might also describe paradigm cases as models. The best model of the moral significance of need is of parents attending to young children. In political theory, Rawls discusses five idealized economic regimes as a means to evaluating real regimes.²¹ The danger with models in any field is the models may abstract from real world details that (often uncharacteristically) are pivotal in the particular case. For example, Rawls misguidedly assumes that if a regime does not have institutions dedicated actively to pursuing certain "political values", it does not obtain them.²²

Caricaturing theories provide a special type of model.²³ Caricature models purposefully isolate and exaggerate some aspect of reality. The models are useful not because the theorist believes reality is relevantly like the model, but because she wants to

²⁰ Immanuel Kant, *The Groundwork of the Metaphysics of Morals*, in Mary J. Gregor (ed., trans.), *Practical Philosophy* (Cambridge: Cambridge University Press, 1996 [1785]), 4:421-440.

²¹ John Rawls, *Justice as Fairness: a Restatement* (Cambridge, MA: Harvard University Press, 2001), pp. 135-140.

²² Rawls, *Justice as Fairness*, p. 137.

²³ Here I closely follow Alan Gibbard and Hal R. Varian, "Economic Models", *The Journal of Philosophy* 75:11 (1978), 664-677.

determine what would happen were the model real. E.g., economist M. Spence wanted to know if it is *possible* that college degrees elicit salary premiums merely because they signal perseverance to employers, not because colleges provide job skills. To determine this, Spence constructed a caricature model: a world in which everyone knows that college degrees do not improve employee productivity.²⁴ He found that the perseverance signal alone would make possible the salary premium, though this does not establish the signal is the cause of real premiums. The use of caricatures is also important in moral theory. David Gauthier's *Morals by Agreement* demonstrates that *homo economicus*, who lacks tuistic preferences, has reason to develop dispositions to cooperate.²⁵ Gauthier does not pretend that actual human beings are relevantly like *homo economicus*. Rather, he wants to show that cooperative morality is rational even for beings that don't care about others. Presumably, this means morality is even more rational for real people who do care. Complaining that Gauthier's model is unrealistic misses Gauthier's point.

Articulative theories attempt to explicate assumptions, sharpen vagueness, elaborate principles, and fill in gaps in moral thinking. They make tacit knowledge explicit. Nearly every moral theory involves articulation. Bernard Gert's *Morality: Its Nature and Justification* is largely an articulative theory meant to describe the commonsense morality most people accept.²⁶

²⁴ Gibbard and Varian, "Economic Models", p. 674n.

²⁵ David Gauthier, *Morals by Agreement* (Oxford: Oxford University Press, 1986).

²⁶ See Bernard Gert, *Morality: Its Nature and Justification* (Oxford: Oxford University Press, 1998), p. ix. Thanks to Bernard Gert for helpful discussions of his project.

Persuasive/rhetorical theories attempt to modify people's behavior. In "Famine, Affluence, and Morality", Singer seeks to induce people actually to contribute to famine relief.²⁷ In principle, a theory could be purely rhetorical. A theorist could tell a "noble lie". One might advance a theory simply to manipulate others.

Justificatory theories defend moral views. (Persuasion and justification are not the same, though they are related.) Most theories are justificatory. They give reasons for substantive moral positions. Gert's *Morality* is not merely intended to articulate, but also to justify commonsense morality. Sidgwick's *Methods of Ethics* does not merely articulate utilitarianism, but attempts to demonstrate that utilitarian principles are correct.

These are some of many purposes theories can serve. It would be nice if one theory adequately served all of these different purposes. However, if no singular theory adequately serves all of our goals, we have a choice. We can abandon some goals out of theoretical loyalty. Alternatively, we can pursue different goals via different theories.

7. DIFFERENT EXPLANATIONS ANSWER TO DIFFERENT PERSPECTIVES

At first glance, explanation is a two-part relationship. Explanations take the form "*E* explains *P*". Quick reflection reveals this will not do. Any teacher knows that to explain is not just to state truths. (As a student, I had teachers who stated plenty of truths but never successfully explained anything.) There are no explanations as such.

Explanations are always for some audience, with certain background needs, assumptions,

²⁷ Peter Singer, "Famine, Affluence, and Morality", *Philosophy and Public Affairs* 1 (1972) 229-243.

abilities, and conceptual schemes. Moreover, the types of explanations offered by moral theorists, scientists, and the like are always a mix of linguistic and mathematical representations. Explanations are at least a four-part relationship of the form "*E* explains *P* to some subject *S* in terms of *T*". For the phenomenon of falling apples, what constitutes a good explanation will be different for a 5-year-old, a high school student, an intelligent layperson, and a graduate student in physics. We should not say that only professional cosmologists have good explanations while all others have pseudoexplanations. That stance would probably backfire. A hundred years from now, when our science changes, we would have to say today's professional cosmologists have pseudoexplanations. Looking backward, I would hate to claim that Isaac Newton, Niels Bohr and Albert Einstein never had good explanations.

Ancient and modern moral theories commonly tried to explain why we should be moral. Gregory Kavka argues that answers to this question have four dimensions. We have to be concerned with 1) the audience of our argument, 2) the people whom we argue should be moral, 3) the conditions in which it is rational to be moral, and 4) those objects of moral assessment (individual acts, character traits, lifestyles) are we trying to justify.²⁸ I would add a fifth dimension: perspective. From what does one want to justify morality, from the perspective of self-interest, natural telos, inherent nature of practical reason,

²⁸ Gregory Kavka, "The Reconciliation Project", in David Copp and David Zimmerman (eds.), *Morality, Reason, and Truth* (Totowa: Rowman and Allenheld, 1984), p. 297-219. See also David Schmitz, *Rational Choice and Moral Agency* (Princeton: Princeton University Press, 1995), chapter 10, especially p. 243.

etc.? Here, explanation is a four or five-part relationship. I suspect that the number of dimensions of explanation varies across the universe of philosophical problems.

Many moral theories are normative explanations of normative phenomena: they say why we ought to do what we ought. Given the variety that can be exhibited along the dimensions of explanation, we should not think any one theory would provide the best explanation for each variation. Changes in audience alone may privilege some theories over others. I would not burden a five-year-old with Kantianism, though Kantianism is an excellent for professional philosophers. However, even among philosophers, there are different needs, background assumptions, etc. It is optimistic—perhaps oppressive—to think that only one explanation will do. It can be oppressive when reasonable theoretical disagreements lead to conflict and resentment.

Sidgwick complained that no theory or philosophical explanation would be fully acceptable to a layperson, because the layperson holds multiple contradictory perspectives on the world. When a “why” question is satisfactorily answered, the layperson switches perspectives and demands a different answer.²⁹ Sidgwick says so much the worse for the layperson. I have the opposite reaction. It is reasonable, normal, and probably inescapable for a person to adopt different perspectives. Responding to human perspectives is philosophy’s task. If no single explanation serves all reasonable perspectives, then philosophy must prepare multiple explanations. This motivates theory multiplism. Theory multiplism accepts that individuals adopt different perspectives in different contexts, allowing that different theories respond to these different perspectives.

²⁹ Henry Sidgwick, *The Methods of Ethics* (Indianapolis: Hackett, [1907] 1981), p. 6.

A change in perspective often occasions a change in values. One way we assess theories is by how well they respect certain cognitive values. Here is a partial list of the cognitive values of moral theory: a) fit with our moral intuitions and considered judgments, b) internal consistency, c) fit with theories outside moral theory, such as economics, physics, and the rest of philosophy, d) ability to be impartially justified, e) ability to be justified from a situated point of view, f) success in explaining new cases, g) rigor and robustness, h) theoretical parsimony, i) generality, j) fecundity in suggesting new philosophical investigation, k) truth or truthlikeness.³⁰ Moral theories will always be subject to assessment by multiple criteria. These criteria cannot be reduced to one master criterion. We are committed to a *plurality of values of assessment*. The criteria can and do conflict. There is no obvious lexical ranking of the criteria. (Even truth does not dominate. A partially false theory that can be understood by smart people is better than a completely true theory so complicated no one can understand it.) The lesson here is that we are asking *for many different things* from our theories.

When endorsing a theory, we effectively take a stand on the comparative importance of the values by which we assess a theory. There is reasonable disagreement on the comparative weight of those values, leading to reasonable disagreement about theories. Perhaps John espouses Kantianism while Erin accepts pluralism simply because they dispute the relative importance of conflict mediation. John is bothered that Rossian pluralism has few resources to distinguish resolvable from irresolvable conflicts. Erin

³⁰ Some of these desiderata are taken from Brad Hooker, *Ideal Code, Real World* (Oxford: Oxford University Press, 2000), p. 4.

worries that while Kantianism has more resources to mediate conflicting duties, it falsely identifies too many cases as resolvable. John prefers imperfect guidance to no guidance; Erin prefers no guidance to imperfect guidance. Both preferences are reasonable.

Disagreements about moral theory need not be disagreements about moral values. Many disagreements over moral theory result from disagreements about the axiology of theories themselves—not of moral values, but the *values by which we judge theories*. A commitment to the irreducibility, partial incommensurability, and plurality of theoretical values is in tension with a strong commitment to the view that just one moral theory (pluralist or monist) merits acceptance. Even if two people agree on the moral truths, even if they agree on all possible moral judgments, reasonable disagreement over the importance of theoretical values means that they need not agree on the best theory of those moral truths.

People assign different weights to cognitive values. This can legitimate their accepting different theories. However, a single person could legitimately assign different weights at different times. As such, it could be reasonable to employ different theories at different times. Some contexts call for simplicity, some for impartiality, and some for rigor. Different theories might be called for in different contexts.

8. WHEN THEORIES DISAGREE

Sometimes theories disagree about what is right or good. Insofar as moral theories "tell us what to do", i.e., are practical guides to right action, theory multiplism

will sometimes mean accepting theories that issue contradictory instructions. Does this mean theory multiplism is incoherent?

Even if theories “tell us what to do”, we need not always listen. Using a theory as a guide requires good judgment and interpretation. Any theory simple enough to be useful will also be misleading.³¹

A typical doctoral student has three or more professors on her dissertation committee. Often one is the official "supervisor" or "advisor", but students often have multiple advisors. Committee members offer writing, research, and career advice. Quite often, they issue contradictory advice. The student must exercise good judgment and interpretation regardless of whether her advisors agree. She must think for herself. When her professors disagree, this shows she should pay attention, for the issue in question is interesting. Normally, it is better to listen to multiple committee members, despite their disagreements, rather than ignore all but one for consistency's sake.

A theory multiplist with contradictory moral theories is much like a graduate student with disagreeing advisors.³² Each theory offers a useful, wise, but incomplete and imperfect perspective on morality. It is often better to have access to these different perspectives than not. Disagreements between the imperatives issued by different theories show the multiplist that she should pay attention. Ultimately, she must exercise good judgment in choosing which theory (or advisor) to listen to in a particular case.

³¹ Cf. Schmidtz, "When Justice Matters".

³² The metaphor of theories as advisors was suggested by Jason Matteson, "Aspiring to Advice", University of Arizona, Ph.D. dissertation (in progress).

Theories offer not just contradictory advice, but also contradictory explanations. This may also seem to make theory multiplism incoherent. However, advisors offer contradictory explanations. The problem of contradictory explanations is the same as the problem of contradictory advice.

If I am right, multiplism already exists in the social sciences, physical sciences, and elsewhere. Given that we use contradictory theories, maps, and models, whether we can properly regard them as true is a difficult question, deserving a paper of its own.³³ However, the key to answering the problem seems to be this. If we are stuck using multiple, imperfect, and often contradictory representations of reality, rather than accusing ourselves of incoherence, we need to reconceive what it means for representations to be true. On my view, good representations (theories, maps, models, etc.) illuminate the truth about certain aspects of the represented subject matter, but being imperfect, also distort and misrepresent other aspects. Truth is spread out over multiple representations and mixed with falsehoods.

9. MULTIPLISM IS NOT REDUCIBLE TO A PLURALIST MORAL THEORY

Theory multiplism is *pluralism about moral theories*. It is not a pluralistic moral theory, such as W. D. Ross'.³⁴ Pluralistic theories hold that there is a multiplicity of basic moral duties and/or values, and these duties and values cannot be subsumed beneath one principle. They claim that morality is not all one thing. Pluralistic theories are theories

³³ For a start, see my "Moral Theories and the Illumination of Truth".

³⁴ W. D. Ross, *The Right and the Good* (Indianapolis: Hackett, [1930] 1988).

about morality. Theory multiplism is a metatheory, a theory *about moral theories*: more than one moral theory merits acceptance. A theory multiplist might accept both Rossian pluralism and utilitarian monism simultaneously. Moreover, a theory multiplist might accept more than one pluralistic theory. E.g., the pluralistic moral theories of Charles Larmore or Michael Walzer are not identical to Ross' theory.³⁵ However, the multiplist might employ each for different purposes. A theory multiplist could even reject all pluralistic theories, finding only monistic theories philosophically useful.

Theory multiplism and pluralist moral theory address different concerns. The views don't compete. One might object that everything I say in defense of theory multiplism can be accommodated by a pluralistic moral theory. E.g., I say that Kantianism is good for dealing with some philosophical problems but utilitarianism is good for others. A committed Rossian pluralist might respond that I am merely advocating a version of pluralism in which there are both utilitarian and Kantian principles. This is not my view. Thus, I need to elaborate the differences between accepting multiple moral theories as opposed to accepting one moral theory with multiple principles.

Deciding to accept a pluralist theory requires another choice—which pluralist theory? This is one reason why theory multiplism cannot be reduced to a pluralist moral theory. Walzer's pluralist theory of justice cuts the valuational universe into separate

³⁵ Charles E. Larmore, *Patterns of Moral Complexity* (New York: Columbia University Press, 1987), chapter 6; Michael Walzer, *Spheres of Justice* (New York: Basic Books, 1983).

spheres of goods, with each good having different principles. David Miller accepts many of Walzer's substantive moral judgments. However, his theory begins with models of human relationships and evaluates institutions by their approximation of the appropriate model.³⁶ David Schmidtz portrays the valuational universe as constituted by more or less separate classes of everyday experiential contexts.³⁷ These theories are not equally successful at dealing with different problems. Perhaps Walzer better articulates the nature of social, collective meaning, but Schmidtz has a better theory of the comparative significance of such meanings.

Committing to a pluralist moral theory involves choosing a manner of dividing the moral landscape. Shall we portray morality as the interplay of seven prima facie duties, as Ross does, of classes of obligations, as Larmore does, or of spheres to be kept entirely separate, like Walzer? One cannot just amalgamate the theories into one giant pluralist theory, or so I will argue.

Pluralist theories hold there are multiple basic sources of reasons. What is right or good in any instance is the outcome of the combined interplay of these reasons. Theory pluralism says that are multiple good theories. Different theories offer different explanations for what is right and good. For some problems, it is best to represent morality via one type of theory, for other problems, via another type of theory. This is not the same as saying that there are multiple basic sources of reasons. Pluralist theories

³⁶ David Miller, *Principles of Social Justice* (Cambridge, MA: Harvard University Press, 1999).

³⁷ Schmidtz, *Elements of Justice*, pp. 4, 17-18.

answer the question what should I do and why by describing competing reasons that are assumed always to be prima facie relevant. Multiplism answers the question what theory best explains this problem by comparing the explanations offered by different theories. It doesn't assume that any reasons are always prima facie relevant.

The leading theories in history, economics, anthropology, psychology, and sociology do not combine into one master pluralist theory. They are separate theories, not one big theory with many explanations and models. The leading theories in physics—general relativity and quantum mechanics—do not (as of yet) combine into one master pluralist theory. General relativity and quantum mechanics are separate theories, not one big theory with many laws. Whether macroeconomics and microeconomics are subsumed beneath one pluralist theory is contested. They appear to be separate theories with different types of explanations. Separate programming languages cannot be combined into one language. They have different grammars. Different maps of the same terrain often cannot be combined into one map. They draw the world differently. One can't create a single map that draws the world two different ways at once. Outside moral theory, multiplism does not reduce to one pluralist theory. Inside moral theory, it does not either.

No moral theory is clearly the best. We cannot be sure whether morality really is monistic, pluralistic, or particularistic. From different perspectives, the moral universe looks different. From some perspectives, it seems perfectly ordered, from others chaotic. From some perspectives, morality seems to be all one thing, but not from others. One reason that multiplism does not reduce to pluralism is that multiplism does not take a

stand on whether morality is all one thing. Pluralism does; it says morality is not all one thing.

Multiplism says that different theories have different insights and so we should accept multiple theories. The pluralist objects, saying that we could just combine the principles of each theory to produce one theory with all the same insights. However, there's no guarantee that combining theories into one results in combining insights. In the case of the social sciences, multiplism about disciplines is often preferable to combining the disciplines. Attempts to combine two disciplines into a single new discipline fail—producing explanatory mush—more often than they succeed. The continual specialization and division of the social sciences is no accident. No “theory of everything” has been found in the social sciences, so social scientists use different theories for different problems, even though the theories contradict each other. Analogously, the reason that pluralist theories are just another class of theories in competition with all the rest is that pluralist theories *don't combine* the insights of monist theories. Rather, they have their own insights and their own failings. Multiplism is a different kind of pluralism from Ross'.

10. WHAT MULTIPLISM SAYS ABOUT THEORY CONSTRUCTION

Theory multiplism is a metatheoretical doctrine concerning how to think about existing theories. It is not a doctrine about how to construct moral theories. This distinction allows multiplism to avoid an objection.

This objection is this: Since theory multiplism suggests that we use different theories for different problems, it thereby recommends that we *build* different theories for different problems. That approach is not progressive. Theory multiplism opposes seeking one unified moral theory that covers all problems. This is akin to assuming that every scientific problem has a special science. In practice, that assumption would stifle the advancement of science. Thus, theory multiplism would stifle the advancement of moral theory. The best approach is to assume (defeasibly) that any given discipline is unified.³⁸

Actually, I agree that seeking unity is the best approach. I recommend that moral theorists continue to do what they've been doing all along: looking for the one best theory, improving existing theories in light of imperfections, and so on. The reason for recommending this is that it is progressive. If we start by assuming that the whole of morality cannot be represented by one moral theory, we will never find unity, even if unity is there. In both the sciences and in moral theory, it is progressive to seek a "theory of everything".

Multiplism reacts to the fact that we do not actually have a theory of everything. We have a number of good but imperfect theories that illuminate various truths about their subject matter while distorting other truths. Thus, multiplism recommends accepting multiple theories as a means to extracting the most truth from the theories. When we confront a problem that our favored theory does not adequately solve, and

³⁸ E.g., see Mario Bunge, *The Methodological Unity of Science* (Cambridge, MA: Springer, 1973).

attempts to modify the theory to solve that problem fail, we may use a different theory for that problem. Philosophers find multiplism distasteful because we still are still burdened with an Enlightenment model of rationality. The sciences regularly accept multiplism because they have dropped that burden.

Theory multiplism is a progressive response to the changing state of moral theorizing and the problems confronting it. It is progressive in that it involves adapting our best theories to solve problems. If ever a theory emerges as a clear winner as measured by all theoretical values (from all perspectives, etc.), theory multiplism recommends adopting that theory. Yet, that has never been the case, and we have no reason to expect it ever to be the case. However, our theories typically are improved by trying to produce a dominant theory. Accordingly, a theory multiplist can endorse the methodological principle that one should seek to produce a single best moral theory for all problems.

CHAPTER FOUR

MORAL THEORIES AND THE ILLUMINATION OF TRUTH

1. THEORY AND METAPHOR

Romeo, upon discovering Juliet dead (or so he thinks), first describes Death as an advancing army that has conquered Juliet. Then, only a few lines later, he says that Death is Juliet's paramour.¹

Suppose a literature student complains that Romeo contradicts himself. Surely, the student says, Romeo is mistaken. Death cannot be both an advancing army and Juliet's paramour. The student is not wrong. In some clear sense, Romeo does contradict himself. On the other hand, there are no grounds for complaint. Very well, Romeo contradicts himself. He is large; he contains multitudes. Shakespeare wants us to understand something through these contradictory metaphors. Had Shakespeare been worried about eliminating literal contradictions, he would have used only one metaphor. Fortunately, he was not worried. The two metaphors may contradict each other, but they illuminate or point to the same basic truths or to different but compatible aspects of the nature of Juliet's death. We come to understand more through the metaphors in conjunction than we would if exposed only to one.

¹ William Shakespeare, *Romeo and Juliet*, in Stanley Wells and Gary Taylor (eds.), *The Complete Works of Shakespeare* (New York: Oxford University Press, 1988) V.3.94-96, and V.3.105.

When I look at moral philosophy, here is what I see. A plurality of moral theories abound. There is significant dissensus. Surely, the point of moral theorizing is to track the truth about morality in some way. However, it is not clear which, if any, of the theories are true. Despite the disagreement among philosophers and the contradictions between their theories, one understands more about morality through *multiple* theories than from just one. If a student asks me which one theory she should read to understand the truth about morality, I tell her to read more than one.

Good moral theories reveal something true about morality. They illuminate a truth. This does not commit us to saying that the theories themselves are true. Similarly, good metaphors can show us something true, but that does not commit us to saying that metaphors are true. It can be illuminating to use contradictory theories and metaphors. The point of this paper is to advocate a certain type of realism about moral theories. I will argue that there can be better and worse theories, even if there is no uniquely best theory, no theory that can be called "the right answer". In my view, truth is not so much *in* the theories. Truth is what a good user *gets from* a good theory.

2. THE MANY MORAL REALISMS

Moral realism is not all one thing. As Geoff Sayre-McCord explains, there are many different ways of being a moral realist.² I would add that there are many different

² Geoff Sayre-McCord, "Introduction: The Many Moral Realisms", in Geoff Sayre-McCord (ed.), *Essays on Moral Realism* (Ithaca: Cornell University Press, 1988), pp. 1-26. According to Sayre-McCord, objectivism, subjectivism, and intersubjectivism can all

things about which to be a moral realist. One can be a realist about particular moral judgments, general moral principles, and/or moral theories. Each of these realisms has a degree of autonomy; in most cases, one can be a realist about one and an antirealist about another. In this essay, I want to advocate a certain type of *theory realism*, while remaining neutral about judgment and principle realism.

There is a distinction between particular moral judgments, general moral principles, and moral theories. A particular moral judgment ascribes a moral status to an individual action, person, state of affairs, character trait, or other object of moral assessment. For example, "That act was wrong" is a particular moral judgment. Moral principles are generalized judgments assessing kinds rather than individuals. E.g., "Killing is wrong" and "Telling the truth is right" are both moral principles. In T. M. Scanlon's characterization, moral principles "are general conclusions about the status of various kinds of reasons for action".³ Principles are complex judgments; the simplicity of "Killing is wrong" belies its complicated meaning and application. There are potentially infinite numbers of particular judgments, and certainly a huge number of moral principles. Moral principles exist at differing levels of generality and universality.

Moral theories are often characterized as high-level (i.e., most general and universal) moral principles. Here I work with a broader conception of moral theory.

plausibly be offered as realist theories, though he believes that only objectivism can ultimately do justice to realist claims.

³ T. M. Scanlon, *What We Owe to Each Other* (Cambridge, MA: Harvard University Press, 1998), p. 199.

When I look at theories, I see many different things. The concept "moral theory" does not seem to have necessary or sufficient conditions. Rather, theories seem to be united by family resemblance. Theories normally contain high-level explanatory (and often very formal) moral principles and defenses of substantive moral judgments. But they also typically provide an account of the scope of moral norms, an explanation of the nature of moral motivation, an answer to the normative question,⁴ a definition of "morality", an account of what distinguishes moral norms from non-moral norms, explications of the nature of moral agency and moral patiency, analyses of morally significant concepts, analytical frameworks including theoretical terms and categories, explanatory structures,⁵ metaethical theses, and more. Theories are complexes or clusters of such explanations. Some moral theories contain all of these features; some contain fewer.

Some may prefer to continue to identify moral theories with high-level principles and say that most of the items on my list are separate theories. In contrast, I conceive of moral theories as being normative explanations of certain normative phenomena. They are normative explanations in that they seek to justify, not merely to describe. To say

⁴ Following Christine Korsgaard, to answer the normative question is to explain what justifies morality's claims on us. See her *The Normative Question* (New York: Cambridge University Press, 1996), pp. 9-10.

⁵ By "explanatory structure", I mean a way of mapping some domain that explains trends and relationships within the domain. Examples of such structures include the deontic square of opposition in deontic logic and the periodic table in chemistry.

that Kant's moral theory is merely the Categorical Imperative misses the richness of his theorizing. However, ultimately, nothing hangs on whether a narrower conception of "theory" is better. The important point is that moral theorists do discuss all of the issues I have listed, so when I ask if one should be a realist about moral theories, I am talking about all of these issues. Crudely put, the relationship between theories, principles, and judgments, as I use the terms, is this: Theories explain principles; principles explain judgments.⁶ (This is not to say that judgments derive from principles or that principles derive from theory.)

This separation of theories, principles, and judgments already exists in the literature. Moral particularists accept that there are correct moral judgments, but deny that moral principles have any explanatory role.⁷ Anti-theorists, on the other hand, may accept that there are correct particular judgments and even moral principles, but deny that there are correct moral theories. E.g., Bernard Williams characterizes theories as follows:

⁶ Even W. D. Ross' moral theory is not merely a defense of a set of principles. It is also an account of how these principles work in moral reasoning, plus an explanation of why other moral principles (of which there may be indefinite numbers) reduce to Ross' listed seven.

⁷ E.g., Jonathan Dancy, *Ethics without Principles* (New York: Oxford University Press, 2004).

An ethical theory is a theoretical account of what ethical thought and practice are, which account...implies a general test for the correctness of basic ethical beliefs and principles..."⁸

Williams is often called an anti-theorist. He accepts and defends basic ethical beliefs and principles, but is skeptical of ethical theory so characterized. The characterization of theory I give is even broader, and an anti-theorist might also be skeptical of theory in that sense. "Killing is wrong" is not quite the same thing as "Wrongness consists of failure to respect humanity in one's actions". The latter principle accounts for the former. But an anti-theorist could accept the former while denying the latter and all of its competitors.

Thus, there are different things about which to be a realist. One could be a judgment, principle, and/or theory realist.

Actually, matters are more complex than that. As Christine Korsgaard notes, there are different senses of realism: procedural and substantive. A procedural moral realist holds that there are right answers to moral questions. (Or, at least, there are better and worse answers.) I.e., for most moral questions, there is a way to get a correct answer.⁹ In contrast, a substantive moral realist holds that there are moral facts to which

⁸ Bernard Williams, *Ethics and the Limits of Philosophy* (Cambridge, MA: Harvard University Press, 1985), p. 72.

⁹ Moreover, procedural realists are committed to the view that there can be correct *positive* answers to questions to moral questions. Error theorists and nihilists, e.g., would say that there is a correct answer to "What I ought I do?", namely, there is nothing one ought to do because there are no oughts.

true moral judgments correspond; the reason that some moral judgment is true is that there are normative objects or properties in the world that make it true.

Procedural realism is an epistemological position, holding that some moral views can be justified. Substantive moral realism is a metaphysical position, holding that there are normative entities or facts. The doctrines can be held independently. For instance, one type of normative skeptic might hold that there are moral facts, but such facts are inaccessible to us.¹⁰ To believe this is to accept substantive realism while rejecting procedural realism. On the other hand, constructivist deontologists accept procedural realism but reject substantive realism.¹¹ They accept that certain moral judgments can be justified, but deny that these judgments are about normative properties in the world. If normativity exists only within the practical perspective, then normative judgments have no special ontological commitments in the speculative perspective. Constructivists can thus be naturalists, agreeing that science is the final arbitrator of what exists.

Thus, there are at least six types of moral realism: substantive judgment realism, procedural judgment realism, substantive principle realism, procedural principle realism, substantive theory realism, and procedural theory realism. Each of these has different commitments. They are charted in Figure 1.

¹⁰ E.g., certain fundamentalist Christians hold that humanity is inherently depraved, so even though there are moral facts, people have no ability to determine them.

¹¹ Another recent example of procedural realism combined with substantive antirealism is Mark Timmons contextualist theory ("assertoric nondescriptivism") in his *Morality without Foundations* (New York: Oxford University Press, 1999).

	Substantive	Procedural
Judgment Realism	Moral judgments, such as, "That act was right," are true <i>because</i> there are normative facts that make them true.	There are correct procedures for making moral judgments. Some moral judgments are correctly accepted.
Principle Realism	Moral principles, such as, "Killing is wrong," are true <i>because</i> there are normative facts that make them true.	There are correct procedures for accepting moral principles. Some moral principles are correctly accepted.
Theory Realism	Moral theories, such as Kantianism, are true <i>because</i> there are normative facts that make them true.	There are correct procedures to accepting moral theories. Some moral theories are correctly accepted.

Table 4.1, Varieties of Realism

In the figure above, by "properties in the world", I mean to include nonnatural and supernatural properties. Thus, divine command theories, Ross' and G. E. Moore's intuitionist theories, and Plato's theory of the Good are varieties of substantive realism. When I say that some moral judgments/principles/theories are correctly accepted, I not only mean accepted, but also asserted and used in reasoning.

Figure 1 also can be used to map commitments. Typically, if one accepts a type of realism, one must accept at least one form of realism that appears in the grid above it.

I accept *multiplistic procedural theory realism*. In my view, certain particular moral judgments and principles are justifiable. (I will not defend this position here.) Let

us call this justified judgments and principles "the moral truths", though I do not mean to endorse substantive realism by calling them that. Rather, I mean "true" in the procedural realist sense. On my view, there are justified truths about morality. However, there is no uniquely correct way of representing those truths. I.e., no one moral theory is best. Still, there are better and worse ways of representing those truths relative to a given user of the representation. Truth is what a good person gets from a good moral theory. Truth is neither in the theory nor in the person. The same person can come to see moral truths using contradictory moral theories. These positions may seem strange. The remainder of this essay is meant to explain and defend them.

3. GRADING PHILOSOPHY EXAMS

Most professional philosophers are college professors with teaching duties. Most spend some time grading philosophy exams. In higher-level classes, the exams are almost never about reciting memorized definitions of various isms; they are critical essays evaluating different ideas and arguments. In this section, I want to articulate how our grading practices reflect a kind of procedural realism. As graders of exams, we are procedural realists in that we think some essays are better than others. Most of us are multiplists, too. We do not think that only one essay is good or justified. Over the course of our careers as teachers, we do not expect there to be one best exam ever.

What do we want from exams? One might say that we want the right answer, but philosophy exams are not like tests in mathematics. If I ask a student what is the first derivative of $f(x) = 6x^2$, she must answer that $f'(x) = 12x$ or she has gotten the answer

wrong. However, even in mathematics classes, teachers require that students show their work. They give credit for showing how one arrived at one's answer, i.e., for displaying one's reasons.

Some philosophy questions do have right answers in this sense. If I ask a student what Rawls' difference principle is, it is easy to tell the student he got it wrong.

However, my exams (and most others' exams) do not consist of such questions. I do not ask students to recite the difference principle; I ask them to *assess* Rawls' reasons for advocating it. I ask them what they think and why.

Essay exams in the humanities and many of the social sciences are often described as "not having a right answer". Many students unfortunately interpret this as meaning that such exams are mere vehicles for opining. In my own experience, some students have even expressed surprise and indignation that they receive bad grades for their opinions. "You asked me for my opinion. Here it is. Why didn't I get an A?" We do not grade opinions so much as *reasons* for those opinions. As I tell my students, good test answers tend to be complete, self-critical, clear, precise, insightful, concise, logical, organized, original, reflective, clever, and simple. Terms are clearly defined and positions well articulated.

Of course, we do grade opinions. Exams have an element of conservatism in them. If a student argues for a bizarre conclusion, she typically must do more to support that conclusion than others must. (On the other hand, she also can gain points for being interesting.) This is standard philosophical practice. It is perhaps unavoidable, but it is not altogether regrettable. On one hand, as teachers we worry about having undue

influence on the views of our students. On the other hand, arguments are always addressed to some audience with a point of view. There is no neutral point of view to address. To be potentially convincing, good arguments must start with relatively uncontroversial premises, but controversy is relative to different people. (A Biblical quotation is not usually a good premise in a philosophy class, but may be in a class on Christian theology.)

In my advanced classes, I have *open question* exams. While the exam is usually on a particular philosopher, I allow and encourage students to produce not only their own answers but also their own questions. For exams like these, the normal evaluative criteria apply. However, students also earn credit for asking good questions. A straightforward answer to a trivial question is worthless, but an honest, if incomplete, attempt to answer an intractable question can merit an excellent grade.

Ultimately, we think that some exam essays are better than others. We are not egalitarians: we do not hold that all essays are of equal value. Nor are we nihilists: we do not say that all essays are of *no* worth, because worth does not exist. As graders, we implicitly accept a type of procedural realism. We think that there are ways of determining whether one exam is better than another. The principles of good essay writing are only partially articulable. (Many of us distribute guides for essay writing to our students. Some famous guides are available on the Internet.) Our best theory of good essays might be a virtue theory or an ideal judge theory: A good essay is one that would be called a good essay by a good grader. A good grader is person that has the following virtues and interests...

Most of us accept a type of multiplism, as well. There is no uniquely good essay answer. Among the students who get and deserve A's, there will be a variety of opinions. In an exam on Rawls, some will defend the difference principle, some attack it, and (sometimes) some transcend it. We could not say the good students approximate the *one true good essay*, and deserve A's in virtue of their approximation. The students are going in different directions. Rather than their essays approximating one ideal essay, they approximate many different ideals.

We do not give out A's for producing true answers. We ask students to tell us what they accept and why. Since the good essays contradict each other, we cannot say that what they have in common is their approximation of the truth. However, truth is not irrelevant to grading. Good essays give us reasons to think their theses are true. The best essays help us to understand something we did not understand before, even if the essay defends some position that we cannot affirm.

As graders, we need not be *substantive realists* about good essays. Thinking that one answer is better than another does not commit us to accepting the existence of normative facts.

4. PROFESSIONAL PHILOSOPHY AS AN OPEN QUESTION EXAM

Being a philosopher is much like taking a lifelong open question exam. Philosophers choose which questions to answer. They often invent their own questions. They do not seek grades, but most do seek to present and publish their work, which entails taking the judgment of other philosophers seriously. The standards for

professional philosophy are harder than the standards for undergraduate essays, but the principles of assessment are mostly the same.

Most philosophers implicitly accept or practice a type of procedural realism about philosophy, including moral philosophy. Regardless of whether they accept normative facts, they hold that there are better and worse theories, arguments, articulations, and so on. They hold that some philosophers do better work than others. They think that some philosophies (even those that advocate epistemological nihilism) merit more serious consideration than others. Like all other disciplines, the practice of philosophy is thoroughly normative, in that philosophers have methods, and judge works as better and worse. Anyone who thinks philosophy is not completely hopeless, who thinks that some philosophies are worth holding, is at least a procedural realist about philosophical theories. Procedural realism about moral theories is one instance of this.

The typical moral philosopher seems not only to be a procedural theory realist, but a *multiplistic* procedural theory realist. They do not evaluate each other as if there is one, true, ideal moral philosophy, and the good philosophers are those who have produced reasonable approximations of that ideal. The philosophers typically have significant disagreements, so if they are approximating ideals, they approximate different ones. Justification among disagreeing philosophical theories mirrors justification among disagreeing student exams.

5. THE PLURALITY OF THEORETICAL VALUES

Philosophers accept many different moral theories. These theories fall into broad types or research traditions, including Kantianism/deontology, consequentialism,

pluralism, contractarianism, and virtue ethics. There are multiple theories inside each tradition, e.g., the virtue ethics of Aristotle versus the Stoics, or the consequentialism of Jeremy Bentham versus John Stuart Mill. One reason (among many) why there are so many different theories is that there is a plurality of *theoretical values*, i.e., the values by which a theory is assessed. Doing justice to these theoretical values does not obligate one to accept exactly one theory.

Here is a partial list of the theoretical values of a moral theory. A good theory provides an excellent normative explanation of morality, fits (or explains away) our moral intuitions and considered judgments, is internally consistent, fits with theories outside moral theory, can be justified from both impartial and personal points of view, is parsimonious, is general, solves a number of theoretical problems, and is fecund in suggesting new avenues of philosophical investigation. These are some of the many criteria by which philosophers have judged theories. There are others I have not listed. Some philosophers would object to some items on the list. Still, any reasonable list will have more than one item.

Typically, moral philosophers have concern for these values, and yet continue to produce different theories. One reason for this is that they weigh the values differently. It is not obvious that all disputes in weights can be rationally resolved, much like it is not obvious that all disputes in the relative weights of moral values can be rationally resolved. The theoretical values can often conflict. E.g., impartial justification and intuitive fit quite often pull in different directions. Utilitarians typically are more committed to the

theoretical value of impartial justification while pluralists are more concerned with intuitive fit.

We might want to say that a good theory is a function of the listed theoretical values. However, in technical terms, functions yield unique results. If f is a function of x , y , and z , there will be exactly one f for every (x, y, z) triad. Moral theories are attempts to do justice to certain theoretical values, but they do not appear to be functions of those values in this technical sense.

I did not include truth in the list above. Someone might claim that the best moral theory is the *true* one. Yet, truth is only one of many desiderata. Suppose the true theory is so complex that no one can understand it. In such a case, if we wanted to understand morality via a theory, we might have to accept a simpler, but partially false theory. A true theory that no one can understand is useless, and thus a bad theory.

Many moral philosophers claim that there is a plurality of *moral* values and thus advocate a pluralistic moral philosophy. Similarly, if there is a plurality of *theoretical* values, we might accept pluralism *about theories*. The moral pluralist says that there are many justifiable basic values and moral principles. The theory pluralist says that there are many justifiable theories. There can be both social and personal versions of such theory pluralism. The social version holds that it is reasonable for different moral philosophers to accept different theories. This is the de facto position of many tolerant philosophers. They may disagree with their colleagues, but they believe that their colleagues' views are more-or-less justified, in a sense. (In the big open-question exam, their colleagues deserve A's.) The personal version of theory pluralism holds that it is

reasonable for one moral philosopher to accept multiple moral theories at the same time.

This version is rare.

6. THEORIES AS EXPLANATORY REPRESENTATIONS

Moral theories exhibit significant variety in their form, content, and function. However, they have a common cluster of questions they attempt to answer. They have many other features in common.

Moral theories are intentional objects. They are *about* something. Their *subject matter* is morality. Moral theories—like other theories—are linguistic representations. They represent the moral through words and concepts. Moral theories should not be identified with morality itself. To put it roughly, "morality" refers to those norms, codes of behavior, character traits, and ends that rational agents should adopt. A moral theory is a certain type of linguistic representation of those norms.

We ask why does morality have its characteristic importance or why should we heed its demands. Moral theories attempt to answer these and other questions. They do not merely articulate the nature of morality, but seek to explain it. Theories are representations that render the moral comprehensible, by providing a new perspective on it.

Since morality and moral theory are different things, it is possible to agree about morality but disagree on the best theory of morality. I.e., you and I could accept the same set of moral judgments, laud the same character traits, seek the same goals, abide by the same norms, and so on. But when we are asked to produce a theory—a linguistic explanatory representation of the moral—we could produce different theories. I produce

some version of rule consequentialism while you produce Rossian pluralism. Not all disagreement between moral theorists is *moral* disagreement. Some is merely theoretical, i.e., over how best to represent and explain the moral. Some of this theoretical disagreement occurs because we attach different weights to theoretical values, such as fecundity, simplicity, generality, and so on.

Theories are representations. There is only one world, one reality, but it does not follow that there is only one good or correct representation of that reality. It is common in an electronics store to see hundreds of televisions each playing the same movie. The picture size, detail, crispness, clarity, and color are different between sets. Some sets are superior to others. But even among the high-end sets, rarely does one dominate on all possible measures. Choosing a TV involves a choice about what type of representation one prefers. It involves economy, as one has to pay more for a better set. The opportunity cost of successive marginal improvements grows exponentially. A slightly better TV might not be worth the expense. Theories are like televisions in these respects.

Televisions have a sort of realism. But it is strange to ask which set has the *real* picture, simpliciter. Some sets are more realistic than others, but even among the good sets, the colors and contrasts tend to be sharper than in reality. This is true of theories as well. Also, no matter how good the picture is, one cannot walk behind the TV to see the back of the image. Every representation, from television to theory, omits something.

Though theories and televisions have these problems, there is still a distinction between good and bad. False theories are like television sets playing the wrong shows. They might have exquisite detail and great beauty, but they do not represent reality.

Alternatively, vague, poorly formulated theories often are like low-fidelity television sets, showing the right show with a poor picture. I agree with postmodernists that theories are artifacts—human constructions. But it does not follow from the fact that something is an artifact that all artifacts are equal, or even that there are always alternative artifacts.¹² A PRS Dragon is a better guitar than a Squier Bullet, a Corvette is a better car than a Gremlin, and a Panasonic LCD flatscreen is a better television than the tubes our grandparents watched.

Of course, televisions are not our only representational device. The movie *Schindler's List* represents the same events, more or less, as Thomas Keneally's earlier novel. There are major differences between film and prose. Paintings, sculptures, music, and even dances can represent the same events as film and writing. Any form of representation comprises multiple genres. E.g., in painting, there is impressionism, expressionism, romanticism, etc. Each form and genre has different strengths. Yet, no form or genre of representation dominates on all measures of quality. It is unclear that they are even commensurate. It makes sense that some people will prefer one type of representation to another. I prefer written to television news. But it does not make sense to say that film is better than photography, or paintings are better than music, tout court.

There are two lessons from these reflections. One, there are many ways of illuminating and revealing the same truth. Two, some theories are more like films, some

¹² For discussion of this fallacious reasoning, see Nicholas Shackel, "On the Vacuity of Postmodern Methodology", *Metaphilosophy* 36:3 (2005), 295-320, esp. pp. 311-314.

are more like paintings, some are more like interpretive dances. Compare an analytic journal article to a Platonic dialogue.

Some theories are better than others, but given that there are different types, we need not claim that one type is better than another as such. Among the philosophical community, there is room for pluralism, variety, and tolerance. Disagreement on philosophical doctrines need not entail that theorizing is irrational anymore than disagreement on lifestyles in liberal society entails that lifestyle choice is irrational. Some lifestyles and doctrines are bad, but many are good.

7. BEING TRUE VERSUS ILLUMINATING TRUTH

In its original meaning, to illuminate is to make visible with light. The concept itself, when applied to ideas, is a metaphor, though now a dead one. To illuminate an idea or a truth is to make something understandable, literally by shining light on it.

Philosophers normally claim that their goal is truth. One possible way of seeking truth is to produce true statements and true theories. Another way is to produce statements and theories that are illuminating, that make the truth understandable. Both options are respectable.

Why seek to illuminate truth rather than state it? Here we should take note of how metaphors work. We understand something when we can see it in different lights, integrating it with experience and previous ways of thinking. In Newtonian mechanics, thinking of gravity as an attractive force "pulling" objects together was metaphorical. Metaphors are meant to illuminate, and do so without being true. (I do not intend to take a stand on the semantics of metaphor. However, most philosophers agree that "living"

metaphors do not state literal truths, though they debate whether there is a separate "metaphorical truth".) A truly striking metaphor allows us to understand truths that language is too crude to state literally. Alternatively, a good metaphor provides a convenient, concise cognitive package. Even when our language is sophisticated enough for us to state all of the truths about something, it might take pages. A good metaphor can package those truths in one sentence.

Take the following famous quotation, also from Shakespeare:

Life's but a walking shadow, a poor player
 That struts and frets his hour upon the stage,
 And then is heard no more. It is a tale
 Told by an idiot, full of sound and fury,
 Signifying nothing.¹³

A few things to notice about this quotation: First, taken literally, Shakespeare contradicts himself. Life cannot literally be a walking shadow, a poor player, *and* a tale. Second, these lines are a linguistic representation of some truths, or at least truth as the disconsolate Macbeth (the speaker) sees it.

Third, and most important, relatively few words are needed to illuminate all of these truths. The lines shed light on something true (or what Macbeth believes to be true), without themselves being literally true. Scholars have written many lengthy essays

¹³ William Shakespeare, *MacBeth*, V.3.23-27.

explaining all of the ideas packed into these few lines. To some degree, the essays never fully explicate the lines. The lines are richer than any explication could be. Good metaphors admit of inexhaustible extension.

Good moral theories might be like metaphors in these respects. Take the Categorical Imperative. The Categorical Imperative concisely packages a huge array of complicated ideas in one convenient package. To some degree, the Categorical Imperative just is a metaphor. Kant tells us that right actions are those that abide by laws in a certain hypothetical kingdom, and wrongness is treating people like hammers.

When scholars try to explicate the quotation from *Macbeth*, they do not claim that the lines *entail* all the sentences in their essays. Nor do they say that the lines are merely *shorthand* for longer essays. Similarly, the Categorical Imperative shines light on particular moral judgments and general moral principles. But it can do so without entailing or being shorthand for all of the judgments and principles. The type of understanding a theory provides may very well be like the type of understanding a metaphor provides.

If so, then this explains why the best varieties of rule consequentialism, deontology, Rossian pluralism, virtue ethics, contractarianism, and so on, can all be good theories, despite their disagreement. Theories are meant to provide understanding. A good theory provides it. Moreover, a good theory *is worthy of acceptance*; it is justifiable. As I argued above, moral theorists implicitly accept procedural theory realism, which is an epistemological notion. I have said that more than one theory can be justified, despite having genuine disagreements. One might object to my view by saying

that since the theories disagree, they cannot all be tracking the same truth. However, contradictory metaphors can be used to illuminate the same truths, or to illuminate different but compatible truths about the same thing. If theories are relevantly similar to metaphors, then contradictory theories can illuminate the same truths or different but compatible truths.

This is not to deny genuine conflict. I am not claiming that theories do not really contradict each other. Nor would I deny that many bad theories fail to illuminate the truth at all. As poets, we could disagree about metaphors. Poet A might understand some truths through the metaphor written by poet B, but still think that B's metaphor is poor or that B should use a different metaphor instead. Similarly, a moral theorist can accept that his competitors' theories shine light on morality, but she can still believe that their theories are not very good and should be replaced.

Theories represent the moral. Because there are many different justified theories, there are multiple good ways of representing the moral, and the representations differ. None is perfect. They contradict each other. We might want to say that representations are each true of the represented in their own way, or we might want to say that representations are not true. Regardless, the reality—the moral—is real in a way that theories—the representations—are not.

8. TRUTH IN THE RELATIONSHIP BETWEEN THEORY AND USER

Metaphors illuminate truths because they prompt us to adopt a perspective from which we can see two distinct things as identical. Life is difficult to understand, but once we see that it is a poor player that struts and frets and then is gone, we understand it

more. Metaphors often equate one difficult to understand (or characterize) thing with something more understandable.¹⁴ Sometimes two difficult notions are rendered understandable by showing their equivalence.

A striking metaphor has depth. A perfect metaphor could be extended without limit. Suppose the perfect metaphor equates X with Y. From the perspective we adopt via the perfect metaphor, there always will be another property of Y that X also possesses. (However, this does not imply that X also possesses every property of Y.) Perfect metaphors do not run out.

In some ways, the success of a metaphor depends upon its audience. Their success is not determined by the semantics alone, at least in the sense that the understanding one gets from a metaphor does not come simply from knowing the definitions of the related terms. To understand Macbeth's cynicism about life when he says that life is a poor player, one needs to know more than the definition of "player". One has to understand something about the motives of actors, the nature of theater, and more. Children know what an actor is, but Shakespeare's metaphor will be lost on them.

John Stuart Mill said that universal idiocy could render any moral standard useless or worse.¹⁵ He was describing a standard to be used in ethical reasoning for

¹⁴ In I. A. Richards' terms, these are the *tenor* and *vehicle*. The tenor borrows attributes from the vehicle. In George Lakoff's terms, these are the *target* and the *source*. See I. A. Richards, *The Philosophy of Rhetoric* (New York: Oxford University Press, 1936) and George Lakoff and Mark Johnson, *Metaphors We Live By* (Chicago: University of Chicago Press, 1980).

making decisions. However, his point applies to both theories and metaphors more generally. A good metaphor is rendered useless or misleading when conjoined with universal idiocy: "Life is a play? You mean life is likely to have majored in theater as an undergraduate?" Similarly, a good theory is rendered useless when conjoined with universal idiocy. The Kantian Categorical Imperative may illuminate much about morality. However, if one deliberately adopts the perspective of an idiot, one will interpret it as claiming that no one can flush toilets.

Moral theories are useful tools for understanding morality. Like any other tool, they can be misused. They require good faith from their users. Moreover, like most any other tool, they require skill from their users. My father is more effective at using a scroll saw than I am. I cannot make fancy designs with the saw, but it is not the saw's fault.

Theories can illuminate the truth. But they do not do so all by themselves. They illuminate only when they are skillfully operated upon by honest users, users who are open to finding what the theories can illuminate. The point of philosophy classes is not just to teach theories and the reasons offered for the theories, but to teach students how to be honest users of theory.

Suppose one buys a book explaining how to do simple home improvements. If one is unable to do the improvements after practicing the techniques described in the book, it could be the book's fault. Perhaps it is not a good guide. However, it could also be one's own fault. One might be bad at following directions, or simply not have the native talent to develop any home improvement skills. These points could hold for a

¹⁵ John Stuart Mill, *Utilitarianism* (Indianapolis: Hackett University Press, 2001), p. 24.

theory, too. If one does not gain any understanding from a moral theory, it might be a bad theory. Or, it might be that one is a bad user of a theory.

Truth is what a good user gets from a good theory or even from multiple good theories. Accordingly, if one wants to learn the truth about morality through moral theories, one must do more than seek the best theories. One must seek to be receptive to what theories can show.

CHAPTER FIVE

THE MORAL COSMOS: ON BEHALF OF GENERALISM IN ETHICS

1. THE DAPPLED MORAL WORLD

Almost no philosopher accepts straightforward falsificationism in the philosophy of science anymore. Why do philosophers act as though falsificationism applies to philosophy? We no longer think that providing counterinstances to scientific principles suffices to show that such principles are bad explanations. Moreover, we recognize that scientific explanations have limited applicability; they do not hold for all domains. We may not ever get a single best theory of the universe—one law for all entities—so perhaps we must make do with a multiplicity of imperfect models to do our explanatory work. Perhaps some things will remain forever unexplained.¹ Why should it be any different in moral philosophy?

¹ In a recent interview, prominent theoretical physicist Stephen Hawking endorsed just such a picture of cosmology. Though he once accepted that science aimed for fundamental unity and simplicity, he now believes that at best we can produce a patchwork of theories and models. These theories contradict each other (one describes the universe as 10-dimensional everywhere, one says is 11-dimensional everywhere), but we recognize that our theories have limited applicability and are essentially incomplete tools. See Leonard Mlodinow, "Prometheus Bound", *Discover* 28:10 (October, 2005),

Morality is a cosmos, not a chaos. That said, just how orderly is this moral cosmos? My view of morality is much like Nancy Cartwright's view of the nature: ...as appearances suggest, we live in a dappled world, a world rich in different things, with different natures, behaving in different ways. The laws that describe this world are a patchwork, not a pyramid. They do not take after the simple, elegant and abstract structure of a system of axioms and theorems. Rather, they look like — and steadfastly stick to looking like — science as we know it: apportioned into disciplines, apparently arbitrarily grown up; governing different sets of properties at different levels of abstraction; pockets of great precision; large parcels of qualitative maxims resisting precise formulation; erratic overlaps; here and there, once in a while, corners that line up, but mostly jagged edges; and always the cover of law just loosely attached to the jumbled world of material things. For all we know, most of what occurs in nature occurs by hap, subject to no law at all. What happens is more like an outcome of negotiation between domains than the logical consequence of a system of order.²

Morality is like this. It consists of different domains largely demarcated by different types of relationships we have with others, with ourselves, and with various things.

64-65. See also Stephen Hawking, "Gödel and the End of Physics", Lecture: Dirac Centennial Celebration, Cambridge: University of Cambridge, 2002.

² Nancy Cartwright, *The Dappled World* (Cambridge: Cambridge University Press, 1999), p. 1.

There are pockets of great precision regarding certain moral principles. Sometimes rightness is subject to no law at all. The talk of negotiation between domains seems apt.

The current debate between generalism and particularism in ethics largely concerns these issues. Particularism is the doctrine that there are no moral principles, but morality gets on just fine without them.³ The particularist David McNaughton avers that principles hinder moral reasoning. Individual cases must be assessed on their own merits using correct perception or judgment.⁴ In contrast, generalism is the thesis that there are correct moral principles. Particularists complain that some forms of generalism make moral judgment seem unnecessary. One could acquire the correct principles and then simply derive correct judgments in each case.⁵ Moral theories attempt to codify morality such that one can rely on theory to make correct decisions, much like one relies on a troubleshooting guide.⁶ However, few moral generalists have held this. Even Immanuel Kant, touted the arch-anti-contextualist, maintained that right action depends largely on context and cannot be mechanically derived from the Categorical Imperative.⁷

³ Jonathan Dancy, *Ethics without Principles* (Oxford: Oxford University Press, 2004), p. 1.

⁴ David McNaughton, *Moral Vision* (Oxford: Blackwell, 1988), p. 190.

⁵ McNaughton, *Moral Vision*, p. 198.

⁶ McNaughton, *Moral Vision*, p. 192.

⁷ E.g., Immanuel Kant, *The Metaphysics of Morals*, trans. Mary J. Gregor (Cambridge: Cambridge University Press, [1797] 1996), 6:469. See Loudon, Robert B., *Morality and*

Particularism is the thesis that there are no "features of actions or other objects of evaluation that are always morally relevant and relevant in the same way."⁸ Another defining statement is not exactly equivalent: "The particularist claims that the moral relevance of any feature depends entirely on context."⁹ Per the particularist, morality is so context dependent that no useful generalizations can be attributed to it. In its strongest form, particularism claims that in moving from the nonmoral features of situations, actions, and character traits in context to their proper moral evaluation, we are dealing with an essentially random phenomenon admitting of no generality.¹⁰ The properly trained moral practitioner can discern the appropriate moral evaluation, but what is discerned via this skill is completely uncapturable using non-moral language. The only thing that right actions have in common is rightness.¹¹

Moral Theory (Oxford: Oxford University Press, 1992), pp. 99-124 for a defense of this interpretation.

⁸ Mark Timmons, *Moral Theory* (Lanham, MD: Rowman and Littlefield, 2002), p. 278.

⁹ Timmons, *Moral Theory*, p. 245.

¹⁰ Frank Jackson, Phillip Pettit, and Michael Smith, "Ethical Particularism and Patterns", in *Moral Particularism*, (eds.) Brad Hooker and Margaret Littlefield (Oxford: Oxford University Press, 2000), p. 86.

¹¹ Dancy accepts that there might be some invariant reasons. Perhaps gratuitous pain always counts against an action; perhaps torture is always wrong. But such cases are so rare, he thinks, that they are the exception that proves the rule. See Jonathan Dancy,

Given the picture of morality I have asserted, it may seem as though I take the particularist's side. Not so. There is a role for moral principles in our patchy moral universe, a stronger one than the particularist permits. On Cartwright's conception of science, science seeks to produce models of the natural patchwork, models that work with precision inside certain domains, and have significant but limited use outside those domains. Similarly, I see morality as being partially capturable by laws and models. As such, this essay is a defense of moral generalism, the thesis that there are moral principles. By extension, it is also a defense of moral theory. My defense proceeds by reconceiving the structure and nature of moral theories, as well as via a programmatic argument for limited generality.

I will be defending a weak, patchy generalism against particularism. I do not intend here to convince the generalist who sees morality as strongly, elegantly unified to accept the weaker version. That is another project. For the purposes of this paper, this generalist and I are comrades. This other generalist may see me as defeating particularism on the particularist's own ground.

2. ARGUMENTS FOR PARTICULARISM

Particularists hold that there are no patterns or principles in moving from the nonmoral descriptions of attitude, actions, and states of affairs to their correct moral

"The Particularist's Progress", in Hooker and Little, *Moral Particularism*, p. 136; Dancy, *Ethics without Principles*, p. 77.

evaluation. There are two common arguments for particularism. One is the Argument by Counterexample; the other is the Switching Valence Argument

The Argument by Counterexample is largely inductive. A generalist proposes some moral principles. For example, the generalist might propose that an action is right just in case it maximizes happiness. That is, there is a law-like generality in going from the nonmoral feature of maximizing happiness to the moral feature of being right. The particularist responds with a counterexample, producing an example of an act that maximizes happiness but that both the generalist and the particularist agree is wrong. The generalist responds by producing a better (hopefully not ad hoc) set of principles, and the particularist responds with another counterexample. And so on. The inductive step is moving from the fact that many counterexamples have been produced to the conclusion that a counterexample is in principle always available.

This argument has flaws. For one, sometimes counterexamples do not matter and do not count against principles. I elucidate this in §3 below. More importantly, as Dancy has observed, it makes it seem as though the debate between generalism and particularism will be decided by attrition—whichever side first relents loses. What we want, Dancy effectively says, are principled, general reasons for accepting particularism.¹²

¹² Dancy, *Ethics without Principles*, p. 2.

Because of his worries about attrition, Dancy employs the Switching Valence Argument, an argument against reasons atomism on behalf of reasons holism.¹³ Reasons atomism is the view that reasons are like little atoms with the same valence—propensity to count in favor of or against an action—in every situation. E.g., a reasons atomist might hold that an act's causing happiness always counts in favor of that action, even if sometimes other features of the act might outweigh this positive contribution and make the act wrong. Dancy champions reasons holism. In the theory of reasons, holism is the doctrine that "a feature that is a reason in one case may be no reason at all, or an opposite reason, in another."¹⁴ Whether promoting happiness, causing harm, or any other feature of an action, counts in favor or against an action depends entirely on context. When we combine the feature of causing happiness with another feature—sadism—the feature of causing happiness no longer contributes positively to the evaluation of the action. In the case of sadistic torture, it's not that causing happiness among sadists has a positive valence, but the goodness of the happiness production is outweighed by the torture. In this context, the feature of causing happiness has a negative valence. The action is made worse for causing happiness. Torture is evil enough, but torture plus the production of sadistic pleasure is worse.

Dancy generalizes from these cases, concluding that all reasons are holistic.

Whether a feature counts in favor or against an action's being done, or a proposition's

¹³ For one instance of the Switching Valence Argument, see Dancy, "The Particularist's Progress". (Dancy does not use these terms.)

¹⁴ Dancy, *Ethics without Principles*, p. 7.

being believed, etc., always depends on context. No features have a simple positive or negative valence. Their valence is determined by their combination with other features, and every action occurs in a context with a unique set of features whose relevance cannot be determined in advance. Dancy then argues that holism about reasons points strongly to particularism. It would be "some sort of cosmic accident" if morality or moral reasons were then capturable by principles. Certainly one would not need principles in order to act morally or be a good person.¹⁵

Dancy is mistaken about how the shifting valence of reasons relates to moral particularism. I will argue in §4 below that if we take his metaphor of reasons *atomism* seriously and look at how actual atoms—the ones studied in chemistry—combine, we will see that the issue of shifting valences does not matter. Indeed, it opens a door for an interesting and largely unexplored form of generalism.

I have one last difficulty with particularism. Particularists paradoxically elevate particularism to a type of generalism (though a generalism operating on a different level). Particularists argue that their theory of reasons holds *invariantly*, that reasons are always particular. In §7, I argue that in some cases, particularism is the best account of reasons, but in other cases, generalism is a better account. Some reasons are particular; others are not. Perhaps Dancy is right that that for some cases, reasons that are reasons elsewhere are not reasons there. (Similarly, laws that hold everywhere else in the universe might not hold in a particular spot. The universe might be patchy.) But perhaps most of the time reasons do function as the generalist describes.

¹⁵ Dancy, *Ethics without Principles*, p. 82.

3. PRINCIPLES AND ANOMALIES

Particularism is sometimes defended via the Argument by Counterexample. It is alleged that there are counterexamples to any possible set of moral principles. Thus, morality is not fully capturable by principles. Even if this is true, and I believe it is, this need not entail particularism.

Consider an analogous case. Most professional philosophers teach philosophy classes; some supervise graduate students. Something philosophers teach their students is how to write philosophy papers. For that purpose, we often produce philosophy-writing guides. There are famous on-line examples. There are published guides as well, such as Joel Feinberg's *Doing Philosophy: A Guide to the Writing of Philosophy Papers*. Feinberg is a good example because he was a superlative writer—a *phronimos* of writing.

Such guides contain *principles* of good writing. Feinberg's guide reads like a version of Rossian pluralism. He enumerates a number of *pro tanto* principles: aim for clarity, simplicity, and economy; avoid padding, repetitiveness, redundancy, misplaced emphasis, and pretension.¹⁶ These principles can conflict. There is no algorithm for determining the outcome of conflicts. Skilled judgment is required.

Should one be a particularist about writing philosophy papers, and thus say that such guides are mostly useless, or even a hindrance to the writing of good philosophy? The Argument by Counterexample can be run with any apparent set of principles of good

¹⁶ Joel Feinberg, *Doing Philosophy: A Guide to the Writing of Philosophy Papers* (Belmont, CA: Wadsworth, 2002), pp. 22-29.

philosophical writing. For example, Feinberg asserts that one should avoid combativeness, preachiness, vagueness, and pretentious, fancy words. These principles seem true. Yet, they have a counterexample. Friedrich Nietzsche violated all of these principles, but his philosophical writing was excellent. It is not that Nietzsche's vagueness and pretension counted against him, but this was overridden by positive points in his philosophizing. Rather, in Nietzsche's special case these points actually counted in favor of his writing. Nietzsche's writing would have been worse if he had tried to be more like Feinberg, even though Feinberg was a superlative writer. However, Nietzsche was an *anomaly*. Hegel's and Derrida's vagueness and pretension were terrible faults. Hegel would have been much better had he internalized Feinberg's principles.

Similarly, in the guides I produce for students, I often advise them to consider and rebut objections. This is common practice in analytic philosophy—one of its best practices. It has counterexamples. Edmund Gettier's "Is Knowledge Justified True Belief?" considers no objections, but it is an excellent paper despite that. *A fortiori*, Gettier's paper was *better* for having not considered objections. Gettier's work sparked an intense debate in epistemology. In its special context, its radical concision is a virtue. However, Gettier's paper is an anomaly.

Anomalies are common in any theory, including theories of good philosophy writing, philosophical theories, and scientific theories. One philosopher of science famously claimed that all theories are born refuted because they inherit the anomalies of

their predecessors.¹⁷ More plausibly, one does not take the presence of an anomaly to be, by itself, a refutation of a theory or a principle. Thomas Kuhn explained that anomalies are always present in scientific theory, but science typically requires a pervasive aggregation of anomalies to prompt a revolution. Larry Laudan writes, "Theories are generally not rejected simply because they have anomalies nor are they generally accepted simply because they are empirically confirmed."¹⁸ Laudan argues this a *good* feature of scientific theorizing. Cartoon falsificationism, where one counterinstance suffices to debunk a theory, is simply a bad picture. Almost no contemporary philosopher of science advocates such radical falsificationism, and scientific practice does not abide by it.¹⁹

Feinberg's principles of writing are good, even though there are anomalous counterexamples. A scientific theory can be good, even if there are anomalies. By analogy, why can't moral principles be good, even if there are always anomalies?

Consider, for example, the principles advocated by the Rossian pluralist. One of the enumerated pro tanto duties is a duty of non-maleficence, which includes a duty not to kill. To my knowledge, no person I have met has ever been in a situation in which it was right to kill another person. With near certainty, I can estimate that no one I know

¹⁷ Imre Lakatos, *The Methodology of Scientific Research Programs, Vol. I* (Cambridge: Cambridge University Press, 2000), p. 5.

¹⁸ Larry Laudan, "A Problem Solving Approach to Scientific Progress", in *Scientific Revolutions*, ed. Ian Hacking (Oxford: Oxford University Press, 1981), p. 144f.

¹⁹ Though searching for disconfirmation still plays an important role.

will ever enter a situation in which it is right to kill. I can specify ex ante a list of situations under which it would be right to kill another: e.g., self-defense, just war, Hobbesian state of nature. If someone I know ever rightly kills another, I can predict with near certainty it will be in a situation on my list. Perhaps, though, he will encounter some anomalous situation, not yet on my ex ante list, in which a person of good judgment would agree that it was right to kill. But this is unlikely.

It seems that a number of other Rossian principles can be expanded this way: "Lying counts against an act except in these cases..." Most of us will give the same list when asked to fill out the principle. The list will be incomplete. However, some philosophical or literary creativity (or bad luck) is needed to find cases our lists do not cover. If that is the case, are we really to reject Rossian generalism in favor of particularism? Rather, it seems that the particularists, with their philosophical ingenuity, are producing anomalies, and anomalies by themselves do not rule out theories or principles.

The particularist may respond that there remains a problem with such generalism. He might say that the generalist principles I have uncovered take the form of Humean laws—they are merely descriptions of contingent generalities in ethics. The right-making features still inhere in the particular. We fortuitously inhabit a universe in which it is almost always wrong to kill another. (Even anomalies have general characteristics. The cases where it is right to kill fall into a few categories). This response has problems. For one, it commits the particularist to rejecting the Humean view of laws in the philosophy of science, and this might be more than the particularist wants to do. Also, it seems that

such laws and principles have explanatory power. Feinberg's principles of good writing certainly do.²⁰

Anomalies do not instantly refute a set of principles, but they cannot be dismissed. The pressure anomalies place on a theory depends largely on whether competitor theories can explain the anomaly.²¹ Particularism does not explain any anomalies, because it is explanatorily empty. By way of parody, consider a particularist scientific explanation: "The reason this phenomenon occurs in this instance is that all of the causally relevant factors produced that outcome." Similarly, a particularist theory of the right says that an action is right in situation *S* in virtue of being favored by the morally relevant features of *S*.²² This is not a competitive explanation that better deals with the anomalies or holes in generalist principles. It is a truism. The particularist explanation of right action holds for Kantianism, utilitarianism, egoism, and every other moral theory.

4. REASONS ATOMISM AND THE PERIODIC TABLE OF REASONS

Reasons atomism is usually rejected because of the Switching Valence Argument. A reason cannot be like an atom, for whether it counts in favor of an action, against it, or

²⁰ Cf. Dancy, *Ethics without Principles*, pp. 111-117.

²¹ See Larry Laudan, *Progress and Its Problems* (Berkeley: University of California Press, 1977), pp. 36-40.

²² Timmons, *Moral Theory*, p. 254.

is irrelevant depends on context. The valence of a reason depends on the reasons with which it is combined. Thus, no general theory of reasons can be given.

The employment of chemistry terminology is fortunate, from my perspective. It illuminates the ways in which the characterization of reasons just given does not count against generalism. Rather, it counts *in favor* of generalism, because it describes reasons as being strongly analogous to actual atoms, the atoms studied in chemistry.

In chemistry, the term "valence" is synonymous with "oxidation number". It is defined as:

1. The combining capacity of an atom or radical determined by the number of electrons that it would lose, add, or share when it reacts with other atoms.
2. A positive or negative integer used to represent this capacity.²³

Actual valences possessed by actual atoms vary according to context. The fluorine atom is unusual in having only one oxidation state: -1. The chlorine atom, which is in the same group as fluorine and has many of the same chemical propensities, has many oxidation states: -1, +1, +3, +5, +7. Some atoms have no oxidation state. E.g., most noble gases do not combine with anything—under normal conditions.

Any former chemistry major recalls the explanatory power of chemical oxidation states. She is aware that quantum mechanics provides a good, if imperfect, explanation of these states. Quantum mechanics is messy, as per Cartwright's dappled world, and its

²³ "Valence", *The American Heritage Dictionary of the English Language, Fourth Edition* (Houghton Mifflin Company, 2004). *Answers.com* GuruNet Corp. Accessed 06 Jul. 2005. <http://www.answers.com/topic/valency>

messiness is increasing with scientific progress. As Ian Hacking remarks, "Every single year since 1840, physics alone has used successfully more (incompatible) models of phenomena in its day-to-day business, than it used in the preceding year."²⁴

The valence/oxidation state of an atom depends on its context, because the way atoms combine depends on what they are being bonded with and the conditions (heat, pressure, etc.) in which the bonding occurs. Valences shift frequently, though not so frequently that we cannot produce the generalist periodic table of the elements. We can also predict what valence an atom will take when placed in context.

Similarly, the valence of many reasons depends on context, e.g., on the reasons with which they are combined. By itself, that does not count against reasons atomism any more than the context-dependency of chemical valence counts against chemical atomism. The particularist argues that reasons cannot be atomistic because they do not have constant valences. But atoms do not have constant valences either. Instead, they have a range of valences that they will take under most specifiable conditions, or at least the conditions in which chemists are interested. The particularist must argue not merely that valences of reasons shift, but that they are radically disunified, such that they are not characterizable at all. The particularist needs to show that nothing like a periodic table of reasons can be constructed.

Philosophers have not tried producing such a table. Ross' pluralism is barely a start. Most moral theorists wish to skip the table in favor of a grand unifying principle beneath it. They would rather do string theory physics than chemistry. Yet, such a table

²⁴Hacking, *Representing and Intervening*, p. 218.

seems like a promising idea. The valences of reasons will of course be qualitative, not quantitative. Morality seems very complicated, but not as complicated as chemistry. Still, a high degree of complexity is not the same as chaos, which the particularist needs. Complexity is opposed to simplicity, not to order. As argued above, there appears to be a significant degree of order in morality, despite the presence of anomalies. Under normal conditions, acting on behalf of the general welfare nearly always possesses positive valence, killing another human being almost always has negative valence, etc. Chlorine nearly always has a -1 oxidation state, etc.

Some evidence that such a theory of reasons would be possible is the plausibility of a new class of pluralistic theories in political philosophy. David Miller's *Principles of Social Justice* and David Schmitz's *Elements of Justice* both argue for a patchy plurality of principles of justice, in which the appropriate principles are largely determined by the types of relationships holding between people.²⁵ Additionally, Michael Walzer's *Spheres of Justice* has a similar form of organization based on the social goods mediating relationships.²⁶ Even if one disputes these philosophers' substantive moral judgments, the *form* of their theories might still seem insightful. One common goal of moral theory is to provide a general account of reasons for action. These "elemental" theories seem promising. Perhaps the model can be further extended to include the whole range of

²⁵ David Miller, *Principles of Social Justice* (Cambridge, MA: Harvard University Press, 2001); David Schmitz, *Elements of Justice* (Cambridge: Cambridge University Press, 2006).

²⁶ Michael Walzer, *Spheres of Justice* (New York: Basic Books, 1981).

ethics. The analog of chemical groups might be types of relationships (relationship with children, with nature, with lovers, with self, etc...), while the periods might correspond to types of situations. Whether such a project could be made to work remains to be discovered. It depends on the success of a generalist moral theory.

The opposition between atomism and holism that the particularist proposes is a mistake. In chemistry, there are atoms, and the properties of atoms depend on context. Some properties (such as atomic number) are stable across contexts, others less so, but a general characterization can be given. Atoms divide rather neatly (but not perfectly) into types of elements. Through experiments, we can uncover their natures, their potentials to combine. Our theory of atoms is both atomistic and holistic in various respects. Why think that a theory of reasons will be any different, given the similarities we have already seen?

Of course, nature as described by chemistry is not perfectly ordered, and neither is the quantum mechanics that can be used to provide a more fundamental explanation of atomic behavior. The natural world is dappled. Still, if moral reasons could be as orderly as chemistry, then what exactly is the particularist supposed to have shown us about reasons?

The oxidation numbers described in chemical theory do not work under extreme conditions. Our simple picture of bonds and chemical states changes when considering the insides of stars, super-compressed solids, or the behavior of atoms approaching absolute zero. In some ways, the order breaks down at the margins. Similarly, the particularist might argue that a periodic table of reasons would collapse at the margins.

He might say that morality goes across possible worlds, and the apparent orderliness of morality is destroyed when we consider possible worlds consisting of extreme conditions. In response, we note that there are possible worlds consisting of extreme conditions from the standpoint of chemistry. The breakdown of theories in such worlds by itself does not defeat our scientific theories. (It is not because scientists don't consider such worlds. Theoretical physics often employs elaborate thought experiments.) Why would it defeat our moral theories?

The best way to convince the particularist to accept at least patchy generalism is to produce a good moral theory of this form. However, more can be said on behalf of this weakly ordered moral cosmos.

5. SINGULARITIES IN THE MORAL FABRIC

We might talk of the "moral fabric", just as physics speaks of the "fabric" of space-time. Space-time is a cosmos, a collection of events that are unified in some way. Einstein's general relativity is a theory of gravitation that can be used to explain the fair orderliness of space-time.²⁷ Mass tells space-time how to curve and space-time tells mass how to move. This statement needs qualification, but part of it should be taken quite literally. The relationship between the whole (space-time) and the particular (individual masses) gives neither priority in explanation.

²⁷ Robert Geach, *General Relativity from A to B* (Chicago: University of Chicago Press, 1978), p. ix.

Einstein's field equations describe the general ordering of space-time. They are complicated and cannot be used for direct calculations. Also, they are largely devoid of content; they need empirical information to fill them out. By themselves, the field equations cannot tell you much about the physics of your office. The equations are consistent with worlds radically different from ours, such as Gödel's universe.²⁸

In practice, scientists employ solutions to the equations that rely on simplifying assumptions about the universe. For example, the Schwarzschild solution describes the space outside a perfectly spherical object as static, stable, and isotropic—very loosely speaking, empty. The Schwarzschild solution is used for many calculations, e.g., in most of the experiments supporting general relativity. While it is useful, e.g., for understanding the anomalies in Mercury's orbit, obviously the Sun is not perfectly spherical, nor is space static, stable, and isotropic as described by the solution. The Robertson-Walker metric is another exact solution to the equations, but it too relies upon simplifying assumptions about the content of space-time.

Scientists first pondered the possibility of black holes by reflecting (via thought experiments) about certain implications of the Schwarzschild solution.²⁹ "Inside" black holes is a singularity. The singularity is not a region of space-time; properly speaking, it

²⁸ Kurt Gödel, "A Remark about the Relationship between the Theory of General Relativity and Idealistic Philosophy", *Collected Works: Publications 1948-1974* (Oxford: Oxford University Press, [1949] 2001), pp. 202-207.

²⁹ Namely, what would happen if the mass of a gravitating region is very dense and condensed within the Schwarzschild region.

is not a place or thing. It is more of a boundary, though one described as having the power to destroy objects approaching it. It is common to speak of space-time as "breaking down" at the singularity. This is not quite appropriate. Regardless, something in the theory breaks down there.

However, this is not a mark against general relativity. It is not as though we bumped into black holes and then had to rush to explain them. Rather, the existence of black holes was a prediction of the theory, and only recently is there empirical evidence of their existence. One might instead say that such breakdowns count in favor of the theory.

What if something similar occurs in moral theory? Consider a view I find attractive. Rossian pluralism proposes a heap of pro tanto duties, principles that in isolation work quite well, but can conflict. E.g., a principle of need for small children seems right, but the principle's explanatory power derives in part from the simplifying assumption of a relationship between parent and child in isolation. In real life, the world is messier than such a principle indicates, but then, the world is messier than what the Schwarzschild solution indicates. Robert Audi suggests that we connect the Rossian heap via the Categorical Imperative. What Ross' duties have in common is that they are ways of respecting humanity. This proposal has the strength of giving content to the (somewhat empty) Humanity Formulation of the Categorical Imperative while giving unity to Rossian pluralism.³⁰ Similarly, the somewhat empty field equations gain salience

³⁰ Robert Audi, *The Good in the Right* (Princeton: Princeton University Press, 2004), chp.

3. This proposal is of course strengthened by the fact that Kant himself, in his

and content through useful solutions describing ideal, isolated, but empirically useful models. As for the idea of a moral fabric—the whole of morality shapes the particular but is shaped by the particular as well.

Are there moral singularities? Perhaps particular principles and/or grand unifying principles like the Categorical Imperative collapse at various points (e.g., in cases of radical conflict or when the universe is at stake.) Previously, I discussed how the existence of anomalies is not necessarily a threat to moral theories or principles. More strongly, if a moral theory predicts a collapse of the principles, much as general relativity predicts the collapse represented by singularity, this is evidence in favor of the theory.

A generalist, pluralistic theory can "predict" moral singularities. Suppose an ethical theory claims that morality comes from the real problems real people face in living their lives, in making choices, in deciding how to relate to themselves, the world, and each other. Morality is a solution to those real problems, but the problems are complicated. Sometimes, unfortunately, the problems have no solution. In cases such as tragic dilemmas, morality breaks down. So much the better for the theory that says so. Miller, Schmitz, Walzer, and other theories have this prediction.

Here are some famous moral singularities. Sophie of *Sophie's Choice* is held in a concentration camp. She is told that both of her children will be killed, or she may choose to save one and have the other killed. She chooses to save her son, but the choice destroys her. The Speluncian Explorers are trapped inside of a cave and must practice

Metaphysics of Morals, unpacks the Humanity Formulation into something like a Rossian list of duties.

cannibalism to survive.³¹ They agree to draw straws to determine who will be eaten. The loser who draws the shortest straw reneges on his agreement. The others decide to kill and eat him anyways. Jean Valjean of *Les Miserables* is a redeemed convict running from a parole violation in a terribly unjust social system. He has adopted a new identity, becoming the mayor of Digne, France, and its principal employer through a factory he has built. He promises to rescue the believed-to-be-sick daughter, Cosette, of an unfortunate woman he has saved. However, before leaving for Cosette, he learns that a mentally disabled man has been captured and is believed to be Valjean. The man will be tried and imprisoned, perhaps executed, if Valjean does not confess. But if Valjean confesses, the townspeople, workers, and Cosette, all of whom depend on him, will be in jeopardy. I have described morality as being a solution to certain types of problems. Each of these situations is beyond morality's capacity to solve. Under incredible weight, the moral order collapses.

6. CETERIS PARIBUS PRINCIPLES

Rossian pluralism lends itself nicely to limited generalism. Ross' duties are often described as pro tanto or ceteris paribus principles. (Ross uses the phrase *prima facie* but admits it is a bad terminological choice.³²) The duties of benevolence, non-maleficence, self-perfection, fidelity, etc., are defeasible. They hold other things being equal. Many

³¹ Lon L. Fuller, "The Case of the Speluncian Explorers", *Harvard Law Review* (1949) 62: 615-645.

³² W. D. Ross, *The Right and the Good* (Indianapolis: Hackett, 1988 [1930]), p. 20.

moral theorists dislike ceteris paribus clauses. Particularists often argue that recourse to ceteris paribus clauses by generalists betrays the generalist program. In the philosophy of science, some have claimed that ceteris paribus laws are not real laws.³³ This is worrisome, because economics, our best social science, comprises many ceteris paribus laws. Very plausibly, *all* of our sciences, including physics, comprise ceteris paribus laws. Cartwright defends ceteris paribus laws. Her reasons are general and apply here as well.

Consider a ceteris paribus claim an oncologist would make:

(C) Ceteris paribus, smoking causes lung cancer.³⁴

Cartwright argues that utterance *C* expresses *C'*:

(C') Smoking has the capacity to cause lung cancer.³⁵

She has argued at length that the laws of physics, e.g., the special force laws, are best understood as descriptions of capacities and powers; otherwise we can attach no sense to many of their claims.³⁶ We cannot go into the details of her argument for the ubiquity of

³³ Jim Woodward, "There Is No Such Thing as a *Ceteris Paribus* Law", *Erkenntnis* (2002) 57: 303-328.

³⁴ Nancy Cartwright, "In Favor of Laws that Are Not so *Ceteris Paribus* after all", *Erkenntnis* (2002) 57: 425-439, here p. 430.

³⁵ *Ibid.*

³⁶ *Ibid.*, *The Dappled World*, and also Nancy Cartwright, *Nature's Capacities and their Measurement* (Oxford: Oxford University Press, 1994).

such laws. If she is right, this counts in favor of limited generalism over particularism, since it allows (in principle) Rossian ethics to be as law-like as physics.

What concerns me here, though, is the issue of verifying *ceteris paribus* principles. One might claim that they are inherently untestable. E.g., you assert a moral principle: "Ceteris paribus, one should be impartial between everyone". I respond, "But what about children? You ought to love your children more than strangers." You reply, "Yes, but that does not count against my principle, because I had a *ceteris paribus* clause." It seems you could make this response for every criticism I have of your alleged *ceteris paribus* principle. The worry becomes that the principle is so defeasible that it does not assert anything. If we were Popperians, we would say it is unfalsifiable.

Cartwright argues that one tests a *ceteris paribus* physical law the same way one would test that law without the *ceteris paribus* clause. We test the unconditional $F = ma$ by arranging an experiment under which only the force we wish to test is operating. We control for interfering forces that would correspond to different force functions. We might purify the experiment by preventing wind or light from disrupting it. In the case of Michelson, horses trotting on the streets outside his Cleveland laboratory introduced a new force (by slightly shaking his building) that disrupted his experiments.³⁷ As Hacking has described, to experiment is to "create, produce, refine and stabilize phenomena".³⁸ Most experiments fail because constructing pure environments takes great skill. After we have produced an experimental system with appropriate quantities and appropriate

³⁷ Hacking, *Representing and Intervening*, p. 257.

³⁸ Hacking, *Representing and Intervening*, p. 230.

shielding from interference, we run it and see what happens. This begins the testing of a law.

What changes must be made to test a *ceteris paribus* version of a law, such as "*ceteris paribus*, $F = ma$ "? None. The apparatus described above already accounts for the *ceteris paribus* clause. To produce any successful experiment, we have to shield from interfering forces. The interesting question is not whether such shielding must occur. All experiments require it.

The interesting question is how we know from what to shield.³⁹ This knowledge requires experience, experimental practice, theory, and luck. We cannot provide a systematic account. However, we can say more. With many *ceteris paribus* laws, part of what gives them content is that the *ceteris paribus* clause selects only a small class of (common) interferers, or that they hold not merely *ceteris paribus* but quite generally. With "Aspirins relieve headaches, if nothing interferes", where we have no idea *ex ante* what the interfering forces might be, we test via random samples to find statistically significant results.⁴⁰

These points apply to *ceteris paribus* principles in ethical theory as well. One way of providing evidence for a principle is to produce a thought experiment under which the principle is the only explanation. Consider DROWNING CHILD:

³⁹ Cartwright, "In Favor of Laws that Are not so *Ceteris Paribus* after all", pp. 432-434.

⁴⁰ Cartwright, "In Favor of Laws that Are not so *Ceteris Paribus* after all", p. 435.

DROWNING CHILD: On your way to the mall, you encounter a child drowning in a pool. It would take almost no effort or time to save the child. If you do not save the child, she will die.

This thought experiment provokes a common intuition: you *must* save the child. Can we discover a principle from it? Since we have been careful to exclude possible defeaters, we can say that *ceteris paribus*, you ought to save drowning children. Perhaps we can say *ceteris paribus*, you ought to act beneficently. Perhaps we can say something stronger. The interesting question becomes just how general and strong a principle we can derive from this thought experiment. Does it hold in only this one case? The particularist might say that with thought experiments like DROWNING CHILD, we need to fill in the details to determine whether it is right to save the child. Yet, the interesting thing is that intuitions start jumping out at us immediately. Thought experiments are themselves *types*, not tokens, and they generate *ceteris paribus* intuitions.⁴¹ By adding more (unusual) detail to DROWNING CHILD, it would be possible to generate the contrary intuition that one ought not to save the child. As a type, DROWNING CHILD is covered by a principle, but additional details can defeat the applicability of that principle.

To test the generality of *ceteris paribus* moral principles, we might do something like a mental random sampling of their applicability in our lives. Ross' principles pass this test. When we consider at random the types of situations we normally encounter, it seems that making a promise is normally a reason for keeping it, that having harmed another is normally a reason for making reparations, that having benefited from another's

⁴¹ Thanks to David Schmdtz for this point.

actions is normally a reason for gratitude, etc. Of course, there are anomalies, singularities, defeaters, gaping holes not covered by such reasons, and sometimes such reasons are impolite enough to switch valence on us! So what?

7. HOW GENERALLY DOES PARTICULARISM APPLY?

Suppose we construct a periodic table of reasons for a Rossian-style pluralist theory. However, suppose also that the table we thus produce in our moral theory provides only 95%, 90%, or even 50% explanatory coverage of the moral life. That is, our characterization of reasons holds most of the time, but not all the time. E.g., our theory might say that feature *F* has positive moral valence, but, in fact, there are many situations in which *F* has negative or no moral valence. Our theory has limits. The remainder of the moral life, and the reasons the moral person considers, outstrips what moral theory can explain. Is this a victory for particularism? It seems to me to be a victory for generalism. 100% explanatory coverage would be a clear victory for the generalist who sees morality as sharply, elegantly unified. 0% would be a clear victory for the particularist. Most of the middle ground is a victory for the patchy generalism I have been advocating.

Maybe particularism is the best account of reasons in those places where our theory does not work. Our best physical theories do not provide a good model for a dollar bill blowing in the wind and thus do not allow us to predict where the bill will land. At best, we say that gravity and other forces affecting it add up somehow or other.

Similarly, there might be cases of ethical appraisal where at best we say that the morally relevant features of the situation add up somehow or other.

Yet, this is not always the case. Wide swaths of morality are neat and orderly. I agree with the particularist that reasons exhibit incredible variety. Some collections of reasons fit nicely into generalist moral theories. Some do not. The particularist insists that reasons are uniformly *not uniform*, such that no general theory of reasons will do. The generalist's reply to the particularist is that reasons are not uniform in their degree of non-uniformity. By analogy, let us grant "fashion particularism": every person has his own personal style of dress. That is true, but most teenagers shop at the Gap and most retirees do not. Similarly, let us grant the bromide that each person is special and unique. Yet, most prefer pay raises to pay cuts. Our degree of individuality is not something we all have in common.

If there is any truth that particularism is meant to track, it is that reasons exhibit incredible variety. Particularism is correct that general truths do not fully explain or account for the particulars in every case. Where particularism goes wrong is by making itself too general. Reasons exhibit too much variety for particularism to be true of most of them.

Good theories have principles with actual counterexamples, not just potential ones. Theories comprise explanatory generalities, but where things do not meet the norm is interesting. The problem with particularism is that by saying that every case is unique, it fails to explain why some cases are morally interesting and others are not. Sometimes lying is morally acceptable, but *that is interesting!* When the murderer comes to the

door, perhaps I have not even a prima facie obligation to tell him I am hiding his sought victims. This is a counterexample to a generalist theory claiming I always have a prima facie obligation to be honest. However, this theory retains an advantage over its particularist competitor. The generalist theory can claim that something unusual happens in the murderer at the door case. A particularist theory cannot.

8. CONCLUSION

Morality appears to admit of enough generality to do traditional moral theory. Moral principles have their place, and are in no worse a situation vis-à-vis exceptions to the rule than many scientific laws. The particularists' arguments do not favor particularism uniquely. Perhaps morality admits of greater unity than I have admitted. I think not, but explaining why is a separate project.

A particularist might accept nearly all of this and still maintain particularism. If so, so much the worse for particularism. If the particularist admits that morality is characterizable by principles of the various forms described above, with dappled, imperfect precision, then what was the fuss about in the first place?

ACKNOWLEDGEMENTS: Thanks to the Institute for Humane Studies for their gracious support in the completion of this project, and to Nigel Ashford, Chris Brown, Chris Freiman, Michael Gill, Keith Hyams, Emir Kamenica, Cara Nine, Alastair Norcross, David Schmidtz, Mark Timmons, and Matt Zwolinski, and the audience at the

American Philosophical Association, Pacific Division 2006 meeting for helpful comments and discussion.

CHAPTER SIX
PHILOSOPHICAL DISSENSUS

"Philosophy begins in wonder. It never ends."

— Robert Nozick¹

Philosophers disagree immensely, in broad, significant ways. Our current top philosophers disagree over doctrines, methods, and even the aims of philosophy. Experts in all fields disagree, but disagreement is more acute and pervasive in philosophy than in most other fields. Philosophers need to account for the fact of dissensus. In particular, they need to determine what dissensus implies about the rationality of their enterprise. As I explain below, there is a simple but powerful argument that purports to show that philosophy is an irrational activity. My goal in this essay is to defend philosophy from its own attack.

1. THE DISTURBING TRUTH

Before turning to the argument against philosophy, let us examine a disturbing truth. Imagine that a conference of prominent philosophers is held. During the discussion period after each talk, the audience expresses *complete agreement* with everything each speaker says. As professional philosophers, we would suspect that something has gone wrong in this scenario. Similarly, imagine a graduate seminar in which students agreed with all of their professor's views. We are inclined to believe

¹ Robert Nozick, *Invariances* (New York: Cambridge University Press, 2002), p. 301.

something has gone wrong. Perhaps the professor is abusing her influence. If the students genuinely agree, they must not be thinking for themselves.

Suppose we invent a machine that allows us to communicate with the philosophers of the future. We learn that in the future, all philosophical questions are regarded as resolved. Philosophers no longer do original work and no longer disagree. The purpose of the professional philosopher is not to search for answers but to remember them. We are surprised to learn this, so we ask the future philosophers what doctrines they espouse, what methods they employ, and for what theoretical goals they aim. Unfortunately, our machine breaks before we receive a response.

We know that future philosophers believe the same things, but we do not know what they believe. Given our limited information, how should we evaluate this future? Perhaps the next hundred years of philosophy will be wildly successful. All questions will be answered in a fully satisfactory way. The arguments for future doctrines will be sound and irresistible. The future consensus is reasonable. Another possibility seems more likely: something terrible has happened. We expect perfect consensus to result from intellectual vice rather than virtue. Perhaps people have been brainwashed by social engineering. Perhaps critical thinking is no longer valued or practiced. Given the past twenty-five hundred years of philosophy, the best (defeasible) explanation for perfect philosophical consensus is that philosophy is practiced dishonestly.

These reflections illuminate a disturbing problem. Deep philosophical consensus appears to result from badly done philosophy. When people genuinely philosophize, they

arrive at different answers. If they were to agree, the most likely cause would be intellectual corruption.

2. IS PHILOSOPHY IRRATIONAL? THE ARGUMENT AGAINST PHILOSOPHY

Most professional philosophers are college professors who must grade student exams and essays. Many have encountered "The Rant", an argument normally presented by a disillusioned undergraduate doing poorly in class.

The Rant

All these philosophers disagree over everything. So, philosophy is just a matter of opinion. There's no truth of the matter, and there's obviously no point in debating it.

The Rant is a bad argument, as it stands. Students normally make the Rant to rationalize bad performance. However, there is a serious version of the Rant. Let us call this the Argument against Philosophy.

The Argument against Philosophy.

The goal of philosophy is to uncover certain fundamental truths. An enterprise is rational only if it achieves its aims. Radical dissensus shows that philosophical methods are imprecise and inaccurate. Philosophy leads people to accept a wide array of incompatible doctrines. Therefore, philosophy is an ineffective instrument for finding truth. It is irrational to use a known-to-be ineffective instrument.

The Argument against Philosophy does not claim that no philosophers have found the truth. If Kant's metaphysics is the truth, then philosophy has arrived at the truth. Yet, it

also has arrived at hundreds of other incompatible doctrines. If philosophy leads to the truth, it is only because it leads almost everywhere.

Here is an analogy. Suppose, thousands of people, each of whom wants to go to São Paulo, randomly board all flights departing Dallas-Fort Worth. Suppose they fill all departing seats, but are not told where they are going. Of these thousands, a few hundred in fact will land in São Paulo. But most will arrive somewhere else. Philosophy seems like this. It may bring some people to the proper destination, but it dumps most somewhere else. Actually, matters are worse than that. Travelers will know whether they have arrived in São Paulo. In philosophy's case, some may indeed arrive at truth. However, they will not have discernibly better grounds for believing this than those who falsely believe they have arrived at truth.

3. SIX INADEQUATE DEFENSES OF PHILOSOPHY

The Argument against Philosophy takes disagreement in doctrines to show that philosophical methods are bad. As stated, it is too strong. "F is a rational field only if its methods do not generate disagreement" is a bad principle of assessment. Principles of assessment are meant to sort good from bad, but every field, including physics and medicine, fails by this test. Is there some better replacement principle, e.g., "F is a rational field only if its methods generate X amount of agreement or more"? It is hard to say what value X should have. This problem weakens the original Argument against Philosophy. However, philosophical dissensus is still worrisome, even if the objection to philosophy cannot be formulated precisely. Here, though, are six inadequate defenses.

Let us start with a facile defense: The Argument against Philosophy undermines itself. The position that philosophy is irrational fails to pass self-inspection. "Philosophy is irrational" is a philosophical position. If philosophy is irrational, so is the view that philosophy is irrational. If philosophical argumentation never establishes any position, then the anti-philosophy position cannot be justified by philosophical argumentation. The Argument against Philosophy refutes the Argument against Philosophy. Yet it is embarrassing if this is the best defense philosophy has.

One could argue that science is less unified than commonly thought. Thomas Kuhn claims that the appearance of unity is largely a myth propagated by ahistorical science textbooks.² This approach may deflate science somewhat, making philosophy seem less inferior in comparison, but it is not enough to show us that philosophy is rational. After all, deflating science also improves the comparative position of astrology, phrenology, and creationism.

One could argue the degree of philosophical disagreement is overstated. Philosophers agree more than they disagree, but they emphasize their disagreements. Our natural desire for consensus drives us to focus on and remove disagreement. On the margins, the academic system appears to reward creativity and innovation. Accordingly, philosophers tend to specialize in fields where they can be most original. Yet, even if the amount of dissensus is overstated, there is not much consensus on the most important questions. Worse, the consensus is often superficial. Most philosophers are mental

² T. S. Kuhn, *The Structure of Scientific Revolutions, 3rd Edition* (Chicago: University of Chicago Press, 1996), pp. 136-138.

physicalists, but they diverge inside physicalism. Most philosophers agree that killing is wrong, but they disagree on their explanations of why.

A fourth type of defense is that offered by Wilbur Urban, former president of the American Philosophical Association. In 1925, Urban attempted to validate the rationality and progressiveness of philosophy by providing a list of its recent accomplishments.³ However, any such list will be contentious. Urban's own list, which I will not recite here, would seem problematic to most contemporary philosophers. Worse, the clearest cases of such progress are often negative. Gettier made progress by demolishing the justified true belief analysis of knowledge. Quine, Putnam, and others eradicated logical positivism. Refuting inadequate past theories clears the path for good answers, but does not thereby give us good answers. Even negative "progress" tends to be reversed, as once dead doctrines (such as Ross' moral theory) are resurrected.⁴

A fifth claim, that philosophy is an end in itself, should not sate us. Perhaps philosophy is worthwhile in itself, but we hope that it is valuable as a method for finding truth.

A sixth way defense would concede that current dissensus shows that philosophical methods are ineffective, but then assert that philosophy could become effective in the future. The reason there is dissensus now is that most or all philosophers fail to use the right methods. We just need to continue working until discover what the right

³ Wilbur M. Urban, "Progress in Philosophy in the Last Quarter Century", *The Philosophical Review* 35:2 (1926), pp. 93-123.

⁴ E.g., Robert Audi, *The Good in the Right*, (Princeton: Princeton University Press, 2004).

methods are. Once we find them, agreement will follow. Indeed, we could even take agreement as a sign that we have discovered the right methods.

However, we do not really expect this. We are more inclined to think disagreement is a permanent fixture. As noted previously, the disturbing truth is that widespread philosophical consensus is more likely to come from irrationality and intellectual corruption than from honest inquiry.

4. WHY DISSENSUS OCCURS

To assess whether dissensus really is evidence that philosophy is irrational, we should determine what produces dissensus. In *The Strife of Systems*, Nicholas Rescher claims "aporetic clusters" are the mechanism for disagreement.⁵ An aporetic cluster is a set of related theses, each of which we have good reason to accept, but when taken together are mutually incompatible. For example, take the following aporetic cluster:

1. All human acts are causally determined.
2. Man can and do make free acts of choice.
3. A genuinely free act cannot be causally determined (for if it is so determined then the act is not free by virtue of this very fact).⁶

There are good reasons to accept each of 1, 2, and 3. But one cannot consistently accept all three. Thus, the common reaction is to accept two of the three premises and reject the

⁵ Nicholas Rescher, *The Strife of Systems* (Pittsburgh: University of Pittsburgh Press, 1985), p. 25.

⁶ Rescher, *The Strife of Systems*, p. 27.

third. Metaphysical libertarians deny 1 but accept 2 and 3, hard determinists deny 2 but accept 1 and 3, while soft determinists deny 3 but accept 1 and 2. With many examples, Rescher argues that most philosophical debates fit the model of the aporetic cluster. Philosophical problems arise because commonsense fundamental beliefs are inconsistent.

We remove the inconsistency by rejecting one of the theses in the aporetic cluster. However, the quandary is determining which to reject. Rescher argues that philosophers choose to reject different theses—and thus establish conflicting schools of thought—because they accept different *cognitive values* or weigh the cognitive values differently.⁷ Cognitive values are the epistemic traits by which we assess a thesis, e.g., coherence, plausibility, generality, importance, informativeness, elegance, etc. A philosopher who more strongly values plausibility and intuitiveness is likely to accept different doctrines from a philosopher who more strongly values systematicity. Rescher argues that differences over the relative weights of cognitive values can be only partially resolved.

Rescher's theory of dissensus is richer than what I have presented here. His explanation of why dissensus occurs and his defense of philosophy are illuminating. However, in this essay, I will take a different approach. To simplify (perhaps oversimplify) Rescher's view: dissensus ultimately springs from irresolvable differences in philosophers' cognitive values. However, I hold that the dissensus occurs *even when* philosophers accept the same cognitive values to the same degree.

In my view, philosophy has two goals. Both goals are rational, but together they generate dissensus. The first goal is epistemic: to illuminate truth. The second goal is

⁷ Rescher, *The Strife of Systems*, pp. 95-115.

Socratic: to ask why. Philosophy seeks to inquire about and often challenge assumptions and fundamentals, including its own. Philosophy questions what is normally taken for granted. Philosophy asks the fundamental questions. Yet, once one tries to answer these fundamental questions, one generates new questions. At times, these two aims synergistically cohere. Inquiring into and challenging assumptions is sometimes the path to truth. More often, challenging assumptions results in disagreement and uncertainty. Philosophy is marked by flux, but to know a body of doctrine is true is to be at rest. Accordingly, philosophy is characterized by competing aspects. On one hand, it seeks to stop at truth; on the other, it seeks never to stop.

One might dispute whether questioning assumptions is a goal of philosophy or just a method for seeking truth. Partly, questioning assumptions is an instrument for building doctrines. However, Western philosophy models itself after Socrates, who questioned for its own sake. We might say that this desire to question is itself a constituent of philosophical wisdom. Determining whether the questioning of assumptions and fundamentals is a separate goal or merely a method is not important. The point is that it is integral to philosophy.

5. PHILOSOPHY AS GADFLY

Socrates described himself as a gadfly stirring up a noble horse—Athens.⁸ My view is that philosophy is its own gadfly, and this is why dissensus occurs.

⁸ Plato, *Apology*, trans. G. M. A. Grube (Hackett: Indianapolis, 1981), 30e.

Socrates was executed partly because he made proud men feel stupid. In Plato's *Apology*, Socrates claims that esteemed authorities could not answer the most fundamental questions about their areas of expertise. For instance, no statesman could explain the nature of justice. They had firm opinions about justice, but Socrates queried the basis of those opinions. He uncovered assumptions and questioned them. We inherit this tradition from Socrates; this is part of what makes us philosophers.

The philosopher is separated from the layperson not for having beliefs about the fundamental issues in human existence. Most people have such beliefs, e.g., via religion. One does not become a philosopher merely by reflecting on those issues. Almost everyone reflects on occasion. Rather, one becomes a philosopher in part by one's willingness to challenge common sense, to engage in critical inquiry into these fundamental questions with some readiness to disclaim even one's most sacred beliefs.⁹ The philosopher need not repudiate common sense or what she holds sacred. Her inquiries may substantiate her pre-critical views. Yet, she must be willing to submit these views to testing. The philosopher does not take things for granted. She asks why and how does one know.

Compare this to science. Kuhn is roughly correct about scientific progress. Scientific revolutions establish new paradigms. A paradigm is a generally accepted conceptual framework and model of a physical domain. Once a paradigm is established, scientists solve puzzles suggested by the paradigm, fill in gaps, gather better

⁹ Jeremy Waldron makes a similar claim in his "What Would Plato Allow?", *Nomos* XXXVII (1995), pp. 138-178, here pp. 170-171.

measurements and constants for the equations important to the paradigm, and so on. This "normal science" does not challenge the basic paradigm. Instead, normal science stretches the paradigm as far as it can go. When a paradigm is threatened by significant anomalies or a better competitor, it is sometimes replaced. Kuhn makes too sharp a distinction between normal and revolutionary science; there are many intermediaries. He fails to provide a good account of change. Yet, he is correct that scientists tend to seek and find consensus on a basic theory. Scientists search for assumptions to hold. Scientists are unsettled by revolutionary science and seek to re-enter periods of normal science. Science is a unity-seeking discipline.

In contrast, philosophers seem unsettled by the settled. They prefer revolutionary philosophy to normal philosophy. Of course, many do adopt and defend joint research programs and general theories. Rawlsian political theory is one recent example. But theories rarely enjoy the level of acceptance found in natural science. While some philosophers attempt to stretch a mutually accepted theory as far as it will go, there are always other philosophers seeking to refute it. Common acceptance of a doctrine is a red flag for other philosophers. Such consensus normally signifies the existence of unquestioned (and questionable) assumptions.

Randall Collins, a sociologist who studies intellectual change, argues that disagreement is selected for via the academic process.¹⁰ In effect, disagreement is an adaptation to a limited attention space. Intellectuals seek attention from other

¹⁰ See Randall Collins, *The Sociology of Philosophies*, (Cambridge, MA: Harvard University Press, 1998), pp. 38-39.

intellectuals, but there is not enough attention to satisfy everyone. One good strategy for receiving attention is to be the first to quarrel with others over some issue. Accordingly, a philosopher is likely to gain prestige by fighting against those doctrines others have not yet thought to challenge.

A paradigm is a set of assumptions. Philosophical method typically avoids settling on assumptions. Thus, philosophical paradigms are unstable. For the purposes of a given paper, a philosopher begins with attractive premises and attempts to support a conclusion. However, she recognizes that the premises themselves are subject to eventual controversy. Even universally accepted premises must be examined in due course. Even if premises seem self-evident, the philosopher must ask what to make of self-evidence.

Additionally, that a philosopher uses one form of argument rather than another is an assumption requiring eventual defense. For instance, many critics argued that Rawls' game-theoretic reasoning as such distorts human nature.¹¹ Such critics do not deny that Rawls' conclusions follow from his premises. Instead, they argue his argument *style* is based on false assumptions. Philosophical method is as controversial as doctrine. Even aims are contentious. Analytic philosophers typically claim to aim for *truth*, but the intelligibility and value of this goal is challenged by many in the pragmatic tradition.

¹¹ E.g., Michael Sandel, *Liberalism and the Limits of Justice*, 2nd Edition (New York: Cambridge University Press, 1998), chapter 1.

More strongly, using arguments to establish a view itself rests on methodological assumptions. Nietzsche and postmodernists challenge mainstream philosophy by eschewing argumentation and presenting evidence for their views.

Any doctrine, method, or aim is based on assumptions. With few exceptions, every philosophical doctrine bears some relationship to every other, be it implication, contradiction, tension, or support. The relationship between distinct aims, methods, and doctrines is similar. No strand in the philosophical web of aims, methods, and doctrines is perfectly isolated. Nothing in philosophy is an island unto itself. As such, every strand is subject to questioning from a thousand angles.

With only some exaggeration, every bit of human knowledge is in some way relevant to every philosophical doctrine. Yet, to have a perfectly coherent philosophy that does justice to all the data is probably beyond human ability. Every strand can be questioned from a thousand angles, but only a superhuman could envision all the angles. One goal of specialization in analytic philosophy was to produce systematic coherence through a division of intellectual labor. Such specialization has failed, as it has resulted in sophisticated particular doctrines at the expense of systematicity. In contrast, great systems of philosophy tend to be weak in their details.

Most likely, there will never be a perfect philosophy worthy of universal acceptance. No philosopher has the ability to check all of her premises and assumptions, but philosophy demands that all assumptions be checked. That is the problem with philosophy. The questions do not stop. Philosophy ends only when there are no pointless questions, but philosophy recognizes few questions as pointless. Philosophers

are as much problem-discoverers as problem-solvers. As Collins illustrates with numerous historical examples, very often, the unsolved problem is introduced by the creative philosopher, as the issues he discusses are non-issues until the philosopher discusses them.¹² Philosophical creativity consists mostly of identifying an unsolved problem and then providing a solution.

6. THE BURDENS OF JUDGMENT

Why should asking why lead to dissensus? The data underdetermine theory choice. (This assertion is rendered tendentious by the fact that no philosopher has produced a good account of theory, data, or explanation.¹³) As Mill said, "Very few facts are able to tell their own story."¹⁴ In accepting a theory, I have to make choices. Not only do the data (intuitions, experiences, conclusions of non-philosophical theories, observations) underdetermine a philosopher's theory choice, but no two philosophers possess exactly the same data. Scientists work with more constrained data sets. Most data are irrelevant to most scientific theories; scientific theories have limited domains. Philosophy aspires to be a total theory of the fundamental. Thus, *everything* is data.

John Rawls argues that modern democratic societies are marked by "the fact of reasonable pluralism", i.e., by unavoidable, reasonable diversity in the religious,

¹² Collins, *The Sociology of Philosophies*, p. 80.

¹³ Cf. Nozick, *Invariances*, p. 111.

¹⁴ J. S. Mill, *On Liberty*, (Indianapolis: Hackett, 1978), p. 19.

philosophical, and moral doctrines citizens accept.¹⁵ Rawls' analysis helps explain why philosophical disagreement will be a permanent fixture when there is free philosophical inquiry. Not all disagreement is the result of irrationality. Reasonable disagreement occurs because of the "the burdens of judgment".¹⁶ Rawls claims that judgment is burdensome because:

1. Evidence is difficult to assess,
2. The evidential weight of various pieces of evidence is itself subject to debate,
3. Our concepts are indeterminate to some degree, and there is reasonable disagreement on interpretation within a certain range,
4. Our different life experiences change the way we perceive evidence.¹⁷

To this I add:

5. Philosophy is subject to the principle of universal relevance: everything is data, but no one has access to all the data.
6. Human psychology is such that consensus and certainty are often evidence of thoughtlessness.
7. Confirmation bias means that the weight we attach to bits evidence is affected by the order in which we receive it.¹⁸

¹⁵ John Rawls, *Political Liberalism* (New York: Columbia University Press, 1996), pp. Xx, 36-37.

¹⁶ Rawls, *Political Liberalism*, p. 55.

¹⁷ Rawls, *Political Liberalism*, p. 56-57.

These factors need not make agreement impossible. Philosophers often do convince one another of certain views. Genuine debate would be impossible if we had no common presuppositions, meanings, and so on. Vagueness exists, but it has its limits. Still, we should not generally expect philosophical disputes to end in agreement, and this lack of agreement need not be irrational. Choosing an outdated car that does not meet one's needs is irrational—the same goes for philosophy. But there are many reasonable choices.

One might object by claiming that the point of philosophy is to eliminate these burdens. If evidence is difficult to assess, philosophy's job is to explain how to assess it. However, any explanation of the proper way to assess evidence will itself be defended by difficult to assess evidence.

Similarly, philosophers are experts at identifying and lessening vagueness. If a philosopher is asked what the meaning of life is, he typically will respond by translating the question into a number of clearer sub-questions.¹⁹ (For instance, is one asking about the significance of life? If so, significance to whom?) However, even if vagueness can be significantly reduced, e.g., by replacing less determinate with more determinate concepts, there are no perfectly determinate concepts. Any attempt to construct an artificial fully determinate language will fail, because we learn such languages by

¹⁸ Thanks to David Schmidtz for this point.

¹⁹ See David Schmidtz, "The Meanings of Life", in David Schmidtz (ed.), *Robert Nozick* (New York: Cambridge University Press, 2002), pp. 199-216, esp. pp. 201-202.

translating them into our natural language. The indeterminacy of natural language infects any artificial language. At the limit, indeterminacy remains.

Lastly, to some degree philosophy seeks to transcend personal experience. Randall Collins describes philosophers as attempting to "produce decontextualized ideas," ideas that are true "apart from any locality, anyone concretely putting them into practice, ...and apart from whoever believes [them]." ²⁰ We are situated in particular bodies and life experiences, but we do not inhabit distinct worlds. Some critics of philosophy claim that our different experiences make any common truth impossible. (Since this view is self-refuting, typically it is gestured towards but not explicitly stated.) However, our unique experiences imply not merely that we have unique evidence, but also unique illusions. Accordingly, while different experiences make reasonable disagreement possible, not all disagreement is reasonable. Part of the point of philosophical conversation is to see through one's own illusions.

7. WHY ASK WHY?

Philosophers are gadflies to each other, to non-philosophers and even to themselves. In part, the rationality of philosophy will depend on the rationality of acting as a gadfly. So is being a gadfly—always being prepared to ask why—rational? The burdens of philosophical judgment make both consensus and justified certainty highly improbable. Why philosophize, then? Why have a philosophy at all?

²⁰ Randall Collins, *The Sociology of Philosophies*, p. 19.

This question presupposes we have a choice in having a philosophy. But we do not. Every normal adult has philosophical beliefs. I would not claim that they have *theories*, but they do possess less robust metaphysical, epistemological, and moral views. Such people may never have reflected on their metaphysical beliefs, and they may have difficulty articulating the beliefs, but they still have them. Everyone has implicit views about the fundamental nature of reality, of the possibility and nature of knowledge, of issues of right and wrong. Human beings cannot help having such beliefs.

Normal human cognition results in everyone having some sort of philosophy. This fact gives us an answer to the first question, why philosophize? We do not have a choice in whether we will have a philosophy or not. But we can choose to philosophize. We can choose to reflect, to do our best to integrate our ideas into a coherent system, to weigh our philosophy against the evidence, and to formulate reasonable abstract principles. Or, we can choose to have our philosophies in an unphilosophical manner, without thinking.

The burdens of judgment also exist for lifestyle choice. One might ask, why bother have a lifestyle? But we do not have a choice about having a lifestyle. Given that we must live some way or other, we should seek to live reasonably, even if making a perfect choice is beyond our grasp.

One's philosophy is one's primary intellectual orientation toward reality. Since one is forced by nature to have a philosophy, it is often rational to pursue the best philosophy one can. Of course, there are costs to this pursuit. Philosophizing takes time away from being with one's family, from sleeping, and from playing guitar. How good a

philosophy one should strive for will vary between people. The opportunity costs and expected returns of pursuing philosophy are different for the philosopher and the layperson.

There can also be philosophical costs to philosophizing. Pursuing the best philosophy need not always result in a better philosophy. Some philosophers should have remained unreflective laypeople. Sometimes philosophers convince themselves of dumb views for dumb reasons. Accordingly, choosing to philosophize is a bet, not a sure thing. However, most philosophers have better reasons for their views than laypeople have for theirs.

According to Rescher, the anti-philosophy skeptic advises us to forget "abstruse theoretical issues" and "focus on [our] practical needs".²¹ Rescher responds by arguing that philosophy is a form of "intellectual accommodation" into the world. Just as we need to physical comfort and security, we need intellectual comfort and security. The need for a good philosophy is like the need for a well-built house, which can withstand bad weather.

Rescher's analogy extends further than he himself may have realized. I would add two points. First, people living in the same city, seeking to build houses that can withstand bad weather, build different houses. The way they solve the problem of physical security and shelter differs, and rationally so—even when they face exactly the same conditions. This is equally true for philosophy. Second, no house can withstand every possible disaster. Every house is vulnerable and can be destroyed. Still, even if no

²¹ Rescher, *The Strife of Systems*, p. 249.

house is perfect, there are more and less vulnerable houses. It is better to pursue the safer houses. The same can be said for philosophies.

8. THE VALUE OF DOUBT

The Socratic method of questioning assumptions inspires doubt about everything. It inspires doubt about philosophy, other areas of human knowledge, and itself. Philosophical speculations can be alienating, generating "melancholy and delirium", until one leaves the study and natural common sense reasserts itself.²² David Hume's reflections did not lead him to dogmatic skepticism, but he does admonish us against claiming that anything is certain, evident, or undeniable.²³ Philosophy can provide reasons for and against views, but it subverts certainty. In the conquest of doubt, philosophy is an enemy.

To be clear, radical dissensus is compatible with the psychological feeling of certainty. A world marred by religious warfare between thousands of conflicting sects could be inhabited entirely by true believers. Still, an uncommitted person who saw that there were thousands of sects would not feel confident about any religion. Similarly, the dissensus in philosophy does not entail that all philosophers have doubts. They could be true believers. However, most are not. A philosophy is not a faith.

²² David Hume, *A Treatise of Human Nature*, (eds.) David Fate Norton and Mary J. Norton (New York: Oxford University Press, 2000), 1.4.7.9.

²³ Hume, *A Treatise of Human Nature*, 1.4.7.15.

If doubt is a bad thing, and philosophy produces doubt, then philosophy would seem to be a bad thing. Consider this view from early twentieth century philosopher Walter Pitkin: "A doubt regarding any matter carries with it (not in mere logic, but in real life) the desirability of the doubt's annihilation."²⁴ Pitkin takes this statement to be incapable of, but not requiring, proof. He compares doubt to pain. Anyone who has experienced pain wants to be rid of it. The only reason to accept pain or doubt is that doing so prevents some greater evil or produces some greater good.

Human beings have limited, imperfect cognitive abilities. Recognition of this fact shows us that we are not entitled to feel certain about very much. A degree of doubt shows respect for our capacities.

Philosophy, the gadfly, is valuable in part as an instrument for producing doubt. Hume argued that a vice is a character trait that elicits disapprobation. On that view, it would seem the disposition to philosophy is a vice. Philosophical reasoning often invites disapproval. Socrates was hated by much of Athens and accordingly executed. Hume himself, though partly in jest, describes his own skeptical thinking as "feeding his despair"²⁵ and "having expos'd [him] to the enmity of all metaphysicians, logicians, mathematicians, and...theologians".²⁶ Philosophy is often viewed as offensive, since it

²⁴ Walter B. Pitkin, "Is Agreement Desirable?", *Journal of Philosophy, Psychology, and Scientific Method* 9:26 (1912), 711-715, here p. 713.

²⁵ Hume, *A Treatise of Human Nature*, 1.4.7.1.

²⁶ Hume, *A Treatise of Human Nature*, 1.4.7.2.

denies uncritical acceptance of sacred convictions. Additionally, philosophy can produce absurd doctrines, doctrines at least as absurd as any vulgar superstition.

Despite this, Hume vindicates philosophy. One of philosophy's chief uses is combating superstition. Superstition is "more bold" than philosophy and "arises naturally and easily", ensnaring people's minds.²⁷

Man, it is true, can, by combination, surmount all his *real* enemies...But does he not immediately raise up to himself imaginary enemies, the demons of his fancy, who haunt him with superstitious terrors and blast every enjoyment of life? His pleasure, as he imagines, becomes in their eyes a crime...Nor does the wolf molest more the timid flock than superstition does the anxious breast of wretched mortals.²⁸

Men terrorize themselves with superstitions. They devise the "monkish virtues" in response to these terrors, denying themselves pleasure, and instead engage in asceticism.²⁹ Moreover, though humanity is tormented by disease, famine, and war, superstition causes people to fear the afterlife. They worry that death, rather than offering respite from pain, will offer them eternal hell instead.

²⁷ Hume, *A Treatise of Human Nature*, 1.4.7.13.

²⁸ David Hume, *Dialogues Concerning Natural Religion*, (Indianapolis: Hackett, 1980), p. 60.

²⁹ David Hume, *Enquiries Concerning the Principles of Morals*, (eds.) L. A. Selby-Bigge and P. H. Nidditch (New York: Oxford University Press, 1975), p. 270.

Superstition does not merely conjure up imagined enemies. It also gives rise to war, intolerance, brutal oppression, and witch-burnings. Superstition seems to be one of the main problems philosophy can solve. Its ability to combat superstition outweighs its own tendency to err. Religious errors are dangerous; philosophical errors are only "ridiculous".³⁰

The philosophical mind is cautious, non-dogmatic, and probing. In contrast, certainty tends to engender obstinacy, righteousness, and intolerance. Doubt may or may not be intrinsically undesirable. Yet, people who lack doubt tend also to be unpleasant, even dangerous. Doubt is good for allaying certain human vices.

9. THE VALUE OF DISAGREEMENT

Not all agreement is rational. Cults are one of our worst epistemic practices, but we see widespread agreement inside the cults. Agreement is not sufficient for rationality. Is there reason to think it is necessary?

Agreement tends to breed psychological feelings of certainty. Opinions seem better validated when popular. Yet, we should not be certain about some issues. Accordingly, disagreement will be instrumentally valuable for restraining certainty.

Disagreement is valuable for improving theories. Times of relatively high philosophical agreement tend to be times of relatively poor philosophy. Agreement on doctrine tends to be worse than agreement on methodology, which is in turn relatively worse than agreement on philosophical aims. For instance, in the scholastic period, basic

³⁰ Hume, *A Treatise of Human Nature*, 1.4.7.13.

doctrines were not frequently challenged. The modern period of philosophy was an improvement in part because there was a freer intellectual climate. Many different doctrines appeared. With regard to method, contemporary philosophical debate is an improvement over debate fifty years ago. Mainstream academic philosophy is no longer as confined to metalinguistic study and conceptual analysis. The number of admissible argument styles and admissible issues has grown, making philosophy more interesting and relevant than in the past. (For example, consider the resurgence of normative ethics.) Lastly, unanimity in philosophical aims seems least problematic. The majority of philosophers throughout history have claimed truth as their goal. However, challenges to this position from pragmatism, postmodernism, and elsewhere, are healthy. The assumption that truth is the correct goal is as subject to philosophical questioning as everything else.

Earlier we imagined complete philosophical agreement occurring in the future. We expected this agreement to issue from intellectual vice, rather than virtue. Our reasons for such pessimism are historical. Periods of intellectual freedom produce myriad doctrines. Widespread agreement is typically the result of intolerant, stifling atmospheres. The Renaissance had greater philosophical diversity than the late Middle Ages. Even in the Modern Period, various philosophers had to flee to Holland to write freely. The appearance of unanimity is more often found in theocratic, ignorant, or totalitarian societies. Philosophical consensus is normally the sign of some problem. It is imposed from without by intolerance (as when Stalin dictates the correct interpretation of Marx) or within by intellectual laziness.

We should welcome disagreement for Millian reasons. As Mill observed, if a view is not "fully, frequently, and fearlessly discussed, it will be held as a dead dogma, not a living truth."³¹ This is a contingent fact about human nature, not about truth, argumentation, or philosophy as such. Perhaps angels or aliens could avoid dogmatism. However, when people are not exposed to a variety of views, they grow unreflective and overly deferential to authority. Unfortunately, consensus breeds intolerance. Without disagreement, truth becomes another superstition.³²

Knowing the arguments for opposing sides is not sufficient. Most intellectuals can rehearse the arguments for rival views. However, they do not feel the force of these arguments; the reasons seem silly to them. Though an intellectual may be respectful and cautious in her writings, she secretly derides the other side. On Mill's view, this is dangerous. Instead, one should throw oneself "into the mental position of those who think differently".³³ Until one truly feels the force of the other side, until one experiences opposing views as genuinely attractive, one cannot hold one's own view in a respectable way.

Mill believed that the number of opposing theories was dwindling. He thought that free discussion was driving toward consensus. Nearly one hundred fifty years later, we see that he was wrong. Discussion has been freer, but the number of opposing theories in almost every field has grown. This would surprise Mill, but in one respect, it

³¹ J.S. Mill, *On Liberty*, p. 34.

³² Mill, *On Liberty*, p. 34.

³³ Mill, *On Liberty*, p. 36.

would please him. Mill worried that future consensus would cause dogmatism and mental apathy. He thought a surrogate for genuine disagreement, such as devil's advocates, would become necessary if people were to have a "living apprehension of truth."³⁴ We might not have a living apprehension of truth, but at least we have some of its preconditions.

Philosophical disagreement is generally a good thing. There is a point to asking questions and seeking answers, even if one cannot be certain and even if one cannot convince everybody else. More importantly, disagreement helps keep us honest.

10. THE PRESUPPOSITIONS OF THE ARGUMENT AGAINST PHILOSOPHY

It is rational to be a gadfly, because it is rational to ask why, to inspire doubt, and to generate disagreement. In establishing that, we have partly defended philosophy. To review, the Argument against Philosophy is this: Philosophical methods lead everywhere, so they are not good methods for arriving at truth. Let us now turn to questioning the assumptions behind the Argument.

Harold Brown articulates what he calls the classical model of rationality. (Brown criticizes the model.) According to Brown, many philosophers implicitly (and often explicitly) assume the classical model is correct. This model has a number of properties:

The Classical Model of Rationality

1. *Universality*: When engaged in rational thinking, everyone with the same information and goals must arrive at the same answer.

³⁴ Mill, *On Liberty*, p. 42.

2. *Anti-Judgment*: Accordingly, there is no room for individual judgment. When judgment is used, the conclusion or choice is nonrational.
3. *Necessitarian*: Rational knowledge shows us necessary connections. A rational argument necessitates its conclusion.
4. *Rule-Based*: Rational thinking involves following rules, i.e., the rational rules.
5. *Algorithmic*: The rational rules must guarantee a unique result in a finite number of steps.
6. *Foundationalist*: Rational thought begins by using the right premises.³⁵

Brown claims that on the classical conception of rationality, being rational is like solving an arithmetic problem. If everyone is rational, he will get the same answer by following the same rules.

Brown argues that skepticism and relativism often derive from the assumption that rational thought must fit the classical model. Yet, most thinking does not fit. Once it is noted that most thinking is not universalistic, rule-based, algorithmic, and so on, it is a short jump to concluding that most thinking is irrational.³⁶ However, Brown argues for the opposite conclusion. Physical science is our paradigmatic rational activity. Since physical science does not fit the classical model, we should conclude not that science is irrational, but rather that the model is incorrect.

The Argument against Philosophy seems to rely upon the classical model. In effect, it notes that philosophy fails to fit the classical model and then concludes with

³⁵ Harold I. Brown, *Rationality* (New York: Routledge, 1990), pp. 6-38.

³⁶ Brown, *Rationality*, p. 130.

skepticism about philosophy. However, aside from some areas of mathematics and deductive logic, almost no thinking fits the classical model. So there is little motivation for holding that philosophy must fit, except for the hope to rationally resolve all debate.

The Argument relies on an oversimplified conception of what philosophers do. It says that each philosopher claims to have the truth, but they disagree with each other. No one has any better claim to it. But this exaggerates the claims philosophers actually make. Most philosophers have learned their lessons from Hume. They are cautious, probing, and non-dogmatic. Philosophers rarely claim, "Here is the absolute truth. I found it. I'm right." Those making such pronouncements are laughed off or ignored. We know better than that. Rather, they say, "I attempted to find the truth. Here is what I came up with. Here are my reasons." Most philosophers acknowledge they cannot solve all of the problems with their views. Thus, they will often say things like, "Here are some of the issues my theory still faces. I hope they can be resolved. But perhaps they cannot. I commit myself to this view, knowing I could be wrong, but acknowledging that I have done the best I can do."

When philosophers take this stance, they stand in different places. They adopt different theories. But there is something eminently humble about this sort of philosophical practice. The skeptic—who demands certainty and agreement—seems rather supercilious in contrast. He holds philosophers to an extravagantly high standard and disdains them for failing to meet it.

Philosophy is not exactly a science, but could it be a pseudoscience? Drawing the distinction between science and pseudoscience has been notoriously difficult. However,

there are some differences between philosophy and clear pseudosciences, such as astrology. Philosophy seeks confirmation but also seeks disconfirmation, i.e., philosophy opens itself up to challenges. It offers reasons for its views. It is open to and seeks to cohere with findings from other fields. It at least attempts to be progressive—finding ways of improving. At their best, philosophers care about evaluating their theories in contrast to competing theories. They take competitors seriously, not as opponents who must be defeated at all costs, but as fellow conversationalists who may be right. In contrast, astrology hides itself from possible disconfirmation, either by ignoring evidence or by producing doctrines so vague that nothing can count against them. Some pseudoscientists claim that their theories cannot be tested because testing interferes with the laws they posit: "I can turn invisible, but only when no one is looking." The pseudosciences ignore contrary views, or they seek to defeat them without taking them seriously. (E.g., the fundamentalist Christian does not think the Muslim might be onto something concerning God.) The pseudosciences also fail to seek coherence with other fields. Pseudoscientists, be they conspiracy theorists or alien abduction experts, offer reasons for their views, but are quick to condemn those that honestly challenge those reasons. The pseudosciences also are non-progressive. Pseudosciences typically believe they have a final, perfect answer, and need not improve. A major difference between scientists and pseudoscientists is that the latter tend to think they are infallible, while the former tend to engage in healthy self-criticism. For these reasons, philosophy does not seem to be a pseudoscience.

11. MAKING DO WITH DISSENSUS

Dissensus is a fact of philosophical life. From our reflections above, we have reason to think that honest, rational inquiry into philosophical issues will always result in some disagreement. Philosophers do not need to learn how to live with dissensus. They are already accustomed to it. However, if dissensus does not threaten philosophical rationality, it would be useful to know what philosophical rationality is.

Brown sketches an alternative model of rationality. His model is intended to describe science, but it generalizes to philosophy and most other fields. If anything, the major problem with the model is that it is incomplete.

Brown first proposes that we take the concept of a rational agent as basic.³⁷ Whereas traditional epistemology evaluates beliefs by examining their conformity to some rules or principles, Brown espouses what amounts to a more character-based approach. Good scientists exhibit the epistemic virtues, including a capacity for reflection and good judgment. (To borrow the language of virtue ethics:) a rational belief is one held by a rational agent when and because he is thinking in a characteristically rational way. To complete this sketch of a theory, one would need an account of the epistemic virtues.

The questioning of fundamentals and background assumptions leads to a regress problem. In principle, we can always ask another question. However, we avoid regress and settle (cautiously) on a doctrine, though we recognize that many questions remained unanswered. We do this by exercising good judgment, by thinking, at last, "That will do

³⁷ Brown, *Rationality*, p. 185.

for now", knowing there is more to be done. Every philosophical essay illustrates this. The author always has undefended premises, background methodological assumptions, cognitive goals, and the like. She defends some of these, but never all, and even her defense in principle needs defense. Still, at some point, a philosopher recognizes that she has done enough and stops writing.

Learning good judgment is difficult. It is partly learned through the reward and punishment scheme of peer-review publishing, comments and questions at conferences, dissertation director feedback, and the like. Philosophers, like their counterparts in other disciplines, largely share the same sense of when a paper is adequately defended. Perfectly and completely defending a paper is impossible. Yet, if perfection is unattainable, there remains a legitimate sense of good enough.

If A claims that X and B claims that not-X, one of their beliefs is true and the other is false. Both views cannot be true. However, justification works differently. Though A and B argue for contradictory positions, it is possible for both to be justified. On the classical model, this would seem impossible. On that model, both A and B are rationally obligated to follow the same algorithm on the same starting premises and arrive at the same conclusion.

Epistemologists typically hold that justified *false* beliefs are possible; one can have warrant in believing something false. Given the fact of dissensus, if philosophy is rational, justified false belief must be common. Depending on how finely we demarcate the various doctrines, probably the majority of philosophers' philosophical beliefs are justified false beliefs. Of course, this assumes that philosophers generally are justified in

holding their doctrines. This should be granted. I might think that my theory of modesty is better than some competitor's, but I acknowledge that both theories have problems. Moreover, I acknowledge that my competitor has good reasons for holding her views. Both of us are seeking the truth, and we are justified because we have good evidence for our views.

To some degree, the process of peer criticism is meant to weed out theories by presenting problems that thereby lessen one's claim to justification. Sometimes this reduces the number of justifiable theories, but just as frequently, philosophers improve their theories and restore their justification. This is one reason why philosophical theories rarely die permanent deaths.

Philosophy is rational because its aims are rational and it pursues those aims in a rational way. The Argument against Philosophy gives us reason to think that philosophers are never certain in having achieved those aims. Yet, philosopher's work could still be good enough, as good as one can reasonably expect, and good enough that it is reasonable to pursue the work rather than avoid it.

12. PASSING SELF-INSPECTION

Philosophy seems able to survive its own attack. Keeping with the spirit of the discussion above, I am not fully certain of my view, and am aware it has its faults.

I have assumed that philosophy aims to illuminate truths. This is not quite the same thing as saying that philosophy aims to produce true statements, or, more plausibly, non-trivial true descriptions of the world. It is not obvious that my account of philosophy's aim is correct. Perhaps most philosophers have aimed to illuminate the

truth, but certainly not all have. Whether this is the best characterization of philosophers' aims is contentious. One might claim instead that the goal of philosophy is simply to ask questions without seeking answers, to defend useful beliefs without reference to truth, or to find communion with some god. These aims conflict with seeking to illuminate truth. Alternatively, perhaps the goal of philosophy is to produce wisdom, to achieve enlightenment, to produce a "philosophical persona",³⁸ to produce a healthy worldview, to make moral progress, or to create good culture. The relationship between these goals and the goal of illuminating truth is unclear. Alternatively, many philosophers have held that the goal of philosophy was self-destruction. Philosophy exists solely to eliminate certain conceptual problems it itself causes; once these are eliminated, philosophy has no use. This contrasts with my view.

In this essay, I began with certain premises that seemed plausible to me. I have employed argumentative and rhetorical forms of a certain nature. I have proceeded with a particular goal in mind. Each of these is a type of assumption.

On the theory presented here, philosophers ought to challenge assumptions. Accordingly, philosophers ought to challenge this theory, including this very sentence. Thus, the doctrine presented in this paper passes self-inspection. Passing self-inspection does not show a view is good. After all, anti-induction passes. But passing is a necessary condition for a doctrine's success. A theory that refutes itself is refuted. A theory that does not refute itself may yet be refuted by something else.

³⁸ See Elijah Millgram, "How to Make Something of Yourself", in *Robert Nozick*, pp. 175-198, here p. 175.

CHAPTER SEVEN

TRUSTING OUR MORAL INTUITIONS

1. MORAL INTUITIONS HAVE CAUSES

We are not as rational as we would like to be. According to evolutionary psychologist Robert Wright, the human brain is "a machine for winning arguments", for seeking "victory, not truth".¹ It aims to convince others, and, in doing so, convinces itself. In one study that is currently gaining much media attention, psychologist Drew Westen found that when members of the U.S. Democratic and Republican parties were confronted with evidence contradicting their worldviews, they immediately denied the evidence and (according to brain scan data) took pleasure in doing so.² Our brains reward us for epistemic vice.

Robert Nozick worries about the phenomenon of "anchoring and adjustment".³ Psychologists ask subjects to estimate how much a person's height deviates from "a

¹ Robert Wright, *The Moral Animal* (New York: Vintage, 1994), p. 280.

² D. Westen, C. Kilts, P. Blagov, K. Harenski, and S. Hamann, "The neural basis of motivated reasoning: An fMRI study of emotional constraints on political judgment during the U.S. Presidential election of 2004", forthcoming in *The Journal of Cognitive Neuroscience*.

³ Robert Nozick, *Socratic Puzzles* (Cambridge, MA, Harvard University Press, 1999), pp. 8-9. Thanks to David Schmitz for reminding me of the salience of Nozick here.

benchmark", for instance, how much a person is taller than five feet or shorter than six feet. The benchmark changes the estimation. Subjects guess a higher number when asked how much the target is shorter than six feet and guess a lower number when asked how much the target is taller than five feet. Nozick suggests this will affect a person's evaluation of ideas: we use our current beliefs as a benchmark and measure new ideas in part by their deviation from ours. We did not evolve to carefully measure and access evidence. We evolved to seek social cohesion with peer groups.⁴

Most of our moral views are in part a product of socialization. Our parents, teachers, and other authorities teach us moral precepts, which we internalize as we mature. We are well aware that had we been born two hundred years ago, we would likely have come to accept certain views that we now regard as false: the superiority of men over women or of whites over blacks. We also recognize that at least some of our current beliefs will probably be exposed in the future as unjustified prejudices.

Nearly every moral theory accepts pre-theoretical moral judgments and intuitions as a source of evidence. By "intuition", I mean an automatic moral judgment or evaluation. An automatic moral evaluation is produced without an occurrent reasoning process. Just as we quickly categorize certain animals we see as dogs without consciously reasoning, an intuition is a spontaneous evaluation of some object, state of

⁴ However, empirical research seems to indicate that philosophers are, in fact, better reasoners and better able to see all sides of a issue than non-philosophers. See Deanna Kuhn, *The Skills of Argument* (New York: Cambridge University Press, 1991), pp. 258-263.

affairs, action, person, etc. For instance, when I hear about the genocides in Darfur, I automatically evaluate them as evil without having to go through a process of evaluative reasoning.

In constructing moral theories, we seem to be stuck using intuitions as data. However, we recognize that at least some of our intuitions could be wrong—the result of prejudice or bad judgment. Our data is not perfect. We would like a method of separating the good data from the bad. A common hope is that we will be able to evaluate the trustworthiness of our intuitions by identifying their source and cause. In turn, by determining which moral intuitions, if any, are trustworthy, we can produce better moral theories.

The question of trustworthiness is especially important when we consider that most of our daily lives consists of automatic action.⁵ We have automatic intuitions, but it is not as though we first have intuitions and then act upon them. Consider most moral omissions. My intuition is that it would be wrong for me to kill my colleagues for pleasure. However, as I go about my daily business, the thought of killing them never occurs to me. Thus, it seems false to say that I am acting on an occurrent intuition. A moral intuition typically only occurs only if I am presented with a situation and asked to make a conscious judgment.

The point of this essay is to defend the use of intuitions in moral philosophy. In sections 2, 3, and 4 of this essay, I examine some recent attacks on common moral

⁵ For an overview of work on this issue, see John A. Bargh and Tanya L. Chartrand, “The Unbearable Automaticity of Being”, *American Psychologist* 54:7, pp. 462-479.

intuitions that purport to show they are untrustworthy. In section 2, I summarize the results of a number of experiments Jonathan Haidt, Joshua Greene, and others have performed to test the psychological nature of moral judgment. In section 3, I describe how Peter Singer and Joshua Greene use the results of these experiments in an argument against commonsense moral intuitions and on behalf of their own counterintuitive moral theory. In section 4, I cast doubt on their argument by showing there are limits to what such experiments can prove.

The remaining sections of this essay turn to a more general defense of intuitions in moral theorizing. In section 5, I describe a broader notion of reflective equilibrium, in which considered moral judgments, general principles, and background theories are not only balanced against each other, but also against methodological principles in moral theorizing. I defend reflective equilibrium because it is one of the most common ways that intuitions are employed in moral theory construction. I argue that the fact that all moral philosophers thus far have used their intuitions as data is evidence that using intuitions is a valid method of theorizing. In section 6, I argue that even if the method of reflective equilibrium fails to generate the same moral theory among all users of the method, the method is still justified. Finally, in section 7, I explain how recourse to intuitions is grounded in a moral philosopher legitimately assuming that she has basic moral competence.

2. INTUITIONS AND EMOTION

Moral intuitions are not magic. We all have a conscience, but our consciences are not God's way of whispering in our ears. Also, our consciences tell us different things.

How do we trust them? If God is not whispering in our ears, then who or what is?

Recent work in psychology and cognitive science has tried to answer that question.

Jonathan Haidt claims that the historically dominant model of moral judgment is "rationalism".⁶ According to the rationalist model, moral judgments are primarily the product of a process of rational deliberation. The moral agent enters an "eliciting situation" which calls for judgment. This situation may trigger certain emotional responses. However, the agent normally channels these emotions into the reasoning process, which generates the final judgment. As Haidt notes, the rationalist model has many prominent defenders in psychology, including Lawrence Kohlberg and Jean Piaget.⁷

In contrast, Haidt defends a "social intuitionist" model of moral judgment. In this essay, I am concerned only with the intuitionist rather than the social part of the model. Haidt and his colleagues conducted experiments that seem to indicate that most people make immediate, unreflective moral judgments, and then (sometimes) seek rational justification for their judgments after the fact. Haidt calls this the emotional dog wagging its rational tail: reason appears only post hoc in moral decision-making. In one particular experiment, subjects are presented with a hypothetical story in which siblings decide to have sex. The siblings use two forms of birth control, making the possibility of genetically damaged offspring negligible. Moreover, the story makes clear that they

⁶ Jonathan Haidt, "The Emotional Dog and Its Rational Tail: A Social Intuitionist Approach to Moral Judgment", *Psychological Review* 108:4 (2001), pp. 814-834.

⁷ Haidt, "The Emotional Dog and Its Rational Tail", pp. 815-816.

suffered no emotional harm (nor does anyone else, since they keep their experience secret); indeed, the sex brings them closer together. However, subjects spontaneously evaluate the sibling's actions as wrong. When asked why, the subjects are often dumbfounded. They cannot cite possible harms that could come from the actions, so many just claim that the action was wrong, period.⁸ This type of reaction is common. Haidt's research indicates that many, perhaps most, moral judgments are automatic responses to eliciting situations. Philosophers may form the only group that systematically reasons about moral judgments; even they do so only when philosophizing.⁹

Joshua Greene and his colleagues have conducted functional magnetic resonance imaging (fMRI) scans of the brain during moral decision-making. Greene's research concerned the psychological causes of people's responses to the well-known thought experiments called the "trolley problems". In the first trolley scenario, a runaway trolley is about to kill five workers on the track ahead of it. There is no way to warn the workers or stop the trolley. However, one can flip a switch that will redirect the trolley onto another track that has only one worker. Most people think that it is morally right to flip the switch, sacrificing one person for the sake of the other five. In a parallel case, one

⁸ Haidt, "The Emotional Dog and Its Rational Tail", p. 814. Jonathan Haidt, Fredrick Björklund, and Scott Murphy, "Moral Dumbfounding: When Intuition Finds Not Reason", Department of Psychology, University of Virginia, 2000, unpublished manuscript.

⁹ Kuhn, *The Skills of Argument*, pp. 258-263.

cannot redirect the trolley. However, one has the option of throwing an obese man onto the track. His body is heavy and large enough that he will stop the trolley when hit, though he will die. (In this thought experiment, one is too light to stop the trolley by sacrificing oneself.) This thought experiment normally generates the opposite intuition: it is morally wrong to throw the corpulent man onto the track in order to save the others. These pair of cases, with their apparently conflicting intuitions, has troubled many philosophers.

Greene hypothesized that people have a strong emotional aversion to causing direct personal harm. The difference in generated intuitions between the two cases, then, comes from our emotional responses. In the first case, the harm is impersonal, caused by flipping a switch. In the second, one has to harm the obese man directly by pushing him onto the track. Greene's fMRI research seems to confirm that the different intuitions are caused by a different emotional response. We have an ingrained aversion to causing direct harm.¹⁰ (I.e., fMRI showed that the case of causing direct physical harm triggers more activity in the brain's emotional centers. In addition, subjects that judge it is acceptable to push the obese man in front of the trolley make such judgments more slowly than when judging it acceptable to flip a switch to redirect the track. This

¹⁰ Joshua D. Greene, R. Brian Sommerville, Leigh E. Nystrom, John M. Darley, and Jonathan D. Cohen, "An fMRI Investigation of Emotional Engagement in Moral Judgment", *Science* 293 (2001), pp. 2105-2108. See also the discussion in Peter Singer, "Ethics and Intuitions", *The Journal of Ethics* 9 (2005), pp. 331-352, here 339-342.

suggests that the subjects must override an emotional aversion before arriving at the judgment.)

The most likely explanation for these emotional responses is that they were selected for via evolutionary processes. Since humans evolved in compact family groups, a strong emotional aversion to causing direct harm was advantageous to survival.¹¹ This aversion facilitated cooperation and lessened fighting and stealing.

3. THE VINDICATION OF UTILITARIANISM?

Some, such as Peter Singer and Joshua Greene, think these results favor utilitarian moral theory. Utilitarianism holds that the rightness of an act is solely a function of its capacity to further some good. E.g., hedonistic act utilitarianism holds that 1) pleasure is the only intrinsic good and pain the only intrinsic bad, and 2) an act is right if and only if it contributes more (than any other available act) to the net pleasure experienced by all sentient beings. Preference satisfaction act utilitarianism holds that an act is right if and only if it maximally satisfies (compared to any other available act) the desires of all desiring beings. Both moral theories seem to enshrine beneficence. Intuitively, they claim that the point of morality is to make the world a better place for the types of beings that can have a point of view.

¹¹ Greene, *The Terrible, Horrible, No Good, Very Bad Truth About Morality, and What to Do About It*, Department of Philosophy, Princeton University, 2002, Ph.D. dissertation, p. 178.

However, act utilitarianism (of either variety) quickly loses its intuitive appeal. Many versions of utilitarianism countenance actions most of us find morally obscene provided that the numbers work out correctly: sacrificing the few for the many, the killing of innocents, disregard for special personal ties and projects. They deny many commonsense moral distinctions, e.g., claiming that there is no difference between my failing to feed a starving Bengali and my actively killing him. Some act utilitarians, bothered by these results, seek to reformulate their theory to better align it with commonsense morality. Other act utilitarians claim that commonsense morality is wrong and that we average moral agents are corrupt.

Before turning to the argument for utilitarianism, let us review Greene's and Haidt's results. These will serve as premises in an anti-commonsense morality, pro-utilitarian argument.

1. Many moral judgments are automatic responses. Subjects often cannot offer reasons for their views, but they strongly believe they are right.¹²
2. The intuitions and judgments that conflict with act utilitarianism tend to be strongly emotion-based. These emotional responses were selected for via evolution. Many commonsense moral views are instinctually ingrained within us because they contributed to reproductive success.¹³

¹² Jonathan Haidt et al, "Moral Dumbfounding".

¹³ Singer, "Ethics and Intuitions", pp. 336.

3. Reasons for these judgments are often constructed after the judgment has been formed.¹⁴
4. Moral judgments according with act utilitarianism tend to activate reasoning centers in the brain.¹⁵

Utilitarians are quick to celebrate these results. They believe the results provide some reason to be suspicious of the intuitions conflicting with utilitarianism.

Here, then, is an argument utilitarians might make. Let us call this the Argument from Emotion.

Argument from Emotion

1. Commonsense moral intuitions (many of which contradict utilitarianism) are merely emotional responses contingently selected for via evolution.¹⁶

¹⁴ Jonathan Haidt, "The Emotional Dog and Its Rational Tail: A Social Intuitionist Approach to Moral Judgment".

¹⁵ See the discussion in Peter Singer, "Ethics and Intuitions", *The Journal of Ethics* 9 (2005): 331-352, esp. pp. 349-350. Singer heavily relies upon Joshua Greene, *the Terrible, Horrible, No Good, Very Bad Truth about Morality and What to Do About It*.

¹⁶ Greene says that these responses reflect "arbitrary features of our evolutionary history". See Joshua Greene, "Reply to Mikhail and Timmons", in Walter Sinnott-Armstrong (ed.), *The Psychology and Biology of Morality* (Cambridge, MA: MIT Press, forthcoming).

2. Therefore, there is no reason to think these emotions reflect some underlying moral reality.¹⁷ A moral theory need not attempt to accommodate them.
3. Utilitarian intuitions, if they can be called that, are more reason-based and are not the contingent results of evolution. They are more likely to be ethical axioms or "rational intuitions".¹⁸
4. Therefore, utilitarianism is better justified than those ethical theories (like deontology) that cater to commonsense moral intuitions.

Peter Singer makes this argument, though he admits it is but a sketch and needs more work.¹⁹ Greene is more forceful. In responding to the objection that it is acceptable for moral judgment to be grounded on sentiment, he says, "Garbage in, garbage out." Evolution-derived emotions (garbage in) form the basis of our moral judgments (garbage out). If we do not want to produce garbage, we will ignore our emotion-based intuitions and proceed a different way in constructing moral theories.²⁰

For those of us who find more extreme utilitarian conclusions repugnant, this argument is worrisome. The psychological experiments of Greene, Haidt, and others

¹⁷ Joshua Greene, "The Secret Joke of Kant's Soul," in *The Psychology and Biology of Morality*, p. 45.

¹⁸ Singer, "Ethics and Intuitions", p. 350-351.

¹⁹ Singer, "Ethics and Intuitions", p. 351.

²⁰ Greene, "Reply to Mikhail and Timmons", in *The Psychology and Biology of Morality*.

deeply challenges the commonsense picture of how moral "reasoning" works. Perhaps our intuitions are merely the outcome of evolution, contingent impulses implanted within us to facilitate reproduction, and nothing more. If so, what legitimate claim can these intuitions have on us? If they have no claim, then utilitarianism's counterintuitiveness is not a mark against it. Indeed, by the lights of the argument, utilitarianism is better justified than competitor theories such as virtue ethics and deontology.

4. INVERSE UTILITARIANISM

The argument, though powerful, cannot be right. Singer says the argument is a mere sketch and needs further work. In this section, I show why the argument needs more work. Singer, Greene, and others seem right to think that empirical research should have some normative implications, i.e., it should have some bearing on our moral theories. Yet, it is unclear that the evidence could show what Greene and Singer think it shows.

Here is the Argument from Emotion in abstract:

The Argument in Abstract:

Theory X is to be favored over Theory Y when it can be shown that

1. Intuitions according with X are more reason-based and less emotion-based than intuitions according with Y,
2. Intuitions and judgments according with X are less the contingent outcomes of biological evolution, and

3. Judgments according with Y tend to be made spontaneously and without reflection, with reasons rationalized only after the judgment is made, whereas judgments according with X tend to be made after a process of deliberation and weighing of reasons.

The Argument in Abstract is intended to show via empirical research that one theory is better justified than another. Each of the conditions for favoring X over Y must be established by experimentation or other scientific investigation. The Argument is supposed to allow scientific results to help determine a theory's acceptability. Given two competing moral theories X and Y, we should be neutral in appraising them until we have run the relevant experiments, assuming both theories are consistent and satisfy other desiderata.)

Suppose that we are trying to choose between two theories, hedonistic act utilitarianism and some other theory X. Assume both theories are consistent, simple, and so on, but their major difference lies in their content. The Argument counts in favor of utilitarianism over deontology. If the experimental results come out the right way, then it would favor X over utilitarianism.

However, it is far from clear that this Argument works. Suppose I advocate a new theory, *inverse utilitarianism*. According to inverse utilitarianism, an act is right just in case, in comparison to other available actions, it causes the maximum net suffering among all sentient beings. Inverse utilitarianism is a moral theory that enshrines maleficence.

Inverse utilitarianism is a very counterintuitive moral theory. It is even less intuitive than the advocacy of state-sponsored murder by Hitler, Lenin, or Stalin. In a contest of intuitive plausibility, utilitarianism beats inverse utilitarianism easily. But, no matter. In a contest of intuitive plausibility, commonsense deontology beats Singer's form of utilitarianism, yet Singer and Greene think that utilitarianism is vindicated because the experimental results come out the right way. When fit into the Argument in Abstract, utilitarianism wins over deontology.

Accordingly, suppose we run Haidt's and Greene's experiments, this time testing between utilitarianism and inverse utilitarianism. After running the experiments, suppose these are our results:

1. Inverse utilitarian judgments are less emotion-based and more reason-based than utilitarian judgments.
2. Reasons for inverse utilitarian judgments occur to subjects before they make such judgments, whereas, in comparison, utilitarian judgments tend to be rationalized after the fact.
3. Inverse utilitarian judgments are not the products of biological or cultural evolution, but utilitarian judgments partially are.

If Singer and Greene's argument (the Argument from Emotion) favors utilitarianism over deontology, then our getting these results would favor inverse utilitarianism over utilitarianism.

Now, for our purposes, it need not matter what the actual results would be. (Though, there are some reasons to think we would in fact get these results.) Even if we do not yet have the results of the tests, we know *ex ante* what certain results will tell us. E.g., even if I have not yet dipped blue litmus paper into some liquid, I know that if blue paper turns red, the liquid is acidic.

Singer and Greene's argument implies that if we get the results above, then inverse utilitarianism should be favored over utilitarianism. However, even if we obtain such results, I doubt that anyone would regard this as a victory for inverse utilitarianism. Thus, it becomes unclear exactly how the Argument works.

One reason we might fail to regard these results as a victory for inverse utilitarianism is that it seems to be a perversion of anything resembling morality. But so what? On Singer and Greene's reasoning, our intuitions on these matters are not trustworthy.

Alternatively, perhaps the results do speak in favor of inverse utilitarianism, but the evidence is overridden by some other consideration. For instance, perhaps an analysis of the function of moral language or some other metaethical concern would override the evidence. Yet, it is not clear that any considerations of this sort exist, unless they appeal to our intuitions.²¹ Most likely, the contrary evidence would be some

²¹ Following R. M. Hare, we might attempt appeal to linguistic intuitions about the use of moral language rather than moral intuitions. However, it is not clear that this distinction can be maintained.

philosophical argument for contrary moral principles, and these would ultimately be rooted in untrustworthy intuitions.

However, it seems to me that evidence, rather than being overridden, *does not count in favor of inverse utilitarianism at all*. I.e., even if we ran fMRI and other experiments and obtained these results, they would not count even slightly in favor of inverse utilitarianism. The contest between inverse utilitarianism and Singer's utilitarianism is not decided by empirical research. There is no contest. Though the results above are logically possible, it is not logically possible that inverse utilitarianism is a better moral theory than Singer's.²² Inverse utilitarianism is simply an inadmissible view. I would expect Singer and Greene to agree. If I am right that inverse utilitarianism is not logically possible, no empirical evidence could count in favor of it, just as no empirical evidence could count in favor of " $2 + 2 = 5$ ".

Alternatively, maybe Singer and Greene would agree that no evidence could count in favor of inverse utilitarianism. They might answer my challenge as follows: The reason that such evidence could favor utilitarianism over deontology, but not inverse utilitarianism over utilitarianism, is simply that evidence—any evidence of any sort—can only count in favor of possible views. Utilitarianism and deontology are moral possibilities, but inverse utilitarianism is not. Thus, psychological evidence can favor utilitarianism or deontology, but not inverse utilitarianism. However, this response

²² At a high level of abstraction, this argument is similar to the Moral Twin Earth argument given in Mark Timmons, *Morality without Foundations* (New York: Oxford University Press, 1999), pp. 59-70.

would not help. The reason we know that inverse utilitarianism is not an admissible view is that it violates our most central, intuitive moral convictions. If inverse utilitarianism is morally possible, then we know nothing about morality and might as well be moral skeptics.

No experiments could increase our confidence in inverse utilitarianism.

Similarly, skeptical thought experiments do little to bolster the theory. Suppose I learn that the reason I find inverse utilitarianism so repugnant is that Descartes' Evil Demon has implanted this reaction in my brain. Even this thought does not improve my appraisal of inverse utilitarianism.²³ Inverse utilitarianism is so perverse a moral theory that it is not an admissible theory at all.

5. WIDENING REFLECTIVE EQUILIBRIUM: METHODOLOGY

Singer is undoubtedly right to worry that some of the intuitions used in constructing moral theories are unjustified prejudices. Even those moral philosophers who see their task more as providing normative explanations than normative guidance should agree. Singer admonishes us against being irresponsible with our data.

Ideally, before we construct our moral theories, would already have separated the good intuitions from the bad. We have no perfectly reliable mechanism for such sorting.

²³ At best, it drives me to skepticism, in which I may admit that it is possible that inverse utilitarianism is possible, but not admit that inverse utilitarianism is in fact possible.

In theory construction, we know we must start somewhere. As with any journey or any task, we must start from where we are.

In one way or another, every moral philosopher appeals to moral intuitions. Even foundationalists like R. M. Hare, who appeals to the logic of moral discourse, or Richard Brandt, who appeals to contingent features of human psychology,²⁴ ultimately must test their principles against moral intuitions and considered moral judgments. Hare and Brandt produce utilitarian theories, and their theorizing inspires enough confidence for them to revise their pre-theoretical judgments that conflict with utilitarianism. However, had Hare's or Brandt's approaches resulted in inverse utilitarianism, this would show they had made mistakes. Even when intuitions are not used as foundational data, they are used as a negative test of a theory. Everyone uses intuitions and everyone engages in a broad sort of coherentism.

If every moral philosopher uses intuitions, it is somewhat strange to attack the use of intuitions. Claiming that moral intuitions should have no evidentiary impact in theory construction amounts to claiming that moral philosophy ought not be done.

A degree of historicism is now rather common in the philosophy of science. That is, most philosophers of science agree that a good theory of scientific rationality and method ought to fit actual scientific practice. Moral philosophers who write on moral metatheory—the theory of proper methods in moral philosophy—ought to take a lesson

²⁴ For an overview of these and other foundationalist approaches, see Mark Timmons, "Foundationalism and the Structure of Ethical Justification", *Ethics* 97:3 (1987), pp. 595-609.

from philosophers of science. If they produce a model of rationality in theory construction but find that few or no moral philosophers actually fit the model, this should be taken not as evidence that most moral philosophers are irrational, but that the model is mistaken. Similarly, we should be skeptical when any ethicist announces that she has discovered that all other ethicists have heretofore used incorrect philosophical methods. Moral philosophers have engaged in diverse projects with diverse aims, but some appeal to moral intuitions, often through reflective equilibrium, is common to all. A good moral metatheory, then, ought not discount intuitions, but rather explain when intuitions are being used well and poorly.

John Rawls advocates reflective equilibrium as a method of theory construction.²⁵ (Though this paper is a defense of intuitions in general, I will discuss reflective equilibrium because it is one of the most common ways of appealing to intuitions when constructing a moral theory.) According to Rawls, not all moral intuitions are on a par. We are more confident of some intuitions than others. Moral theory construction involves balancing our considered intuitions about particular cases against general principles. Sometimes principles are modified in light of intuitions, and vice versa. We continue the process of modifying our judgments and principles until the pressure to

²⁵ John Rawls, *A Theory of Justice*, 2nd Edition (Cambridge, MA: Harvard University Press, 1999), pp. 17-19.

change members of either set is equalized.²⁶ Then we have arrived at "reflective equilibrium".

It is normal to distinguish between narrow and wide reflective equilibrium. Narrow reflective equilibrium is reached through considering individual cases and general principles. Judgments about individual cases are revised when they conflict with principles in which we have greater confidence, and general principles are revised when they conflict with judgments about individual cases in which we have greater confidence. Thus, narrow reflective equilibrium balances 1) considered moral judgments and confident intuitions about particular cases with 2) general principles. Wide reflective equilibrium adds a third set of considerations: relevant background theories. These theories include philosophical theories of personhood and the role of morality in society, as well as relevant economic, psychological, sociological, and anthropological theories.²⁷ Singer and Greene's invocation of experimental psychology as evidence against deontological intuitions falls within the method of wide reflective equilibrium.

However, in practice, reflective equilibrium is wider than many people have recognized. Whereas wide reflective equilibrium is cast as balancing three items (considered moral judgments, general principles, and background theories), we should

²⁶ For an excellent summary of this method, see also T. M. Scanlon, "The Aims and Authority of Moral Theory", *Oxford Journal of Legal Studies* 12:2 (1992) 1-23, here pp. 2-3.

²⁷ Norman Daniels, "Wide Reflective Equilibrium and Theory Acceptance in Ethics", *The Journal of Philosophy* 76:5 (1979), pp. 256-282, here pp. 258-262.

cast it instead as balancing at least four. The fourth item is proper philosophical method, i.e., what sorts of arguments and data count (and in what way) when producing moral theory.

We are not balancing Greene's and Haidt's psychological results against our considered moral judgments or general principles. Rather, we are balancing judgments about the *evidentiary force* of such results against our considered moral judgments and general principles. The inverse utilitarianism thought experiment shows us certain limits on *what counts as evidence* in evaluating the trustworthiness of moral claims. Singer and Greene try to undermine judgments that conflict with utilitarianism by producing evidence that these judgments do not track any moral truth. They argue against considered moral judgments by trying to show that using such moral judgments in theory construction is bad methodology. In effect, what we are asked to do falls within reflective equilibrium. We balance our confidence in our moral judgments against the evidence that these judgments could be mere prejudices. In the hypothetical case in which inverse utilitarianism is favored by Greene and Haidt's research, our confidence in our basic moral outlook outweighs the undermining force of the Argument from Emotion.

Everyone recognizes that philosophical methods constrain theory and moral beliefs. However, theory and moral beliefs in turn constrain philosophical methods. We should not accept the methodological view that certain experiments could prove inverse utilitarianism right. If any philosopher proposes a new method, and this method leads to

inverse utilitarianism, then the method has failed. Thus, methods are testable by our moral intuitions.

6. DIVERGENCE, REFLECTIVE EQUILIBRIUM, AND APPEAL TO INTUITIONS

Reflective equilibrium is a method employed by many moral theorists.

However, reflective equilibrium does not appear to lead to agreement on moral theory or convergence on a basic moral outlook. (In general, appeals to moral intuitions do not generate much convergence on theory.) One might take this as evidence that the method is not being employed properly. If Hilary Putnam and Robert Nozick disagree, it must be that at least of one of them is misusing the method. But this seems implausible.

Philosophers like Putnam and Nozick at their very best continue to disagree.

Alternatively, one might take this as evidence that even if reflective equilibrium is a method employed by all, it must be supplemented by some yet to be discovered method that will at last produce rational convergence. The problem with the method is that it is too indeterminate. This too seems implausible. Previous attempts to discover a rational method for producing convergence have failed.

This is taken by some as evidence that the method is ineffective. For example, Michael Smith argues for a form of moral rationalism. In his view, an action is right in certain circumstances just in case we would desire to perform that act in those

circumstances were we fully rational.²⁸ Smith intends to produce an objectivist account of moral rightness. For us to have a normative reason to perform some action, it must be the case that we would all desire to perform the same action in the same circumstances were we fully rational.²⁹ Smith himself claims that his view requires convergence between rational agents. Smith first specifies what it is for an agent to be fully rational, and then claims that all rational agents in identical circumstances would agree on the same norms, desire the same things in the same circumstances, and so on. If rational agents would not converge, then moral skepticism must be true.³⁰ Smith intends reflective equilibrium to be the method for arriving at rational moral beliefs.

Smith does not think that moral skepticism is true—he thinks that convergence would be reached. However, some critics have argued that no convergence is forthcoming. E.g., Terence Horgan and Mark Timmons argue that the disagreement between Robert Nozick and Hilary Putnam shows that Smith is wrong—convergence is not reached as agents become more rational (in Smith's sense of "rational").³¹ If Nozick

²⁸ Michael Smith, *The Moral Problem* (Malden, MA: Blackwell, 1994), p. 184. In addition, the act must satisfy certain platitudes about moral substance.

²⁹ Smith, *The Moral Problem*, p. 167.

³⁰ Smith, *The Moral Problem*, pp. 200-201.

³¹ Terence Horgan and Mark Timmons, "Troubles for Michael Smith's Metaethical Rationalism", *Philosophical Papers XXV:3* (1996), pp. 203-231, here p. 212.

and Putnam genuinely disagree, and their disagreement is not the result of irrationality or a lack of information, then on Smith's view, moral skepticism follows.

This reasoning can be employed against reflective equilibrium and the use of moral intuitions more generally. Critics often complain that reflective equilibrium and the appeal to intuitions cannot be rational because they do not guarantee convergence on any beliefs. The assumption seems to be that we have a dilemma between convergence and skepticism.

More broadly, throughout philosophy's history, philosophers have tried to specify "decision procedures" or sets of conditions under which rational decisions and judgments can be made. But many people criticize these by pointing out that unless we load in questionable assumptions, we do not get unique results. If the procedure does yield unique results, then it must be because questionable assumptions "load" the decision. Still, there is an assumption that a properly formulated rational decision procedure must have a unique outcome; otherwise, the procedure is not rational. This again means we have a dilemma between convergence and skepticism.

A. *Rational Convergence*: All fully rational agents in the same circumstances will converge on the same beliefs and judgments.

B. *Skepticism*: There is no truth of the matter or the truth is inaccessible to us.

However, it is not obvious why we cannot escape the dilemma by accepting C.

C. *Rational Divergence*: In the pursuit of truth, fully rational agents, even if they possess the same information and are in the same circumstances, will continue to disagree about some central beliefs and judgments. This disagreement does not entail that the agents are unjustified, nor does it entail there is no truth of the matter.

C would strike many as a worrisome view. In particular, they would worry that C collapses into D:

D. *Relativism*: Truth is relative to the agent. Thus, if two fully rational agents disagree, they both have beliefs that are *true for them*.

D is unacceptable for familiar reasons, which I will not rehearse here. The difference between C and D is metaphysical. A person who holds C but not D accepts that two disagreeing agents might both be *justified* (i.e., they both have proper epistemic warrant for holding their beliefs), but not both *correct* (i.e., they cannot both have true beliefs).

Is C true? It had better be. Philosophy in general, not just moral philosophy, is marked by divergence. If philosophy is not irrational, then C must be true. But if philosophy is meant to find truth and truth is not relative, then D must be false as well. In particular, what philosophers want (or should want) is to combine C with E.

E. *Standards*: It is rational to search for and demand reasons for different philosophical positions. Some positions are better justified than others. Some arguments are better than others. Some philosophical views are unjustified.

If C and E are false, the philosophy in general is a failed enterprise.

One reason that reflective equilibrium results in disagreement is path dependence. Suppose you and I have identical intuitions. However, in constructing our moral theories, I begin by considering one intuition (e.g., about rights) and you begin with another one (e.g., about obligations). Starting with different premises changes our propensity to accept new principles as we theorize. Though we have the same pre-theoretic intuitions, our different starting points lead us to different theories. Had I started with a different premise, I would have different tests for the acceptability of later principles.³² However, there is rarely one obvious correct starting point. Starting points are largely optional.

This is meant to be a defense of reflective equilibrium and of the use of intuitions more generally. One of the most common charges against the reflective equilibrium is that it leads to disagreement rather than agreement. In particular, it is claimed that when thinkers have different starting premises, they will likely stop at different final theories. However, there is no reason to hold that thinkers with the *same* starting premises would stop at the same theory either. Reflective equilibrium is achieved by balancing many sets of considerations, but there are multiple ways of achieving such balance. Balancing involves choice; there are usually multiple, equally compelling ways of achieving

³² Thanks to David Schmitz for this point.

balance. If reflective equilibrium is a rational method, then rational disagreement *at the limit of inquiry* must be acceptable. But those who find this unacceptable should note that disagreement seems to be at the limit of all philosophical inquiry, regardless of whether reflective equilibrium is the method used. *Vis-à-vis* the problem of disagreement, reflective equilibrium in moral theory is no worse off than philosophical methods in general.

7. SELF-TRUST AND THE ASSUMPTION OF BASIC COMPETENCE

In order to pursue knowledge, we must assume that we know something to begin with. We must grant our initial beliefs and our intuitions some degree of credibility. We need not assume that we cannot turn to have been wrong. Indeed, in the case of scientific knowledge, many of our starting beliefs about nature did turn out to be false. Rather, the necessary assumption of theorizing—any theorizing—is that we are not completely hopeless. We cannot find our way out of darkness with the assumption that we are blind. Philosophy is ultimately grounded on self-trust, even as philosophy is the discipline most likely to challenge that trust.

In what way must we trust ourselves? For one, we trust that we have enough moral competence to know what morality is about. We know that a theory describing the proper way to play chords on guitar is not a moral theory. Thus, we trust that we can distinguish morality from non-morality, even though we know the border between the two is fuzzy. We trust that inverse utilitarianism is not a moral possibility. *Ex ante*—before we theorize—we know that any philosophical method that leads us to the principle

of reverse utility is a bad method. But if we know this before we theorize, then we must have pre-theoretical moral knowledge.

Might we be wrong about these inverse theories? Perhaps inverse utilitarianism is not logically possible, but it is logically possible that inverse utilitarianism is logically possible. Better yet, we might cast the possibility of inverse utilitarianism in terms of epistemic possibility. I know that inverse utilitarianism cannot be true, and perhaps I know that I know, but I do not know that I know that I know. There is a possibility of error, but it is remote.

Why should I trust myself? What justifies my assumption that I have a basic competence about morality, a competence that permits me to engage in moral theorizing? It is true that without that assumption, I cannot theorize, but the question is whether I can theorize. Unfortunately, there does not seem to be anyway to prove that assumption without using that assumption. I have reasons to think I am a trustworthy moral agent: I was raised in a good social environment by parents that were of (as far as I can tell) decent moral character. I am aware of the nature of prejudice and actively seek to eliminate it within myself. I seek and consider views opposing my own. On John Stuart Mill's recommendation, I even try to get inside opposing views, to feel what it would be like to accept them. I have an extensive general educational background, including work in certain sciences and social sciences. I submit my views to the criticism of experts (as far as I can tell they are experts). I have undergone philosophical training in a top philosophy graduate program. And so on. Before I theorize, I believe that I have these reasons for trusting myself, and this belief survives theorizing. But if I made no such

assumption of my own trustworthiness, it would be impossible for me to prove that I am trustworthy.³³ Thus, my trustworthiness survives one test: The assumption that I am generally trustworthy does not undermine itself. But it does not pass another: I cannot prove that assumption without making the assumption. This is the best I can do. There are limits to what I can prove.

Moral philosophers begin with the assumption that they exhibit basic moral trustworthiness. Normally, this assumption does not undermine itself. Because the assumption is not self-undermining, they allow themselves to use moral intuitions as evidence in moral theorizing. Reflective equilibrium, in which intuitions play a part, is a process of self-criticism, but it is constructive criticism, built on the trust that one can get better.

³³ Thanks to Keith Lehrer for discussions regarding trustworthiness and justification. See Keith Lehrer, *Theory of Knowledge, 2nd Edition* (Boulder: Westview, 2000), pp. 138-144, 201-203, 209-212.

CHAPTER EIGHT

THE PRACTICAL VALUE OF MORAL THEORY

1. SHOULD THEORY BE PRACTICAL?

To paraphrase Marx, the point of analyzing morality is not to interpret the world, but to change it. I would be pleased if my moral theory could improve the world. Yet, I do not regard theories as failures even though they normally have little direct effect on the world.

Moral theory is no substitute for virtue. Nor is it a cure for vice. This seems so obvious that is almost not worth saying.¹ However, both friends and foes of theory come close to acting as if moral theory is meant to substitute for virtue. E.g., philosophers have tried to produce texts explaining the authority of morality with the hope of convincing the immoral to heed morality's dictates. Immorality persists, despite the prevalence of such texts. Some might take that as evidence that the explanations are wrong. Others take an even more skeptical stance: they claim that moral theories are worthless. They say moral theory must be a practical affair for it to have any worth.

Many people say such things, but are there any reasons to believe it?

Mathematics is as purely a theoretical subject as possible. However, thanks to

¹ The converse, virtue is no substitute for moral theory, seems less obvious to many people, but I believe it is also true. See my "Moral Theory and Its Discontents", this volume.

engineering, many branches of mathematics turned out to be immensely practical. Still, many branches have no practical application, as of yet. Yet, we do not count this against mathematics. It would be nice if moral theorizing turned out to be practical, to be instrumental in transforming the world into a better place. But suppose it does not. In contrast to the case of mathematics, does this show moral theory is worthless?

What is the practical value of moral theorizing? In what way is moral theory practical? I will argue that, for the most part, I do not know and neither do you. Additionally, I will argue that neither you nor I are in a good position to find out. (I am assuming the reader is a fellow philosopher.) The practicality of moral theory, for the most part, cannot be settled without empirical investigation. Moral theory's usefulness is matter of psychological, economic, and sociological contingency. We can't determine it from the armchair. However, should we come to learn that moral theory is not practical, it is still a worthwhile theoretical activity. Pure theory, apart from its practical effects, is a valid enterprise.

2. HOW COULD THEORY RELATE TO PRACTICE?

Frances Kamm lists several different ways someone might want a theory to relate to practice.² She does not pretend her list is complete. I will provide such a list here, noting that many items on the list are hers or modifications of her suggestions.

² Frances Kamm, "High Theory, Low Theory, and the Demands of Morality", *Nomos* XXXVII (1995), pp. 81-107, esp. p. 82.

A theory might tell us what to do or to value. A theory might contain a few principles, which, when combined with a description of context, entail certain evaluations. David Brink notes an important distinction here.³ On one hand, these principles might merely be a *criterion of rightness* or *goodness*, a theoretical construct that need not be consciously employed in ethical decision-making. On the other hand, they may be a self-consciously employed *method for making decisions*. When people speak of theories offering a "decision procedure", this is often what they have in mind. A theory articulates a code or algorithm that replaces moral judgment and tells us what to do and what to care about.

A theory can have implications, not directly about courses of action, but the right way to deliberate in making moral decisions. Whereas some theorists have attempted to offer moral principles from which one can deduce the right course of action, others attempt to describe the correct states of mind needed to make moral decisions. These are less ambitious than "decision procedures". Decision procedures attempt to bypass good judgment. Theories about proper deliberation allow that moral decisions require judgment, but try to explain what good judgment is.

A theory might be testable by its real world implications. Even if a theory does not have an algorithm, it may imply that certain actions ought to be done. However, most people have moral knowledge prior to theory; indeed, it is this moral knowledge that we

³ David O. Brink, "Utilitarian Morality and the Personal Point of View", *The Journal of Philosophy* 83:8 (1986), pp. 417-438.

rely upon in constructing and testing theories.⁴ Suppose we construct a theory of justice, thinking it an excellent theory, until we realize that the theory implies that killing innocents for the sake of expedience is morally acceptable. This gives us pause, and it normally gives us reason to reject the theory.

A theory can revise our pre-theoretical beliefs about what to do or value.

Sometimes when our pre-theoretical considered judgments and the implications of our theory conflict, we take this as reason to reject our pre-theoretical judgments. We begin theory construction in part by trying to accommodate "considered judgments"—those moral judgments about which we believe we are unbiased and justifiably confident. Our theoretical investigations may remove our confidence. To paraphrase T. M. Scanlon, moral theory clarifies our reasons for holding our considered judgments.⁵ Sometime this clarification strengthens our conviction, but sometimes it reveals confusion, bias, self-deception, or simple lack of support.

*A theory might be purposefully and properly isolated from practice.*⁶ John Rawls, for instance, argues that just liberal democratic societies respect an ideal of public

⁴ Shelly Kagan denies that moral knowledge is possible without theoretical justification.

I argue that this view is mistaken in chapter 1 of this volume. See Shelly Kagan, *The Limits of Morality* (New York: Oxford University Press, 1991), pp. 11-13.

⁵ T. M. Scanlon, "The Aims and Authority of Moral Theory", *Oxford Journal of Legal Studies* 12:1 (1992), pp. 1-23, here p. 16.

⁶ See Kamm's discussion on pp. 87-88.

reason.⁷ While individuals will have differing comprehensive moral, philosophical, and religious views, they generally keep these views isolated from public debate. Instead, they rely upon a constructed public morality, to which they subscribe because it bears a certain relationship to their own private moral codes. If Rawls were right, then even if, say, Scanlon's contractualism were the correct moral theory, a Scanlon in political office would normally not justify his policy proposals via his contractualism.

A theory can be self-effacing in practice. Attempts to "live by" the theory might fare badly, as measured by that theory. For example, even if act utilitarianism is correct, and thus we ought to choose whatever action maximizes aggregate utility, it might turn out that *attempting* to maximize utility normally fails. The best way to maximize utility might be a policy of ignoring utility altogether.

More generally, accepting a theory might affect people's character. Accepting or rejecting certain theories might be a mark for or against one's goodness as a human being. E.g., it's hard to see how one could accept Nazi political theory without this detracting from one's moral character. Another way of interpreting this relation: There is the Platonic hope that a proper explication of the nature of morality will *cause* people to become moral. I.e., if we show people what morality really is, they cannot help but desire to promote the good. This view holds that vice is the result of misunderstanding.

A theory might posit ideals, but these ideals are not realizable in practice. In eudaimonist moral theory, there is the *phronimos*, the fully virtuous person who exhibits

⁷ John Rawls, *Political Liberalism* (New York: Columbia University Press, 1993), esp. pp. 212-254.,

complete moral self-mastery. This is an ideal we can approach and approximate, but not reach. Alternatively, some fundamentalist Christian ethicists, such as John Calvin, hold Jesus as an ideal, but claim that we are so depraved that (but for God's grace) we cannot even approach the ideal.

A theory might resolve disagreements. Sometimes people look to theories because their considered moral judgments conflict. Philosophers sometimes hope that theoretical investigation will remove this conflict.

A theory might have unanticipated practical benefits. Certain branches of mathematics turned out to be useful in strange ways. So might a moral theory. We might discover that for some people, the life of contemplation involved in thinking about moral theory over a period of years lowers blood pressure, reduces accident rates, counters tendencies to substance abuse, and so on. Analogously, some controversial studies have shown that listening to Mozart can temporarily boost performance on certain spatial reasoning tasks.⁸ Mozart's aim was to produce beautiful music, but this is an unexpected practical benefit of his work.

These are some ways moral theory could be practical. The next question is in which ways is moral theory in fact practical. The answer is a mixed bag. Moral theories cannot be practical in certain respects. As for whether they are or should be practical in other respects, we do not know.

⁸ F. H. Rauscher, G. L. Shaw, & K. N. Ky, "Music and Spatial Task Performance", *Nature* 365 (1993), p. 611.

3. WHY THERE CANNOT BE A UNIVERSAL DECISION PROCEDURE

A classic argument against act utilitarian decision procedures runs as follows. Suppose the right way to make decisions is to calculate the expected utility of every feasible course of action and then do whatever has the highest expected utility. However, before doing so, one must first calculate how long one should spend calculating. Calculating possible outcomes takes time away from better projects, such as feeding Bengal. Before making this calculation, though, one must calculate how long one must calculate how long one must calculate...

This objection can be generalized to any theory that purports to have a decision procedure. Decision procedures are supposed to be algorithms for performing correct actions. The problem is that *deliberating, using a decision procedure, and making a decision* are themselves actions. Just like other actions, they are morally evaluable. In some cases, *using a decision procedure* will be judged wrong *by its own lights*. The reason is the same as with the utilitarian case. Calculating takes time away from relieving famine. According to the act utilitarian decision procedure, it's better to relieve famine than to deliberate about relieving famine. But any theory's decision procedure faces the same problem. The act of deliberating via the decision procedure can conflict with performing other actions or promoting other values. There are always cases where a given decision procedure fails to pass self-inspection.

Here is a strange, but clear, case of this. Suppose a gunman grabs your child and puts a gun to her head. He says, "I'll shoot her unless you give me your money. But I'll also shoot her if you use a decision procedure." A theory might imply the proper course

of action: one should hand him the money without using a decision procedure. Here one must use one's judgment, not a decision procedure.

Constructing more realistic cases is always possible, but will depend on the particular values of a theory. In general, a decision procedure aims to produce certain values, such as general happiness, and to prohibit certain actions, such as neglect. To find a decision procedure's limits, one need only imagine the cases where using the decision procedure inhibits general happiness or amounts to neglect.

Decision procedures cannot escape this problem by having a clause in the algorithm that says, "In cases where using this algorithm is morally wrong, do not use it." Once an agent has gotten to that line of the algorithm, it is too late.

Some decision procedures do not succumb to such counterexamples. The *null decision procedure* says, "Do anything". It is not ruled out, but it is not really a decision procedure (moral or otherwise). The *enshrined decision procedure* has a lexical ordering of rules. Its first rule is, "Use this decision procedure, no matter what." Its next set of rules could be anything, e.g., some utilitarian calculus. However, the enshrined decision procedure is unreasonable, because rule one is unreasonable regardless of the content of rule two. The enshrined decision procedure cannot plausibly be said to track moral, aesthetic, prudential, or other value, because it makes *itself* most important.

4. TESTING IDEALS BY PRACTICALITY⁹

⁹ This section incorporates and revises material from my "Modesty without Illusion", *Philosophy and Phenomenological Research* (forthcoming).

Many theorists hold an ideal of virtue, but also hold that *perfect* virtue is unattainable. What then is the point of having ideals, if we cannot achieve them? I will defend the role of regulative ideals in moral thinking. Below I describe what ideals are and are not for. However, the view that this is the proper function of ideals depends in part on certain optimistic assumptions about human psychology. As we'll see in §6, if human nature turns out badly, the proper function of ideals might be quite different.

Ideals must be attractive, rather than alienating. Some philosophers have argued that a world of perfectly impartial moral agents is an ideal, because presumably in such a world widespread altruism will produce the most good. But this is an ugly ideal. When normal people contemplate the things that give their lives meaning and worth—children, personal projects, relationships, small personal associations—they see that meaning and partiality go together. Few people find attractive a world in which parents do not have a special bond with their children. Perfect impartiality obliterates the special intimacy of lovers.

On the other hand, the ideal of self-perfection pursued by the wisest person is attractive. Under all but the worst of circumstances, the phronimos leads a life without regret, confusion, or alienation. The phronimos can feel deep love for his family. As Plato describes him, his soul is in harmony. He has achieved a happy life, and a life in which virtue is a constitutive part.

We must be able to approach the ideal. For an ideal to serve its regulative role in our behavior, there must be coherent notions of approaching or retreating from that ideal.

If the ideal is so foreign to us that nothing we can do is a step towards it, it cannot regulate our behavior. It is no longer an ideal; it is just alien.

How closely can we approximate a proper ideal? It does not make sense to say that the ideal runner could run the speed of light. Such an ideal is so infeasible as to be unattractive, not worth aiming for. Rather, ideals, moral or otherwise, should be asymptotes. With enough effort, one can come inspiringly close to the ideal. The best will still see the distance between themselves and the limit.

Attempts to live by the ideal ought generally to have good results. Attempts to make an entire, large-scale society adhere to the principle of distribution, "To each according to his needs, from each according to his abilities", have a robust history of resulting in distrust, free riding, moral hazard, and often starvation and death.¹⁰ Thus, it is absurd to say it is ideal in theory but bad in practice. An ideal in theory might not be completely attainable in practice, but attempts to practice ought to have good results. Attempting to become the phronimos normally improves one's character. Increasing one's understanding of the phronimos typically makes one both more aware and more receptive to salient moral reasons. People who attempt to live by an image of the good normally improve their character.

Ideals are not for condemning people. The function of ideals is to attract one to the good and foster self-improvement. They are not for reproaching others or to foster a

¹⁰ See David Schmitz, "Taking Responsibility" in David Schmitz and Robert Goodin, *Social Welfare and Individual Responsibility* (Cambridge: Cambridge University Press, 1998), especially pp. 31-33, 53-56.

sense of self-righteousness. The ideal person is not self-righteous; she does not savor scorn. Ideals are also not for self-condemnation. The best of us should honor the part of glass that is full rather than obsessing in shame over the empty part.

5. I'M A UTILITARIAN; THEREFORE, I'M VIRTUOUS

"I accept utilitarianism; therefore, I'm virtuous." This statement is clearly false, and few dedicated utilitarians endorse it. Unfortunately, many come close to saying the contrapositive: Bob is not a utilitarian; therefore, Bob is not virtuous.

Philosophers sometimes condemn other philosophers as immoral because of theoretical disagreements. This is usually nonsense. I might be a Kantian and you are a utilitarian. We both agree that I should feed my cats. When asked why, both us recognize that the immediate reason that *I*—rather than you—should feed them is because they're my cats. I have a special relationship to them. I owe them my support. We also agree that the reason they should be feed at all has to do with the fact that they need food. When we step back in the attempt to explain these and other obligations, we construct different theories.

This disagreement over the best moral theory often does not entail a disagreement about morality itself.¹¹ Moral theories are a type of condensed representation of the moral. An expressionist painting, a Romantic painting, and photograph can be representations of the same desert landscape. Even if the expressionist, the Romanticist,

¹¹ See chapters 1, 3, and 10 of this volume. Also, cf. David Schmitz, *Elements of Justice* (Cambridge: Cambridge University Press, 2006), pp. 5-6, 227.

and the photographer disagree on the best way to convey the sublime in the landscape, they agree in an important way on what they're representing. The expressionist doesn't think he's representing a rainforest; the photographer doesn't believe it's a marsh.

Sometimes subscribing to a certain type of theory is immoral. I would say, at this point, it really does count against one's character to advocate dictatorship. At the limit, some theories end up not representing the moral at all, and the theorists subscribing to those theories lose sight of the good. A person whose moral theory calls for the maximization of pain doesn't really have a moral theory or a conception of morality. But these are exceptions.

Accepting a moral theory does not make us good. Theories are not that powerful. As far as we know, they cannot transform a person's character. But neither does *rejecting* a moral theory—even the correct one!—make us bad. Theories are not so straightforwardly bound up with moral practice that their absence renders us evil.

Similarly, no one gets points for having a "more demanding" theory. My theory may imply that the bare moral minimum required of everyone is moral perfection, while your theory leaves plenty of room for personal projects. Under normal conditions, having this more demanding conception of morality does not affect my character. It does not move me closer to virtue. The test of character is what one does, not what one argues one must do. Arguing for theory is too easy for it to count on behalf of one's character.

Maybe these points are obvious and no one disagrees. If so, there's still a point to explicitly stating that the thesis "One has to endorse moral theory X to be virtuous" is

false. Students believe this is what theorists think. If we do not think this, we should make that public.

6. THE LIMITS OF PRACTICALITY AREN'T SET BY THE THEORIES

Generally, the practicality of moral theory is an open empirical question. What moral theory is able to help us *do* cannot be settled by a priori armchair deduction. From the armchair, we can rule out a few hopeless theses. But the rest is subject to practical investigation, something at which philosophers are not particularly good. Can reading Kant help save marriages? Perhaps a creative counselor can put Kant's theory to that use.¹² We cannot answer the question by a priori analysis of the concept "moral theory".¹³

Thrasymachus challenges Socrates and claims that morality is for suckers. Socrates attempts to explain why morality is good for the soul, why the virtues benefit their possessor. One reason Socrates might do this is in the hope of improving Thrasymachus.

¹² Lou Marinoff, *Plato, Not Prozac* (New York: Harper Collins, 1999), pp. 123-125.

¹³ Many philosophers of science claim that one major lesson from Kuhn is that one cannot construct proper accounts of scientific method a priori. E.g., see Richard Giere, *Science without Laws* (Chicago: University of Chicago, 1999), pp. 53-55. I think this point applies to theories of method in all disciplines, including mathematics, music theory, and moral theory.

Suppose Thrasymachus sees a good moral theory, recognizes its authority, and concedes to Socrates that virtue in fact benefits its possessor. Would this concession be sufficient for Thrasymachus to become a good person? Thrasymachus might say, "Better to be an immoralist dissatisfied than a Socrates satisfied". Even if we show Thrasymachus that virtues really do benefit the possessor, this does not entail that he will be virtuous. Even if he *believes* and *understands* the beneficial nature of the virtues, this need not change his behavior. Perhaps he prefers to harm himself, or he prefers snide immoralism to eudaimonia.

Suppose we convince the immoralist of his error. Thrasymachus can believe he is in error but not change his ways. Convincing him to take this error seriously, and thus change his ways, is another issue.¹⁴ Even that is not enough. We might even convince him of the thesis, "You ought to take your errors seriously and change your ways". He still might not change his ways. Whether he changes depends deeply on certain psychological/motivational facts about him, not whether he affirms the truth of propositions like "I ought to change and I'm not kidding". Similarly, I can convince my students that they must study to pass exams, but that does not imply that they will actually study. I can further convince them of this proposition: "One must *take seriously* the thesis 'You must study to do well' and *actually study* in order to do well." They can believe that too, but that also does not imply they will study. When the time comes to choose between studying and partying, my cogent arguments may echo in their minds,

¹⁴ Cf. Peter Singer, "Famine, Affluence, and Morality", *Philosophy and Public Affairs* 1 (1972) 229-243, here p. 242.

but the temptations of partying win. My theory of what it takes to pass exams is not thereby shown incorrect, but the limits of its practicality are exposed.

Some philosophers claim such things are impossible. A person cannot generally believe X is right without being (at least somewhat) motivated to X. This view, *internalism*, has many defenders. I claim that it's *possible* for a person to believe X is right and not be motivated to X. Thus, some would label me an *externalist*. Not so. When I say it's "possible", I do not intend to be making a claim about psychological possibility. Rather, the operative modality here is epistemological possibility. My view is that *for all we know*—and we philosophers don't know—one can sincerely believe that X is right and not care.¹⁵ Whether human psychology is such that accepting a norm produces motivation is an open, empirical question, one which philosophers have little expertise to answer. Psychology may not even have much use for categories such as "believe", "accept", "desire", and "motivate".

Moral theorists are good at determining whether a theory has merit, whether there are good reasons to accept the theory. They can identify premises, evaluate them to some degree, and determine what degree of support the premises give the conclusion. Whether *truth*, argumentative *cogency* or *validity*, or other theoretical virtues translate into

¹⁵ Some psychological work seems to indicate that one can be a competent moral judge but lack any motivation to act morally. See Antonio R. Damasio, *Descartes' Error: Emotion, Reason, and the Human Brain* (New York: Grosset/Putnam, 1994), pp. 34-51, esp. p. 45.

behavior modification is an open question.¹⁶ It is something for philosophically informed psychologists, experimental economists, and educators to *discover*.

Here is a sketch of a test of the practicality of a theory. Experimental economists run games with real monetary stakes. The games often simulate various exchanges and types of relationships. We might run a game in which players have the possibility of acting on or ignoring moral reasons. A control group would play the game straight. An experimental group would play only after being taught a moral theory that explains the salient moral reasons. They could read a highly regarded article, with the article reinforced by a qualified teacher. We then check the game behavior of the theory-exposed players against the control group, looking to see whether exposure to theory increased (or decreased) the players' responsiveness to moral reasons.

7. THE THEORETICAL AND PRACTICAL AIMS OF THEORY MIGHT CONFLICT

It is often said that moral theories have both practical and theoretical aims.¹⁷ We hope these aims will cohere, i.e., satisfy one will at least not come at the cost of satisfy another. However, they very well could conflict.

¹⁶ James P. Sterba claims that his philosophical arguments are actually motivating. He claims that if you find someone with one of four conceptions of justice, his arguments will cause them to adopt a new conception and actually become just. See his *How to Make People Just* (Totowa, NJ: Rowman and Littlefield, 1988), p. 147.

¹⁷ E.g., Mark Timmons, *Moral Theory* (New York: Rowman and Littlefield, 2002), pp. 3-4.

Permit me certain liberties with *The Republic*. Let us say that Socrates theorizes about virtue and flourishing for both theoretical and practical reasons. Socrates wants to show Glaucon that virtue is necessary for flourishing, even though Glaucon sincerely declares that he will remain virtuous should virtue not be in his interest. For Glaucon, it is a purely theoretical question. He wants the truth. Socrates wants to convince Thrasymachus, but more importantly to change Thrasymachus' behavior and character. For Thrasymachus, the issue is practical.

There is no reason to suppose that the theoretical aim (of showing the connection) will succeed in producing the practical aim (of modifying Thrasymachus' behavior). Perhaps Thrasymachus is so perverted that these aims conflict. In order to change Thrasymachus' behavior, maybe Socrates must tell him a noble lie. Maybe Thrasymachus will be a better person only if he falsely believes some god will punish him for having bad attitudes. Perhaps that will motivate Thrasymachus to engage in a pattern of self-deception and brainwashing whereby he comes to value morality for its own sake.¹⁸ One can be brainwashed into true belief. Perhaps no amount of *convincing* and *argument* will modify Thrasymachus at all. Why think that we can talk him out of his immorality, even if we are right? Maybe Thrasymachus needs drugs and shock therapy. Maybe he will only change if we beat him up. To change Thrasymachus, we might even have to violate our own judgments and norms.

¹⁸ Similarly, Blaise Pascal argued that certain practices of self-deception and repetition could induce theism in the atheist.

The theoretical and practical aims of moral philosophy do not reduce to each other. They are not guaranteed to cohere in all cases. They could conflict. Success in one need not entail success in the other.

In the battle of ideas, most of us hope that truth will win. However, truth has to be marketed. A sound moral theory or theory of justice, to gain acceptance, often has to be conveyed with vigor and clarity. The idea must be in some respects timeless but presented in a timely manner. Bad ideas are also being marketed, perhaps better marketed. Henry Hazlitt once despaired over the economic ignorance and misconceptions held by laypeople and intellectuals alike. Bad ideas, such as "war is good for the economy", often win because they are "half-truths", attractive to one's conceit, and rewarding of intellectual laziness.¹⁹ Understanding usually requires more effort than *misunderstanding*. Accordingly, *misunderstanding* often wins.

The limits of the practicality of a moral theory are determined by 1) the complexity of the theory itself, 2) the mode of presentation, i.e., the rhetoric of the theory, and 3) the motivational set and psychological nature of the audience. Motivational and other psychological factors mediate between theory and practice. A false theory presented the right way might induce moral behavior. A true theory might not. We philosophers generally don't know.

Many philosophers have argued that utilitarianism is self-effacing. The crudest version of utilitarianism claims that we ought to maximize aggregate happiness in every

¹⁹ Henry Hazlitt, *Economics in One Lesson* (New York: Random House, 1979), pp. 18-19.

act. But suppose that *believing this thesis* does not maximize aggregate utility. Suppose that belief in utilitarianism, for some reason, turns out to be morally corrupting, such that utilitarians have worse character. Since bad character tends to have bad results, it would better maximize utility if no one believed utilitarianism. Does that show utilitarianism is false? No, it merely shows that by its own lights, we are morally obligated not to believe it. Utilitarianism is an account of the nature of the moral. The moral could be such that knowing its nature tends to make one less moral. Philosophers can provide good reasons for a particular conception of morality. Yet, what effect believing that conception will have depends on human psychology. Knowing the truth is not necessarily good for us.

Across campus from the philosophers, other professors teach music theory. Music theory nicely parallels moral theory. It is normative and evaluative. Theorists agree on many of the basic categories and judgments, but disagree over mid and high-level explanatory principles. The principles of good composition cannot be used as an algorithm for producing sublime music. But they do help us evaluate and understand music. This lack of an algorithm for writing masterpieces is not a count against music theory. Anyone who complained, for these reasons, that music theory is a worthless and practice primary probably has not taken any music theory.

Suppose learning music theory impeded one's playing or composition, i.e., one would be a better player or writer without it. Many blues guitarists claim that this is so. Still, if it were the case, music theory—a normative theory—would continue to have a point. It would remain crucial for *understanding* music, but not for *performing* it. It would become the domain of critics, but not musicians or composers. In contrast,

morality demands that we all be performers and not mere critics. If moral theory impeded performance, we would have moral reasons against producing it. But we would continue to have epistemic reasons for producing it. Moral reasons aim at the moral; epistemic reasons aim at the truth.

8. HOW VALUABLE ARE THEORIES? THEORISTS?

If the practical value of moral theory is an empirical question, is there any empirical evidence of its value? There is some, but it should be taken hesitantly.

The dominant theory of economic value in modern economics is marginalism. According to this theory, under normal, more or less competitive market conditions, the price of a good is its marginal utility, i.e., the value of that additional unit of the good to the consumers of the good. Price is function of demand at the margin, which itself is partly a function of the scarcity of the product. The principle subsumes wages; wages are the price of labor. (Applied to wages, the theory is called the marginal productivity theory of wages.)

We might consider what the price of a moral theorist's labor tells us about her practical value. Most professional philosophers are employed as university faculty. According to the College and University Personnel Association, the average salary for philosophy and religious studies faculty in 2004-2005 was \$61,470.²⁰ (Much better than

²⁰Scott Smallwood, "Faculty Salaries Increase at Faster Clip This Year, Survey Indicates", *The Chronicle of Higher Education*, Vol. 51:32 (04/15/2005), accessed online (last accessed 03/28/06).

what Socrates made.) The amount is higher when community college professors are excluded. These data measure all philosophers, not just moral theorists. We can make an informed assumption that the salaries of moral theorists do not diverge greatly from the norm. (The exception is within certain fields in applied ethics.) So perhaps we could estimate the practical value of the typical moral theorist at around \$61,000. This figure underestimates the total compensation of the average moral philosopher, since it does not measure additional non-salary benefits. (These typically add 25% to salary.)

This number doesn't differentiate between teaching, administrative duties, and the actual production of moral theories. Few theorists are paid simply to write theory. Most are required to teach to earn any income. However, tenure, promotions, and salary increases are predominantly tied to publishing. Producing moral theory, rather than teaching, tends to increase one's salary. The top stars in moral theory earn \$100,000 a year or more. The better a theory as measured by the profession, the more its author makes. A large percentage of income is tied directly to the production of published theory. Most professional philosophers, like most academics, publish very little.²¹ Most publications are written by a tiny minority of the field, and this minority is paid substantially more than the majority. Accordingly, the figure of \$61,000 is a low

²¹ For an account of stratification in productive ability in the natural sciences, see Derek J. De Solla Price, *Little Science, Big Science, and Beyond* (New York: Columbia University Press, 1986). For an account of the relevance of Price's findings to philosophy, see Randall Collins, *The Sociology of Philosophy* (Cambridge, MA: Harvard University Press, 1998), pp. 42-46.

estimate. The marginal productivity of active moral theorists—those that consistently produce and publish theory—is higher.

In addition, moral theories are often published in book form. John Rawls' *A Theory of Justice* has sold well over 250,000 copies since its 1971 release. It currently retails at \$24.95 on Amazon.com. Should we say the practical value of Rawls' *Theory* in book form has been millions of dollars? After all, the price is a measure of what the book is worth to consumers. Amazon sells Rawls-critic Robert Nozick's *Anarchy, State, and Utopia* for \$25.00. Does this mean that Nozick's work is slightly more valuable a theory-in-book-form per unit? Is Rawls' theory better because his total sales revenues are higher?

There are obvious objections to this way of measuring the practical value of a theory. Even if the moral philosopher makes \$61,500/year plus benefits, it's not clear what this number means. Above, we said that one way a theory could be practical is by providing us with moral guidance. Are theorists producing \$61,500/year on average in guidance? Let's say I am an assistant professor currently making \$45,000/year. I write *The Best Moral Theory Ever*. This book receives great accolades, and I am immediately recruited by the top Ph.D.-granting universities, promoted to associate and then full professor, and provided a salary of \$125,000/year. Does this mean that the value of the theory is \$80,000/year, plus book royalties? (Or does that measure the value of my working as a theorist?) Even if the answer is yes, this doesn't address the issues we presented above. It does not tell us whether the theory resolves disagreements, cures immorality, or whatnot.

The market for college professors is not quite like the market for engineers. Many college professors are government employees, and there are major difficulties in estimating the economic value of government employees. Education is paid in part through subsidies, donations, loans, and so on. It's not quite like a typical good. Whatever salary is a measure of, the information contained the price is distorted.

Most importantly, some prices could not be said to reflect practical value. A price normally measures how much the marginal buyer values a good. But prices do not tell us in what way the good is valued. A Van Gogh painting might sell for \$2 million, but that does not indicate its practical value. The Van Gogh's monetary worth reflects something different, perhaps its value (and/or scarcity) as an object of aesthetic enjoyment or contemplation. Economists would feel uncomfortable differentiating between such categories as practical and contemplative value. Given what economists are trying to do, they have good reasons for such discomfort. But given what we are trying to do, we have good reasons to make the distinction.

One key insight of economics is that market prices bear and convey information about what people value. Market prices must tell us something about the value of moral theorizing. What they tell us is not at all clear.

9. COULD MORAL THEORY BE A FAILED ENTERPRISE?

I argued that the practical value of moral theory is largely an empirical issue. We philosophers don't know what its value is. But suppose we do the experiments, determine what the market prices mean, etc., and determine that moral theories have no practical value. Does this entail that moral theorizing is a failed enterprise?

Some might claim that theories *must* be practical. Morality, after all, is practical. Claiming necessity is a mistake. Moral theories are primarily for explaining morality, just as music theory explains music. Applied music theory helps us perform. Perhaps applied moral theory could help us perform as well, but most people perform without it. Folk musicians worldwide produce complex, beautiful music. They acquire norms from imitation, feel, and intuition. Many of these norms have evolved in a mixed Lamarckian/Darwinian fashion. Musicians develop a good ear, spontaneity, and rhythm. They exercise and sharpen their innate faculties. They learn discipline. These folk musicians know how to perform musically, and they have reasons for playing the way they do. Though they do not have theory, they get on fine without it. Music theory is the domain of the few. All of this holds for moral theory as well. If so, then is it really necessary that moral theory be practical?

Suppose someone claims to have discovered a test for the acceptability of a moral theory. He claims to have isolated a trait, P, necessary for a theory to be good. Any theory lacking P is worthless. Excited, we ask him how different theories fare. He tells us that he has tested all theories against P, and *all theories fail*. Perplexed, we ask him what P is. He answers, "Practicality."

When someone proposes a test of moral theory, but all theories fail that test, we have a choice to make. We can decide that all theories are bad or we can decide that the test itself is bad.

In my view, there is a test of whether having P is a *genuine test*, regardless of whether P is practicality, simplicity, fecundity, or something else. The purpose of having

tests is to measure performance against some standard. More specifically, tests sort good from bad. A test that nothing does or could pass is no test at all. It might indicate that what we thought was an obvious test is not. If theories cannot pass a certain test, say, practicality, this raises the question of what test they can pass. The worth of moral theory requires that some theories pass some tests that are themselves worth passing.

Theories of all types are attempts to solve problems. (A lack of understanding is the typical problem.) Solving certain sorts of problems, then, is a typical test of a theory. Many people want moral theories to solve practical problems. However, some of the problems we hope a theory will solve may turn out to be pseudo-problems, i.e., they are not real problems at all. (To say someone has a pseudo-problem means that she mistakenly believes she has a problem.) Other problems we hope to solve by one type of theory may turn out to be solved by another type of theory.

Here are some historical and scientific examples. Citizens of Salem, Massachusetts, once worried about their witch problem. Some may have wanted a witch removal theory to solve it. But witches do not exist. There was no witch problem, and so solving the witch problem isn't a test of *any* theory. In the Middle Ages, scientists thought that comets were atmospheric disturbances, and thus thought meteorology should explain them. They were wrong in holding the ability to explain comets as a test of meteorological theories; it was a job for astronomy.²² Similarly, we may want a set of explanatory principles that convert the vicious to virtue. But if no theory yet has

²² Larry Laudan, *Progress and Its Problems* (Berkeley: University of California Press, 1977), p. 21.

succeeded, it may well be that we're wrong in thinking it's the job of a moral theory to do so.

Many former philosophical problems are now best seen as pseudoproblems. Philosophers and theologians spent thousands of years attempting to explain the existence of evil in a world governed by an omnibenevolent, omnipotent God. Leibniz claimed that while evil exists, this is the best of all possible worlds. Spinoza claimed that evil does not exist; we are mistaken in believing it. However, the best current ontologies do not include omnibenevolent, omniscient deities. Pain needs (and has) a biological explanation. From a contemporary secular point of view, there is no theological problem of evil, because there are no theological problems, period.

If a moral theory turns out to be practical, then that gives us some reason to hold other, competitor theories should be practical. Theories compete via their virtues. If no theory of type M exhibits a particular virtue, we cannot be sure that this virtue really belongs to type M theories.

Perhaps moral theories were invented with the intention of being practical. Again, suppose they are not practical. Are they failures? Post-It glue and Teflon were considered failures because they did not do what their inventors intended. Sometimes we invent things for one reason, but keep them for another. The *contemplation of the nature of morality* apart from the end of actual moral improvement may be an end in itself.

What *theoretical* virtues do we know moral theories have? This merits a long discussion, beyond the scope of this paper. But there is a tentative answer: Reflect not just on your favorite moral theory, but all of those you have encountered. What have you

learned from them? How much understanding have you gained? The virtue of moral theories is that for most of us, the answer is quite a lot.

REFERENCES

- Aristotle, *Physics* (Cambridge, MA: Harvard University Press, 1986).
- Audi, Robert, *The Good in the Right* (Princeton: Princeton University Press, 2004).
- Ayer, A. J. (ed.) *Logical Positivism*, (Glencoe, IL: The Free Press, 1959).
- Bargh, John A., and Chartrand, Tanya L. "The Unbearable Automaticity of Being", *American Psychologist* 54:7, 462-479.
- Becker, Gary, *The Economic Approach to Human Behavior* (Chicago: University of Chicago Press, 1978).
- Beckert, Jens, *Beyond the Market* (Princeton: Princeton University Press, 2002).
- Brennan, Jason, "Modesty without Illusion", *Philosophy and Phenomenological Research* (forthcoming).
- Brink, David O., "Utilitarian Morality and the Personal Point of View", *The Journal of Philosophy* 83:8 (August 1986), pp. 417-438.
- Brown, Harold I., *Rationality* (New York: Routledge, 1990).
- Bunge, Mario, *The Methodological Unity of Science* (Cambridge, MA: Springer, 1973).
- Carmo, José, and Brown, Mark (eds.), *Deontic Logic, Agency, and Normative Systems* (Berlin: Springer Verlag, 1996).
- Cartwright, Nancy, *How the Laws of Physics Lie* (Oxford: Oxford University Press, 1983)
- Cartwright, Nancy, *Nature's Capacities and their Measurement* (Oxford: Oxford University Press, 1994).
- Cartwright, Nancy, *The Dappled World* (Cambridge: Cambridge University Press, 1999)
- Cartwright, Nancy, "In Favor of Laws that Are Not so *Ceteris Paribus* after all", *Erkenntnis* (2002) 57: 425-439.
- Clarke, Stanley G. and Evan Simpson, "Introduction", in (ed.) Stanley G. Clark and Evan Simpson, *Anti-Theory in Ethics and Moral Conservatism*, (Albany: SUNY Press, 1984).

- Clarke, Stanley G. and Evan Simpson (eds.), *Anti-Theory in Ethics and Moral Conservatism* (Albany: SUNY Press, 1984).
- Cohen, Joshua, "The Arc of the Moral Universe," *Philosophy and Public Affairs* 26:2 (1997): 91-134.
- Collins, Randall, *The Sociology of Philosophies* (Cambridge, MA: Harvard University Press, 1998)
- Copp, David, and Zimmerman, David (eds.), *Morality, Reason, and Truth* (Totowa: Rowman and Allenheld), p. 297-219.
- Cowen, Tyler, *In Praise of Commercial Culture* (Cambridge, MA: Harvard University Press, 2000).
- Cowen, Tyler, *Creative Destruction* (Princeton: Princeton University Press, 2002).
- Cowen, Tyler, *Good and Plenty: The Creative Successes of American Arts Funding* (Princeton: Princeton University Press, 2006).
- CUPA, *National Faculty Salary Survey, Executive Summary* (2005).
- Damasio, Antonio R., *Descartes' Error: Emotion, Reason, and the Human Brain* (New York: Grosset/Putnam, 1994).
- Dancy, Jonathan, *Ethics without Principles* (Oxford: Oxford University Press, 2004).
- Daniels, Norman, "Wide Reflective Equilibrium and Theory Acceptance in Ethics", *The Journal of Philosophy* 76:5 (1979), pp. 256-282.
- Darwall, Stephen, "How Should Ethics Relate to (the Rest of) Philosophy?: Moore's Legacy", *The Southern Journal of Philosophy* (2003), vol. XLI, pp. 1-29.
- Davidson, Donald, "What Metaphors Mean," *On Metaphor*, ed. Sheldon Sachs (Chicago: University of Chicago Press, 1978).
- Dummett, Michael, *Truth and Other Enigmas* (Cambridge, MA: Harvard University Press, 1978).
- Edidin, Aron, "Philosophy: Just Like Science, Only Different", *Philosophy and Phenomenological Research* 45:4 (1985), 537-552.
- Feinberg, Joel, *Doing Philosophy: A Guide to the Writing of Philosophy Papers* (Belmont, CA: Wadsworth, 2002).

Friedman, Milton, "The Methodology of Positive Economics", in Daniel M. Hausman (ed.), *The Philosophy of Economics* (Cambridge: Cambridge University Press, 1984).

Fuller, Lon L., "The Case of the Speluncian Explorers", *Harvard Law Review* (1949) 62: 615-645.

Gauthier, David, *Morals by Agreement* (Oxford: Oxford University Press, 1986).

Geach, Robert, *General Relativity from A to B* (Chicago: University of Chicago Press, 1978).

Giere, Ronald N., *Science without Laws* (Chicago: University of Chicago Press, 1999).

Gert, Bernard, *Morality: Its Nature and Justification* (Oxford: Oxford University Press, 1998).

Gibbard, Alan, and Varian, Hal R., "Economic Models", *The Journal of Philosophy* 75:11 (1978) 664-677

Gödel, Kurt, "A Remark about the Relationship between the Theory of General Relativity and Idealistic Philosophy", *Collected Works: Publications 1948-1974* (Oxford: Oxford University Press, 2001 [1949]), pp. 202-207.

Greene, Joshua, *The Terrible, Horrible, No Good, Very Bad Truth about Morality and What to Do About It*, Department of Philosophy, Princeton University, 2002, Ph.D. dissertation.

Greene, Joshua, "The Secret Joke of Kant's Soul," in Walter Sinnott-Armstrong (ed.), *The Psychology and Biology of Morality* (Cambridge, MA: MIT Press, Forthcoming).

Greene, Joshua, "Reply to Mikhail and Timmons", in Walter Sinnott-Armstrong (ed.), *The Psychology and Biology of Morality* (Cambridge, MA: MIT Press, forthcoming)

Greene, Joshua D., Sommerville, R. Brian, Nystrom, Leigh E., Darley, John M., and Cohen, Jonathan D., "An fMRI Investigation of Emotional Engagement in Moral Judgment", *Science* 293:14 (2001) 2105-2107.

Hacking, Ian (ed.), *Scientific Revolutions*, (Oxford: Oxford University Press, 1981).

Hacking, Ian, *Representing and Intervening* (Cambridge: Cambridge University Press, 1983).

- Haidt, Jonathan, "The Emotional Dog and Its Rational Tail: A Social Intuitionist Approach to Moral Judgment", *Psychological Review* 108 (2001) 814–834
- Haidt, Jonathan, Björklund, Fredrik, and Murphy, Scott, "Moral Dumbfounding: When Intuition Finds No Reason", Department of Psychology, University of Virginia, 2004, unpublished manuscript.
- Hare, R. M., *The Language of Morals* (Oxford: Oxford University Press, 1952).
- Hausman, Daniel M. (ed.) *The Philosophy of Economics* (Cambridge: Cambridge University Press, 1984).
- Hawking, Stephen, "Gödel and the End of Physics", Lecture: Dirac Centennial Celebration, Cambridge: University of Cambridge, 2002.
- Hayek, F. A., *The Constitution of Liberty* (Chicago: University of Chicago Press, 1960).
- Hazlitt, Henry, *Economics in One Lesson* (New York: Random House, 1979).
- Hill, Jr., Thomas E., "Ideals of Human Excellence and Preserving Natural Environments", *Environmental Ethics* 5 (1983): 211-224.
- Hooker, Brad and Margaret Little (eds.), *Moral Particularism* (Oxford: Oxford University Press, 2000).
- Hooker, Brad, *Ideal Code, Real World* (Oxford: Oxford University Press, 2001).
- Horgan, Terence, and Timmons, Mark, "Troubles for Michael Smith's Metaethical Rationalism", *Philosophical Papers* XXV:3 (1996), pp. 203-231.
- Hume, David, *Enquiries Concerning the Principles of Morals*, (eds.) L. A. Selby-Bigge and P. H. Nidditch (New York: Oxford University Press, 1975).
- Hume, David, *Dialogues Concerning Natural Religion*, (Indianapolis: Hackett, 1980).
- Hume, David, *A Treatise of Human Nature*, ed. David Fate Norton and Mary J. Norton (Oxford: Oxford University Press, 2000 [1739]).
- Hursthouse, Rosalind, *On Virtue Ethics* (New York: Oxford University Press, 1999).
- Jackson, Frank, Pettit, Phillip, and Smith, Michael, "Ethical Particularism and Patterns", in *Moral Particularism*, (eds.) Brad Hooker and Margaret Littlefield (Oxford: Oxford University Press, 2000).

Jonsen, Albert R. and Toulmin, Stephen, *The Abuse of Casuistry* (Berkeley: University of California Press, 1988).

Kagan, Shelly, *The Limits of Morality* (New York: Oxford University Press, 1991).

Kamm, Frances, "High Theory, Low Theory, and the Demands of Morality", *Nomos* XXXVII (1995), 81-107.

Kant, Immanuel, *The Groundwork of the Metaphysics of Morals*, in Mary J. Gregor (ed., trans.), *Practical Philosophy* (Cambridge: Cambridge University Press 1996 [1785]).

Kant, Immanuel, *The Metaphysics of Morals*, trans. Mary J. Gregor (Cambridge: Cambridge University Press, 1996 [1797]).

Kavka, Gregory, "The Reconciliation Project", in David Copp and David Zimmerman (eds.), *Morality, Reason, and Truth* (Totowa: Rowman and Allenheld), p. 297-219.

Kitcher, Philip, "Explanatory Unification and the Causal Structure of the World," *Minnesota Studies in the Philosophy of Science, Volume XIII*, edited by P. Kitcher and W. Salmon (Minneapolis, University of Minnesota Press, 1990).

Klee, Robert, *Introduction to the Philosophy of Science* (New York: Oxford University Press, 1997).

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

Kuhn, Deanna, *The Skills of Argument* (New York: Cambridge University Press, 1991).

Kuhn, Thomas S., *The Structure of Scientific Revolutions, 3rd Edition* (Chicago: University of Chicago Press, 1996 [1962]).

LaFollette, Hugh (ed.), *The Blackwell Guide to Ethical Theory* (Malden, MA: Blackwell, 2000).

Lakatos, Imre, and Musgrave, Alan (eds.), *Criticism and the Growth of Knowledge* (New York: Cambridge University Press, 1970),

Lakatos, Imre, *The Methodology of Scientific Research Programs, Vol. I* (Cambridge: Cambridge University Press, 2000).

Lakoff, George, and Johnson, Mark, *Metaphors We Live By* (Chicago: University of Chicago Press, 1980).

Larmore, Charles E., *Patterns of Moral Complexity* (New York: Columbia University Press, 1987).

Laudan, Larry, *Progress and Its Problems* (Berkeley: University of California Press, 1977).

Laudan, Larry, "A Problem Solving Approach to Scientific Progress", in Ian Hacking (ed.), *Scientific Revolutions*, (Oxford: Oxford University Press, 1981).

Laudan, Larry, *Science and Values* (Berkeley: University of California Press, 1984).

Levitt, Steven D. and Duggan, Mark, "Winning Isn't Everything: Corruption in Sumo Wrestling", *American Economic Review* 92:4 (2002), pp. 1594-1605.

Lehrer, Keith, *Theory of Knowledge*, 2nd Edition (Boulder: Westview Press, 2000).

Lomasky, Loren, "Nozick's Libertarian Utopia", in David Schmidtz (ed.), *Robert Nozick* (Cambridge: Cambridge University Press, 2002).

Louden, Robert B., *Morality and Moral Theory* (New York: Oxford University Press, 1992).

Marinoff, Lou, *Plato, Not Prozac* (New York: Harper Collins, 1999).

Matteson, Jason, *Aspiring to Advice*, Department of Philosophy, University of Arizona, Ph.D. dissertation (in progress).

McCloskey, Donald N., *If You're So Smart* (Chicago: University of Chicago Press, 1990).

McDowell, John, *Mind, Value, and Reality* (Cambridge, MA: Harvard University Press, 2001).

McNamara, Paul, "Must I Do What I Ought? (Or Will the Least I Can Do Do?)", in José Carmo and Mark Brown (eds.), *Deontic Logic, Agency, and Normative Systems* (Berlin: Springer Verlag, 1996), pp. 154-173.

McNaughton, David, *Moral Vision* (Oxford: Blackwell, 1988).

Mill, John Stuart, *On Liberty* (Indianapolis: Hackett, 1978 [1859]).

Mill, John Stuart, *Utilitarianism* (Indianapolis: Hackett, 2002 [1861]).

Miller, David, *Principles of Social Justice* (Cambridge, MA: Harvard University Press, 1999).

- Millgram, Elijah, "How to Make Something of Yourself", in David Schmidtz (ed.), *Robert Nozick* (New York: Cambridge University Press, 2002).
- Mlodinow, Leonard, "Prometheus Bound", *Discover* 28:10 (October, 2005), 64-65.
- Moore, G. E., *Principia Ethica* (Buffalo: Prometheus Books, 1988 [1902]).
- Neurath, Otto von, "Protocol Sentences", in A. J. Ayer (ed.) *Logical Positivism*, (Glencoe, IL: The Free Press, 1959), pp. 199-208.
- Nietzsche, Friedrich, *On the Genealogy of Morals* (New York: Vintage, 1989 [1887]).
- Norton, Brian G., *Toward Unity among Environmentalists* (Oxford: Oxford University Press, 1991).
- Nozick, Robert, *Anarchy, State, and Utopia* (New York: Basic Books, 1974).
- Nozick, Robert, *Philosophical Explorations* (Cambridge, MA: Belknap Press, 1981).
- Nozick, Robert, *The Examined Life* (New York: Simon and Schuster, 1989).
- Nozick, Robert, *The Nature of Rationality* (Princeton: Princeton University Press, 1993).
- Nozick, Robert, *Socratic Puzzles* (Cambridge, MA: Harvard University Press, 1999).
- Nozick, Robert, *Invariances* (Cambridge, MA: Belknap Press, 2001).
- Nussbaum, Martha, *The Fragility of Goodness* (New York: Cambridge University Press, 1986).
- Nussbaum, Martha, *Upheavals of Thought* (New York: Cambridge University Press, 2001).
- Paine, Lynn S., "Managing for Organizational Integrity", *Harvard Business Review* 72:2 (1994), 106-117.
- Pitkin, Walter B., "Is Agreement Desirable?", *Journal of Philosophy, Psychology, and Scientific Method* 9:26 (1912), 711-715.
- Plato, *Apology*, trans. G. M. A. Grube (Hackett: Indianapolis, 1981).
- Pojman, Louis P., *Ethics: Discovering Right and Wrong* (Belmont, CA: Wadsworth, 1999).

- Posner, Richard A. *The Problematics of Moral and Legal Theory* (Cambridge, MA: Harvard University Press, 1999).
- Price, Derek J. De Solla, *Little Science, Big Science, and Beyond* (New York: Columbia University Press, 1986).
- Prichard, H. A., "Does Moral Philosophy Rest on a Mistake", *Mind* XXI:81 (1912), 21-37
- Putnam, Hilary, *Ethics without Ontology* (Cambridge, MA: Harvard University Press, 2004).
- Rauscher, F. H., Shaw, G. L., and Ky, K. N., "Music and Spatial Task Performance", *Nature* 365 (1993), 611.
- Rawls, John, "Outline of a Decision Procedure in Ethics", *Philosophical Review* 60:2 (1951), 177-197.
- Rawls, John, *A Theory of Justice* (Cambridge, MA: Harvard University Press, 1971).
- Rawls, John, *Political Liberalism* (New York: Columbia University Press, 1996).
- Rawls, John, *Justice as Fairness: a Restatement* (Cambridge, MA: Harvard University Press, 2001).
- Rescher, Nicholas, *The Strife of Systems* (Pittsburgh: University of Pittsburgh Press, 1985).
- Rescher, Nicholas, *Philosophical Standardism* (Pittsburgh: University of Pittsburgh Press, 1994).
- Rescher, Nicholas, *Philosophical Dialectics* (Albany: SUNY Press, 2006).
- Richards, I. A., *The Philosophy of Rhetoric* (New York: Oxford University Press, 1936).
- Rose, Carol, "Possession as the Origin of Property", *University of Chicago Law Review* 52 (1985) 73-88.
- Ross, Steven L., "Taking Posner Seriously", *The Philosophical Forum* XXXII:1 (2001) 1-23.
- Ross, W. D., *The Right and the Good* (Indianapolis: Hackett, 1988 [1930]).

Sandel, Michael, *Liberalism and the Limits of Justice*, 2nd Edition (New York: Cambridge University Press, 1998),

Sayre-McCord, Geoff (ed.), *Essays on Moral Realism*, (Ithaca, NY: Cornell University Press, 1988).

Sayre-McCord, Geoff, "Introduction: The Many Moral Realisms", in Geoff Sayre-McCord (ed.), *Essays on Moral Realism* (Ithaca: Cornell University Press, 1988), pp. 1-26.

Sayre-McCord, Geoff, "Hume and the Bauhaus Theory of Ethics", *Midwest Studies in Philosophy* 20 (1996), 280-298.

Scanlon, T. M., "The Aims and Authority of Moral Theory," *Oxford Journal of Legal Studies*, 12:1 (1992), 1-21.

Scanlon, T. M., *What We Owe to Each Other* (Cambridge, MA: Belknap Press, 1998).

Schedler, George, "Does Ethical Meat Eating Maximize Utility?", *Social Theory and Practice* 31:4 (2005), 499-511.

Schmidtz, David, *Rational Choice and Moral Agency* (Princeton: Princeton University Press, 1995).

Schmidtz, David, "Taking Responsibility" in David Schmidtz and Robert Goodin, *Social Welfare and Individual Responsibility* (Cambridge: Cambridge University Press, 1998).

Schmidtz, David, "The Meanings of Life", in David Schmidtz (ed.), *Robert Nozick* (New York: Cambridge University Press, 2002).

Schmidtz, David, (ed.), *Robert Nozick* (Cambridge: Cambridge University Press, 2002).

Schmidtz, David, *Elements of Justice* (New York: Cambridge University Press, 2006).

Schmidtz, David, "When Justice Matters", *Ethics* (forthcoming 2007).

Schmidtz, David and Goodin, Robert, *Social Welfare and Individual Responsibility* (Cambridge: Cambridge University Press, 1998).

Schmitt, Carl, *The Concept of the Political*, trans. George Schwab (Chicago: University of Chicago Press, 1996 [1932]).

Schumpeter, Joseph A., *Capitalism, Socialism, and Democracy* (New York: Harper Perennial, 1962).

Shackel, Nicholas, "On the Vacuity of Postmodern Methodology", *Metaphilosophy* 36:3 (2005), 295-320.

Shakespeare, William, *The Complete Works of Shakespeare*, Stanley Wells and Gary Taylor (eds.) (New York: Oxford University Press, 1988)

Sidgwick, Henry, *The Methods of Ethics* (Indianapolis: Hackett, [1907] 1981).

Simon, Herbert A., *Models of Thought* (New Haven: Yale University Press, 1979).

Singer, Peter, "Moral Experts", *Analysis* 32 (1972), 115-117.

Singer, Peter, "Famine, Affluence, and Morality", *Philosophy and Public Affairs* 1 (1972) 229-243.

Singer, Peter, *Practical Ethics* (New York: Cambridge University Press, 1979).

Singer, Peter, "Ethics and Intuitions", *The Journal of Ethics* 9:3-4 (2005), 331-352.

Sinnott-Armstrong, Walter, and Timmons, Mark (eds.), *Moral Knowledge?* (New York: Oxford University Press, 1998), pp. 293-325.

Sinnott-Armstrong, Walter (ed.), *The Psychology and Biology of Morality* (Cambridge, MA: MIT Press, forthcoming).

Sorrell, Tom, *Moral Theory and Anomaly* (London: Blackwell, 2000).

Smith, Michael, *The Moral Problem* (Malden, MA: Blackwell, 1994).

Sterba, James P., *How to Make People Just* (Totowa, NJ: Rowman and Littlefield, 1988).

Sturgeon, Nicholas, "Moral Explanations" in Geoff Sayre-McCord (ed.), *Essays on Moral Realism*, (Ithaca, NY: Cornell University Press, 1988).

Timmons, Mark, "Foundationalism and the Structure of Ethical Justification", *Ethics* 97:3 (1987), pp. 595-609.

Timmons, Mark, "Outline of a Contextualist Moral Epistemology", in Walter Sinnott-Armstrong and Mark Timmons (eds.), *Moral Knowledge?* (New York: Oxford University Press, 1998), pp. 293-325.

Timmons, Mark, *Morality without Foundations* (Oxford: Oxford University Press, 1999).

Timmons, Mark, *Moral Theory* (New York: Rowman and Littlefield, 2002).

Tommasi, Mariano and Ierulli, Kathryn (eds.), *The New Economics of Human Behavior* (New York: Cambridge University Press, 1995).

Toulmin, Steven, "Does the Distinction between Normal and Revolutionary Science Hold Water?", in Imre Lakatos and Alan Musgrave (eds.), *Criticism and the Growth of Knowledge* (New York: Cambridge University Press, 1970).

Urban, Wilbur M., "Progress in Philosophy in the Last Quarter Century", *The Philosophical Review* 35:2 (1926), 93-123.

Waldron, Jeremy, "What Plato Would Allow", *NOMOS XXXVII* (1995), 138-180.

Walker, Margaret Urban, *Moral Contexts* (New York: Rowman and Littlefield, 2003).

Walzer, Michael, *Spheres of Justice* (New York: Basic Books, 1981).

Weintraub, Roy E., *Microfoundations: The Compatibility of Microeconomics and Macroeconomics* (New York: Cambridge University Press, 1979).

Westen, D., Kilts, C., Blagov, P., Harenski, K., and Hamann, S., "The neural basis of motivated reasoning: An fMRI study of emotional constraints on political judgment during the U.S. Presidential election of 2004", *The Journal of Cognitive Neuroscience* (forthcoming 2007).

Williams, Bernard, *Moral Luck* (Cambridge: Cambridge University Press, 1981).

Williams, Bernard, *Ethics and the Limits of Philosophy* (Cambridge, MA: Harvard University Press, 1985).

Williams, Bernard, *Making Sense of Humanity* (New York: Cambridge University Press, 1995).

Woodward, Jim, "There Is No Such Thing as a *Ceteris Paribus* Law", *Erkenntnis* 57 (2002), 303-328.

Wright, Robert, *The Moral Animal* (New York: Vintage, 1994).