

2007

The genealogy of morals: contemporary empirical accounts

Franklin Donald Worrell

Louisiana State University and Agricultural and Mechanical College

Follow this and additional works at: https://digitalcommons.lsu.edu/gradschool_theses



Part of the [Arts and Humanities Commons](#)

Recommended Citation

Worrell, Franklin Donald, "The genealogy of morals: contemporary empirical accounts" (2007). *LSU Master's Theses*. 1536.

https://digitalcommons.lsu.edu/gradschool_theses/1536

This Thesis is brought to you for free and open access by the Graduate School at LSU Digital Commons. It has been accepted for inclusion in LSU Master's Theses by an authorized graduate school editor of LSU Digital Commons. For more information, please contact gradetd@lsu.edu.

THE GENEALOGY OF MORALS: CONTEMPORARY EMPIRICAL ACCOUNTS

A Thesis

Submitted to the Graduate Faculty of the
Louisiana State University and
Agricultural and Mechanical College
in partial fulfillment of the
requirements for the degree of
Master of Arts

in

The Department of Philosophy and Religious Studies

by

Franklin Donald Worrell

B.A., Louisiana State University, 2005

B.A., Louisiana State University, 2005

May 2007

Acknowledgments

I want to offer my utmost thanks to both Jen O’Conner and Margaret Toups, who facilitated and encouraged this project more than they know. Mrs. O’Conner provided constant emotional and technical support far above and beyond the call of duty throughout my entire career at LSU. Mrs. Toups was always ready with a suggestion that proved critical and with a smile. I thank these two women from the bottom of my heart.

It is impossible to look back on my philosophical career at LSU without noting the profound influence of two gentlemen: Husain Sarkar and Jon Cogburn. Jon Cogburn is single-handedly responsible for getting me hooked on philosophy—I still look back on his Honors Introduction to Philosophy and smile. Husain Sarkar taught me what it was to truly understand a text and made me believe that I could accomplish a long research project, which I proved with his help when I completed my undergraduate honors thesis. Sincerest thanks and deepest esteem to both these men.

I want to acknowledge the profound impact my dear friends have had on my career. Dustin, Miranda, Marty, Jenni, Angela, and Monique have listened to countless lectures on my ideas and have paid the utmost attention, even at the most inconvenient of times. They have encouraged me and motivated me at crucial moments. However, the greatest influence on my life and the people most deserving of thanks are my parents: Frank and Patty. They encourage me constantly, and I wish at least to acknowledge their influence, love, and aid.

Last but not least, I wish to thank the members of my committee, Ed Song, Greg Schufreider, and Jon Cogburn, for committing time and resources to this project. I especially want to thank Alison Denham for her devotion that made the present work possible.

Deepest and sincerest thanks to all!

Table of Contents

ACKNOWLEDGMENTS.....	ii
ABSTRACT.....	iv
CHAPTER 1 INTRODUCTION.....	1
1.1 Moral Realists and Moral Evolution.....	2
1.2 The Coming Terrain.....	3
1.3 Slote's Argument Regarding Unidirectionality.....	5
1.4 A Note on the Varieties of Realism.....	6
1.5 Progress or Evolution?.....	10
CHAPTER 2 MORAL REALISM AND THE ARGUMENT FROM MORAL EVOLUTION..	12
2.1 Overview.....	12
2.2 Sturgeon <i>Et Al</i>	12
2.3 Railton's Social Unrest Analysis.....	22
2.4 Conclusion.....	27
CHAPTER 3 NICHOLS AND THE AFFECTIVE RESONANCE ACCOUNT.....	28
3.1 Introduction.....	28
3.2 Nichols and the Sentimental Rules Account.....	28
3.3 Nichols' Metaethical Commitments.....	31
3.4 An Excursion into Epidemiology.....	34
3.5 Patterns of Moral Evolution.....	40
3.6 The Best Explanation (According to Nichols).....	42
3.7 The Attack on Realism.....	46
3.8 Conclusion.....	50
CHAPTER 4 CONCLUSION.....	51
4.1 Introduction: So Who Won?.....	51
4.2 The First Use and the Charge of Circularity.....	51
4.3 The Second Use and Lack of Content.....	55
4.4 Directions for Future Research.....	59
4.5 Conclusion: What about Moral Evolution/Progress?.....	60
WORKS CITED.....	61
VITA.....	63

Abstract

In the late twentieth century, moral realists began to resurrect a type of argument that emerged during the Enlightenment. These realists appealed to moral progress as evidence for moral facts, and their arguments took the form of inferences to the best explanation. Recently, the argument style has emerged again. This time, the inference to the best explanation is being used by empirically-informed sentimentalists to argue that their theories can provide accounts of moral evolution that have greater explanatory and predictive power than the accounts offered by the moral realists. This thesis examines the arguments to the best explanation of such moral realists as Nicholas Sturgeon, Michael Slote, Michael Smith, Peter Singer, and Thomas Nagel. The views of these moral realists are confronted with the substantial empirical evidence provided by Shaun Nichols to bolster his Sentimental Rules account, which is a variety of sentimentalism. Nichols attempts to expand epidemiological approaches to cognitive anthropology to accommodate his research on affect-backed norms. I elucidate Nichols' research and his own inference to the best explanation of the data he examines as well as his attack on the accounts provided by the moral realists.

After examining this substantive piece of the debate over what actually counts as the best explanation of moral evolution, I argue that the inference to the best explanation is actually being employed in two distinct uses by these theorists. The first use presupposes a metaethical thesis regarding the nature of moral facts and renders the inference circular. The majority of the moral realists examined employ the inference in this fashion. The second use is not circular, but leaves the theorist with a very restricted ability to fill out the content of moral beliefs and moral facts.

Chapter 1

Introduction

It is evident that we are at a primitive stage of moral development. Even the most civilized human beings have only a haphazard understanding of how to live, how to treat others, how to organize their societies. The idea that the basic principles of morality are known, and that the problems all come in their interpretation and application, is one of the most fantastic conceits to which our conceited species has been drawn. (The idea that if we cannot easily know it, there is no truth here is no less conceited.) Not all of our ignorance in these areas is ethical, but a lot of it is. And the idea of the possibility of moral progress is an essential condition of moral progress. None of it is inevitable. (Nagel 186)

At the out, my choice of this passage from Thomas Nagel's *The View from Nowhere* to open this thesis may rightly seem mysterious. My goal is to examine a specific type of argument that attempts to explain moral progress. Arguments of this type attempt to offer the most powerful or best explanations of moral progress. Such arguments to the best explanation are used to bolster metaethical theories, most notably moral realism. These arguments have received renewed attention from a group of empirically informed 'sentimentalists', who have attempted to best the explanations provided by moral realists.

The present work will examine the arguments as presented, first, by a few notable moral realists and, second, by Shaun Nichols in his 'Sentimental Rules' account of moral judgment. Nichols presents new anthropological evidence regarding the genealogy of morality, and he hopes that his theory will provide a more plausible explanation of moral evolution. My goal is not to condemn moral realism, nor is it to attack the analysis offered by Nichols. Instead, I want to focus my attention on the type of argument being presented by the two sets of theorists. When the evidence is all in, we will be in a position to ask whether or not it is a legitimate tactic for either school to apply the inference to the best explanation as an attack on its opponents. After the nature of the debate is framed, I hope to show that inferences to the best explanation are

being used in two ways by moral philosophers. I maintain the first use is circular and those employing the inference to the best explanation in this fashion are begging the question. The second use is not susceptible to the charge of question-begging, but it leaves us in a poor position to say anything regarding the substantive content of moral facts and beliefs.

1.1 Moral Realists and Moral Evolution

In the latter half of the twentieth century moral realists resurrected an argument that emerged during the Enlightenment. They appealed to moral evolution as empirical, factual evidence supporting moral realism. Their arguments took the form of inference to the best explanation. An inference to the best explanation is a style of argument that examines some set of data or findings and presents a theory to explain the data. If this theory explains the data better than any rival theory, then it is deemed more successful and explanatorily powerful. Thus, an inference to the best explanation tries to establish two things: that the theory in question actually explains a set of data and that the theory provides a better explanation than any other plausible theory. The moral realists claimed that the truth of moral facts could explain the evolution of morality and the historical trend for moral communities to become more inclusive. One task of the present work is to bring contemporary anthropological research regarding the genealogy of morality to bear upon such arguments to the best explanation.

Shaun Nichols in his *Sentimental Rules: On the Natural Foundations of Morality* attacks the soundness of the moral realists' arguments. He draws from so-called "epidemiological" accounts of cultural transmission to argue that his Sentimental Rules account (a variety of moral sentimentalism) provides a better explanation of moral evolution and the historical widening of

the moral community. His model has both explanatory and predictive powers that the realist models lack.

1.2 The Coming Terrain

I will proceed in Chapter 2 by laying out the arguments of the moral realists. It is necessary to divide such arguments into two camps. The first group includes Nicholas Sturgeon and related arguments by Michael Slote, Michael Smith, Peter Singer, and Thomas Nagel, whose arguments do not come attached with as broad of an explanatory apparatus as Peter Railton's argument from "Moral Realism," which sits alone in the second. It is necessary to offset his argument from the rest because he offers a much more in-depth and substantial analysis of moral evolution than do many other moral realists and substantially spells out one of the gears driving moral progress.

Chapter 3 begins with the relevant research, collected by Nichols, on the fitness for cultural transmission of our moral norms. His account has much in common with other "epidemiological" accounts from anthropology and focuses on the increased likelihood of survival by those norms prohibiting an action that elicits negative affect. After the relevant research has been elucidated, Nichols' attacks on the moral realists will be explained and evaluated. Next, Nichols' own argument to the best explanation will be evaluated. He maintains that his sentimental rules account of morality is in a better position to explain the phenomena regarding moral evolution than are those of the moral realists.

I agree with Nichols in that his account possesses greater explanatory power, but I maintain that he is using the inference to the best explanation in a different way than the moral realists are. In Chapter 4, I argue that there are two uses of the inference to the best explanation

of moral evolution at play in the surveyed literature. The first use, the one employed by the first group of moral realists, is circular, and cannot be successfully launched at an opponent, because it presupposes a metaethical thesis regarding the direction of fit between moral facts and moral beliefs. Since the relevant theorists are helping themselves to this metaethical thesis and then using their arguments to support it, their reasoning is circular. Railton and Nichols use the argument style to a different effect, since they are providing purely empirical explanations of empirical phenomena. While their arguments are not circular, they are left in a difficult position. It seems as though the only option for both theorists is to use a naturalistic reduction to explain moral terms. Neither Railton nor Nichols employs moral terms in their explanation, so the options left open for them to built moral terms into the content of our moral beliefs is a difficult task. It is of extreme interest to note that Nichols does not recognize the different ways arguments to the best explanation have been used in the literature, and so he supposes that he has established himself victorious over the realists. While his argument is not circular (and therefore not susceptible to one of the failings of the realists' arguments) he is left in a difficult position to explain moral terms.

Before all this progresses, however, it may be informative to start with an argument to the best explanation that can be seen as a model for the arguments that the moral realists want to make. In his 1971 paper "The Rationality of Aesthetic Value Judgments" Michael Slote constructs an argument to the best explanation regarding aesthetic judgments that he uses to argue for what he calls the "Thesis of Rationality in Aesthetics," which is the belief that claims of artistic merit "can be objectively confirmed, verified, or supported in the way that claims in science and mathematics are" (821). The Thesis of Rationality in Aesthetics, simply put,

explains that statements of aesthetic value are objective. To support this claim, Slote offers an argument attempting to explain cases of unidirectionality of aesthetic preference.

1.3 Slote's Argument Regarding Unidirectionality

Slote begins by observing that most well-versed listeners of classical music prefer Mozart to Bruckner and that such unidirectionality of preference is also apparent between other pairs of composers and figures working in other artistic genres (Rationality 823). As Nichols acknowledges (152), Slote is quick to point out the nature of his argument:

When scientists attempt to explain a given fact or phenomenon, they generally consider various possible alternative explanations of the fact or phenomenon, and accept one of these explanations as correct only if it is clearly more reasonable, in terms of certain standards of scientific methodology and according to the available evidence, than any of the alternative explanations that they have been able to think up. (Slote, Rationality 824)

Drawing an analogy with explanatory power in the sciences, he argues that his explanation of the unidirectionality of aesthetic preferences should be accepted only if it is the best explanation.

Having drawn the analogy with the sciences and having clearly situated the nature of his argument as inference to the best explanation, Slote continues, offering his explanation of the phenomenon:

But, assuming the existence of such unidirectionality, one explanation of it is of particular importance for our purposes here: the hypothesis, namely, that the more people study and are exposed to music, the more they like what is good in the field of music and the less they like what is mediocre or bad in the field of music, and that Mozart is, in fact, a greater, a finer, a better composer than Bruckner. (Slote, Rationality 824)

He explains the apparent unidirectionality of aesthetic preference for Mozart over Bruckner by claiming that Mozart actually is a superior composer to Bruckner and that over time increased study of music will cause someone to like what is better in music (i.e., Mozart) and prefer it to

what is not as good in music (i.e., Bruckner). He then asserts the importance of this account being the best explanation of the phenomenon of unidirectionality of preference if his argument is to succeed:

If we could show that this explanation of the unidirectionality that exists with respect to (the music) of Mozart and Bruckner is superior to all its various alternatives and also show that it is unreasonable to deny the existence of that unidirectionality, we could perhaps show that there is good reason for us to think that certain particular aesthetic value judgments are true, and this would give us good reason to believe in the truth of the [Thesis of Rationality in Aesthetics]. (Slote, Rationality 824)

His argument will go through if he can establish two premises: first, that the unidirectionality he seeks to explain actually exists; second, that his is in fact its best explanation. In the rest of his article he attempts to show that his argument is sound and that both these premises go through. Slote analyzes several alternative plausible explanations for the trend in preferences that he identifies but finds each of them unacceptable and/or less explanatory.

For present purposes, much of the rest of Slote's argument is unimportant. What matters is that Slote has used an inference to the best explanation to get to a type of realism regarding aesthetic judgments, his naturalistic dispositional theory of aesthetic value terms (Rationality 837-8). He is also careful to note that his argument could be used to defend value terms in other families besides aesthetics. The goal of the next chapter is to analyze arguments that do just that—extend the argument style Slote has applied and use it to defend a type of realism in ethics. Before we can embark on this endeavor, a few qualifiers are in order.

1.4 A Note on the Varieties of Realism

The philosophers I collect under the heading “Moral Realists” constitute a rather heterogeneous group: the ethical theories they espouse are diverse, and each has its own

particularities. Therefore, I wish to clarify what I mean by “moral realism” and thus show that my collection of these figures into one group is a legitimate move on my part. It will also be of use in later discussions to differentiate Nichols’ view (and other sentimentalist accounts generally) from moral realism in a substantial, illuminating manner. Nichols also sees the necessity in such a clarification and offers one thus:

The label “moral realism” is used for markedly different positions. In some cases, “moral realism” maintains only that some moral claims are true...On this construal, moral realism is perfectly consistent with a thoroughgoing relativism, according to which moral claims are true but relativized to an individual or a culture...In other places, “moral realism” is used to pick out the view that moral claims are not just true, but that they are true apart from any particular perspective and independent of people’s beliefs about right and wrong...The appeal to moral progress is typically used in defense of this stronger form of moral realism. (Nichols 150-1)

As Nichols demarcates the playing field of realism, the arguments he picks out are those that hold not only that moral claims carry a truth-value but their truth-value is not dependent upon a particular individual or culture.

Something stronger may be said of the theories that the various moral realists examined here espouse. I want to present and explain three conditions, and for my purposes any and all theories that count as ‘realist’ ones must meet them. Theories counting as realists hold:

- 1) Moral claims are truth-apt.
- 2) Some moral claims are true.
- 3) Moral claims have a truth-value in a substantial sense of truth.
- 4) The truth-value of moral claims is independent of our actual moral beliefs and practices of endorsing and rejecting on either the individual or cultural level.¹

¹ Conditions 1, 2, and 4 are relatively standard in the literature. I introduce condition 3 as a further differentiation between sentimentalism and realism, even though it is closely connected (as I will explain it) to condition 4. However, condition 3 is a separate substantial conclusion, because it fleshes out the notion of the truth predicate to which each theorist will have to commit.

These four claims say essentially the same thing that Nichols spells out, but they merit closer inspection than he affords them. The first condition is simple enough; it is a basic statement of cognitivism. Moral claims are capable of being true or false. The second condition is simply a denial of error-theory, and grants that moral claims are not only truth-apt, but that some of them will be true.

The third condition requires a bit more in terms of explanation. It appeals to the sense of truth that the theorist must invoke in his or her discussion of the bivalence of moral claims. I borrow the distinction of minimal versus substantial truth-aptness from A. E. Denham's *Metaphor and Moral Experience*, where she discusses the truth-aptness of moral discourse (100-5). She begins with Paul Horwich's disquotational schema ('p' is true if and only if p)² and Crispin Wright's argument that the disquotational schema fails to respect the intuitive distinction between truth and warranted assertability. From this failure emerges the idea of a minimal conception of truth:

The minimalist view is that when a predicate has been shown to have the relevant features and to have them for the right reason, there is then no further question about the propriety of regarding it as a truth predicate. Minimalism is thus at least in principle open to the possibility of a *pluralist* view of truth: there may be a variety of notions, operative within distinct discourses, which pass the test. (Wright 24-5)

Under a minimalist conception of truth, different truth predicates could apply to different discourses, and any candidate is plausible so long as it meets certain minimal constraints, including "[accommodating] the 'logic' of assertion, and...[coinciding] in normative force with (but...potentially divergent in its extension from) the norms of warranted assertability operating over a discourse" (Denham, 103). So long as the standards of correctness of a discourse satisfy the minimal constraints, then the standards of correctness of that discourse may count as

² See Horwich's *Truth* for a full account of his disquotational schema and the deflationist theory of truth.

standards of truth within that discourse, and claims made in that discourse will be truth-apt (Denham 103). According to Denham, moral discourse satisfies these requirements for truth-aptness (103).

Denham, though she thinks that moral discourse is at least minimally truth-apt, has some substantial further worries:

Suppose we agree, at least in principle, that one can define a truth-predicate on moral discourse. Does this automatically earn for it the status of 'cognitive'? Or does this status require more than a (perhaps wholly minimal) truth-predicate? ...In particular, it may be felt that the way in which a discourse is secured of minimal truth is what matters to its status as cognitive or non-cognitive. For a moment's reflection will reveal that the conceptual discipline required for minimal truth is secured by convergence in subjects' judgements, and this convergence might be explained by either a convergence of our traditional cognitive responses (say, our perceptual and rational responses) or a convergence of our feelings or sentiments. Suppose, for example, that by some miracle we all came to a meeting of minds with respect to our fashion preferences...That would render fashion discourse minimally truth-apt. But would it make it cognitive? Would it make one's fashion commitments a proper subject of epistemic evaluation? It is at least not obvious that it would. (Denham 104)

The truth-aptness of a discourse is not enough to secure the claim that the discourse is cognitive.

To ensure that moral discourse demarcates a cognitive domain, one needs to have a truth-predicate of the proper type, and mere convergence of judgments, as Denham hopes to have shown, is not substantial enough to make the discourse cognitive (104-5). She explains:

...[A]t least some truth-predicates may require more than convergence: a more 'substantial' truth-predicate may require that the convergence of belief be explained in one way rather than another. Specifically, we may suppose that a more substantial truth-predicate may only be defined on a discourse in which convergence is explained by subjects' ability accurately to represent states of the world, rather than the propensity to respond to its states in some other ways (e.g. in affective or conative ways). (Denham 105)

In denying that mere convergence of affect allows moral discourse to be cognitive, Denham points out that our conception of moral 'truth-makers' (i.e., the truth conditions for a claim in moral discourse) must be robust enough to make the discourse cognitive. All the moral realists

examined here provide substantial conditions for the truth of moral claims, whereas in the case of Nichols (a sentimentalist), it is unclear whether he can delineate truth conditions strong enough for moral discourse to count as cognitive in the requisite sense. Thus, another condition that can help us delineate moral realism from other theories is whether the truth conditions for moral claims are merely those minimally required for truth-aptness or if they are substantial enough to ensure more than that.

The fourth condition that I have presented is an extension of the third, as it introduces additional qualifiers on the truth conditions that a realist theory can proclaim. This condition further distances the moral realists from Nichols and other sentimentalists by removing the truth-value of moral facts from a dependency on our actual patterns of endorsement and beliefs. The moral realist can claim that even though we in fact believe an action to be morally permissible, we can actually be wrong about this claim. Under certain varieties of sentimentalism, by contrast, it is unclear that we can call certain beliefs mistaken, since moral facts are dependent upon moral beliefs. This condition eliminates various forms of relativism from counting as ‘realists’ as that label is being used here. The theorists I discuss—Sturgeon, Slote, Smith, Singer, Nagel, and Railton—all fit this broad construal.

1.5 Progress or Evolution?

Before I begin, I want to offer one further note clarifying the debate. It should be noted that moral realists tend to describe the expansion of the moral community and other such phenomena as being “moral progress.” This is a thick concept for them, since on their views, moral beliefs are not just changing, but are changing and progressing towards a closer correlation with the moral facts: moral beliefs change and over time more closely approximate the moral

facts. Sentimentalist theories (among whom I count Nichols) tend to describe the same phenomena as “moral evolution.” They are not committed to a theory-laden end-point or *telos* of moral change. Avoiding the thick concept of “moral progress” they describe the phenomena as “moral evolution,” and the changes in norms do not necessarily say anything about the end point of morality or necessarily lead to an end point of moral change. It is not until the final chapter that the different senses of these phases becomes important: there I will argue that the two distinct terms in fact pick out substantially different conceptions of the phenomena. We are now in a much better position to begin our analysis of inferences to the best explanation.

Chapter 2

Moral Realism and the Argument from Moral Evolution

2.1 Overview

This chapter will examine various arguments to the best explanation provided by a number of different philosophers who fit into the loose category of moral realism employed here. Special attention will be paid to two figures in this chapter. The first is Nicholas Sturgeon. Sturgeon's argument is one of the most developed and often referenced in the literature, and many of the other arguments I will examine do not substantially divert from it. The second argument worthy of greater attention is the account provided by Peter Railton. Railton's argument differs from the others in explanatory power, and Nichols will highlight this difference in his attack on realist inferences to the best explanation. Nichols' concerns justify a closer examination of Railton's hypothesis than those of some of the other realists surveyed here.

2.2 Sturgeon *Et Al*

Sturgeon's argument from moral evolution, which has been picked out as particularly interesting among the lot,³ can be found in his article "Moral Explanations" in which he attempts to defend the view that moral claims can have explanatory power against an argument from Gilbert Harman. Harman's argument comes from *The Nature of Morality: An Introduction to Ethics*. Here, Harman asks if "moral principles can be tested and confirmed in the way scientific principles can?" (3) Both ethical investigation and scientific investigation take the form of thought experiments, but unlike thought experiments about ethics, scientific thought experiments can be tested against the way the world is (Harman 4). Harman asks, "You can observe someone

³ See Michele Moody-Adams, "The Idea of Moral Progress," page 83. Moody-Adams herself endorses the idea of moral progress bolstering the case for moral realism. Her paper deals insightfully with some of the ways in which moral progress can occur.

do something, but can you ever perceive the rightness or wrongness of what he does?” (4) Is the rightness or wrongness of the action judged on a perception of its rightness or wrongness? He asks a further question:

If you round a corner and see a group of young hoodlums pour gasoline on a cat and ignite it, you do not need to *conclude* that what they are doing is wrong; you do not need to figure anything out; you can *see* that it is wrong. But is your reaction due to the actual wrongness of what you see or is it simply a reflection of your moral “sense,” a “sense” that you have acquired perhaps as a result of your moral upbringing? (Harman 4)

What these questions drive home is what Harman calls “a basic philosophical problem about morality” (vii). This problem is morality’s “apparent immunity from observational testing” (Harman vii) and “the seeming irrelevance of observational evidence” for morality (Harman viii).

Sturgeon summarizes Harman’s argument as saying “that reference to moral facts appears unnecessary for the *explanation* of our moral observations and beliefs” (50). Harman’s goal is to push anyone who wants to avoid ethical skepticism into giving reductive definitions for our ethical terminology and make moral observations and beliefs dependent upon non-moral facts (Sturgeon 50). Sturgeon latches onto cases of moral character to show that moral facts can have explanatory power:

My Hitler example suggests a whole range of extremely common cases that appear not to have occurred to Harman, cases in which we cite someone’s moral character as part of an explanation of his or her deeds, and in which that whole story is then available as a plausible further explanation of someone’s arriving at a correct assessment of that moral character. (Sturgeon 63)

Hitler’s moral character can be used in an explanatory way to describe his actions: Hitler was a morally depraved individual, and since he was a morally depraved individual, he was capable of ordering the execution of millions.

Sturgeon cites another case of Selim Woodworth who volunteered to lead a rescue party and failed miserably, his incompetence resulting in deaths that could have been prevented (63).

A researcher attributes Woodworth's actions to his moral character, which he claims is despicable. Sturgeon comments:

I cite this case partly because it has so clearly the structure of an inference to a reasonable explanation. One can think of competing explanations, but the evidence points against them. It isn't, for example, that Woodworth was a basically decent person who simply proved too weak when thrust into a situation that placed heroic demands on him. He volunteered, he put no serious effort even into tasks that required no heroism, and it seems clear that concern for his own position and reputation played a much larger role in his motivation than did any concern for the people he was expected to save. (Sturgeon 64)

In the absence of any additional evidence, it seems that Woodworth's behavior can be successfully explained by simply saying that he was not a morally upstanding individual. His moral character—his moral depravity—caused him to be unconcerned with the lives that were dependent upon his success and more concerned with his own comfort and social standing.

Sturgeon attempts to expand his argument by posing a question:

But do moral features of the action or institution being judged ever play an explanatory role? Here is an example in which they appear to. An interesting historical question is why vigorous and reasonably widespread moral opposition to slavery arose for the first time in the eighteenth and nineteenth centuries, even though slavery was a very old institution; and why this opposition arose primarily in Britain, France, and in French- and English-speaking North America, even though slavery existed throughout the New World. (Sturgeon 64)

Sturgeon asks why the moral standing of the institution of slavery, which had existed for millennia throughout the world, was first questioned only a few hundred years ago.

Furthermore, given that slavery was a widely practiced phenomenon, why did opposition arise in only a few countries and their colonies? He asserts the following:

There is a standard answer to this question. It is that chattel slavery in British and French America, and then in the United States, was much worse than previous forms of slavery, and much worse than slavery in Latin America. (Sturgeon 64)

In this answer, the change in attitude is attributed to a moral fact. Thus, the moral belief in the legitimacy of slavery was challenged as a result of the moral fact that slavery as an institution in the relevant areas was worse in its practices than slavery in other areas and other time periods.

The moral facts explain the change in moral beliefs.

Further challenging Harman's criticism of the explanatory power of moral facts, Sturgeon makes an explicit appeal to inference to the best explanation:

For it is natural to think that if a particular assumption is completely irrelevant to the explanation of a certain fact, then the fact would have obtained, and we could have explained it just as well, even if the assumption had been false. But I do not believe that Hitler would have done all he did if had not been morally depraved, nor on the assumption that he was not morally depraved, can I think of any plausible alternative explanation for his doing those things. Nor is it plausible that we would all have believed he was morally depraved even if he hadn't been. (Sturgeon 65)

Sturgeon makes the case that the best explanation for what Hitler did is that Hitler was a morally depraved person. Likewise, the best explanation for the historical shift in moral feeling towards slavery changed as a result of the moral facts surrounding the institution of slavery as it was practiced.

Sturgeon's argument thus takes the form of inference to the best explanation. Moral beliefs regarding slavery changed drastically a few centuries ago. One possible explanation for this change is that practice of slavery was morally much worse than it had been in previous centuries. This explanation is the best among other possible, less plausible interpretations. Therefore, the moral facts of the matter caused the change in moral opinions of slavery. One can see this argument as an argument from moral evolution to moral realism (in the limited sense) because it takes the process of moral evolution (changes in norms regarding slavery) and infers from the most plausible explanation of this evolution that moral facts exist.

A look at some of the other instances of realist arguments from moral evolution will help clarify the issue. One such argument comes from a later article by Michael Slote, “Is Virtue Possible?” in which he identifies some of the inadequacies of the virtue ethics of Murdoch and McDowell. Slote conducts what can be considered a very brief, very skeletal anthropology of the history of moral opinions surrounding slavery and then asks:

If virtue requires the disposition to act justly, do the right thing, with regard to important issues, can we sensibly suppose that the failure of the ancient world to attain and act upon a correct moral view of slavery merely reflects personal weaknesses or deficient methods of moral training that, unluckily, happened to be fairly universal in those days? Presumably not. (Slote, Virtue 72)

Slote is criticizing the espoused views of both Murdoch and McDowell that virtue is dependent upon the individual’s character, and certain character traits are the primary obstacles to virtue for an individual. Are personal character deficiencies really the best explanation for the failure of the ancient world to see the moral depravity of slavery? Slote wants to argue that there are bigger, more impersonal forces that work against the individual’s becoming virtuous:

[The ancient belief that slavery is natural and inevitable] is at least partly to be explained by the fact that slavery was a universal phenomenon. Just as ignorance of the alternative terms used by other languages can make matters of linguistic convention seem to be inevitable facts of nature, so too can ignorance of alternatives to a given social arrangement instill the belief that that arrangement is natural and inevitable and thus beyond the possibility of radical moral criticism. So if the ancients were unable to see what virtue required in regard to slavery, that was not due to personal limitations (alone) but requires some explanation by social and historical forces, by cultural limitations, if you will. And if we today can see the wrongness of slavery, that is in part because we have the benefit of knowledge that makes slavery seem less natural and inevitable. For unlike the ancients we know of “experiments” in living without slavery; we possess an historical record of societies, among them our own, where slavery has been absent and people have survived and flourished, nonetheless. (Slote, Virtue 72-3)

Slote argues that the failing of the ancient world to see the wrongness of slavery maybe the result, at least in part, of the lack of plausible alternatives that the ancients believed open to them. The ancient world experienced slavery as a universal practice, one from which no society that

they were aware of seemed to abstain. The abolition of slavery just was not something that occurred to them and if it had, it would not have been plausible.

Slote's argument thus far involves seeing that certain historical contingencies may have made virtue impossible for past generations. This condition holds with the case of slavery. Slote then turns his eye to our times. Are there historical factors or influences that might make present day peoples blind to some moral truth? Slote asks this question:

Is there any less reason to believe the same about current-day moral thinking? If not, then there is reason to believe that moral virtue as traditionally understood is not accessible even today, but rather, like various kinds of scientific truth, is to be attained, if at all, only at the historical limit of human cultural endeavour, in a long run that no individual may ever encompass. (Slote, *Virtue* 75)

True virtue may be off-limits even to us. In fact, there may be some moral truths only knowable at the theoretical limit of human culture—at the absolute end of our advancement. Here, Slote also picks up on the analogy between moral philosophy and science, whereby moral philosophy gets closer to the truth of the matter over time just as science seems to do the same. After expressing some concerns regarding such analogies, he concludes:

And if new and promising ideas not about moral concepts but about actual right and wrong can thus emerge through the individual insight or vision of particular moral philosophers (or others), does this not locate another reason why virtue may not be attainable in those too-often-neglected factors of historical development that make some sort of analogy between science and moral thinking entirely plausible? (Slote, *Virtue* 76)

Developments in moral philosophy may help us get closer to the facts regarding actual right and wrong, just as developments in science help get us closer to the truth of the matter in the examination of the universe. Slote, in a very subtle way, is offering an argument to the best explanation of moral progress: moral philosophy progresses and moral concepts change because we are approaching the fact of the matter about right and wrong as we progress. This argument sits nicely with his argument from "The Rationality of Aesthetic Value Judgments."

Michael Smith in *The Moral Problem* offers an argument similar to the ones under discussion here (187-9). Smith attempts to reach objectivity in his dispositional theory of value by appealing to the idea of a rational convergence of moral opinions, so his goals for the argument are somewhat different from those of Sturgeon and Slote. After offering a synopsis of his view regarding the *a priori* of the truths of morality, he begins his argument:

However, for all that, it seems to me that we should none the less have some confidence in the legitimacy of moral talk. For, in short, the empirical fact that moral argument tends to elicit the agreement of our fellows gives us reason to believe that there will be a convergence in our desires under conditions of full rationality. For the best explanation of that tendency is our convergence upon a set of extremely unobvious *a priori* moral truths. And the truth of these unobvious *a priori* moral truths requires, in turn, a convergence in the desires that fully rational creatures would have. (Smith 187)

Rational argument tends to lead to agreement regarding moral judgments. From this fact, Smith infers the possibility of a convergence of desires among fully rational agents. This convergence will be to a set of moral truths that are known *a priori*. Furthermore, that these moral truths are true necessitates the convergence of desires among fully rational agents.

Smith notes the real existence of significant moral disagreement, but he bolsters his argument with three points. First, he indicates that there is extensive moral agreement on many issues (Smith 188). This idea is easy enough to establish. For example, few of us find the need to argue the case against wanton killing and thieving (in the absence of extenuating circumstances). His second point is more significant for my purposes here. Smith points out:

[W]e must remember that in the past similarly entrenched disagreements were removed *inter alia* via a process of moral argument. I am thinking in particular of the historical, and in some places still current, debates over slavery, worker's rights, women's rights, democracy and the like. We must not forget that there has been considerable moral progress, and that what moral progress consists in is the removal of entrenched disagreements of just the kind we currently face. (Smith 188)

Smith is claiming that moral progress is just the resolution of entrenched moral disagreement. Smith is arguing that the best explanation of moral progress is that, through the course of history, people have resolved moral conflict by reaching a rational convergence upon *a priori* moral truths arrived at by productive argument. His third point is that much moral disagreement can be attributed to factors inhibiting or standing in for rational exercise, such as religion (Smith 188-9).

Smith is arguing, when he makes his second point above, that moral progress occurs when moral disagreements are resolved. This resolution is the effect of rational argumentation, which could get us closer to the fully rational truths that occur at the theoretical point of full convergence among fully rational agents. His argument form is again inference to the best explanation. Resolution of moral conflict is taken as evidence that morality would be arrived at under conditions of full rationality.

Arguments from moral evolution can also be used to justify very different forms of moral realism as the following argument from Peter Singer shows. In *The Expanding Circle: Ethics and Sociobiology* Peter Singer examines how moral communities seem to expand over time. He asks:

So the shift from a point of view that is disinterested between individuals within a group, but not between groups, to a point of view that is fully universal, is a tremendous change—so tremendous, in fact, that it is only just beginning to be accepted on the level of ethical reasoning and is still a long way from acceptance on the level of practice. Nevertheless, it is the direction in which moral thought has been going since ancient times. Is it an accident of history that this should be so, or is it the direction in which our capacity to reason leads us? (Singer 113)

We are still reckoning with the aftermath of expanding the circle of ethics to include all human beings, according to Singer. Moral philosophy, he claims, has been pushing us in this direction since its beginnings, and it is completely legitimate to ask how far the expansion of the moral community will go:

Where does this process end? Taking the impartial element in ethical reasoning to its logical conclusion means, first, accepting that we ought to have equal concern for all human beings...The circle of altruism has broadened from the family and the tribe to the nation and race, and we are beginning to recognize that our obligations extend to all human beings. The process should not stop there...The only justifiable stopping place for the expansion of altruism is the point at which all whose welfare can be affected by our actions are included within the circle of altruism. This means that all beings with the capacity to feel pleasure or pain should be included; we can improve their welfare by increasing their pleasures and diminishing their pains. The expansion of the moral circle should therefore be pushed until it includes most animals. (Singer 119-20)

Singer is a utilitarian, and he famously thinks that the moral community should be expanded to include all sentient beings capable of feeling pleasure and pain. He offers an explanation of moral progress—here, the historical expansion of the moral community to include all human beings—to justify his utilitarian ethics. Implicitly, he is claiming that since all human beings feel pleasure and pain, they must be considered of the same moral standing as ourselves, our family members, our fellow citizens, etc. Reason leads us to recognize this, and the moral community expands as a result. He also predicts, based upon his utilitarian affiliation, that the moral community will continue to expand until all animals capable of feeling pleasure and pain are given moral standing.

Thomas Nagel has a brief discussion of moral progress in his *The View from Nowhere* and, though his argument form is slightly different from that of some of the other authors considered here, it is relevant to what we are presently considering. He says:

Although the methods of ethical reasoning are rather primitive, the degree to which agreement can be achieved and social prejudices transcended in the face of strong pressures suggests that something real is being investigated, and that part of the explanation of the appearances, both at simple and at complex levels, is that we perceive, often inaccurately, that there are certain reasons for action, and go on to infer, often erroneously, the general form of the principles that best account for those reasons. (Nagel 148)

Despite the primitive nature of our devices for investigating ethics, that we can achieve some measure of accord even when confronted by emotionally charged barriers to agreement implies, according to Nagel, that there is something real that we are striving to reach. There are reasons to act in certain ways, and we infer from those reasons general principles that accommodate them, sometimes mistakenly. One explanation of moral agreement is the approach to moral principles. Nagel clarifies:

Again let me stress that this is not to be understood on the model of perception of features of the external world. The subject matter of our investigations is how to live, and the process of ethical thought is one of motivational discovery. The fact that people can to some extent reach agreement on answers which they regard as objective suggests that when they step outside of their particular individual perspectives, they call into operation a common evaluative faculty whose correct functioning provides the answers, even though it can also malfunction and be distorted by other influences. It is not a question of bringing the mind into correspondence with an external reality which acts causally on it, but of reordering the mind itself in accordance with the demands of its own external view of itself. (Nagel 148)

Coming to these moral principles is a matter of abandoning one's particular perspective and seeing oneself from the outside. Nagel is a rationalist, and moral principles are rational constraints on motivation. Our moral judgments aspire to objectivity, and this can be explained by reference to the mind's ability to see itself from the outside in. We infer principles from motivating reasons and use these principles to explain the reasons. Moral beliefs are, to some extent, explained by reference to a further understanding of our own minds and reasons for acting.

Further on in *The View from Nowhere*, Nagel makes an analogy that should be familiar to us by now:

I do not think it is utopian to look forward to the gradual development of a greater universality of moral respect, an internalization of moral objectivity analogous to the gradual internalization of scientific progress that seems to be a feature of modern culture. (Nagel 187)

Here, Nagel, like many theorists before him, draws an analogy between moral philosophy and science. Nagel compares the internalization of scientific study with what he hopes will occur with moral philosophy. So Nagel can be seen as espousing a view similar to that of the other moral realists described here: as time passes, we will (hopefully—and Nagel attaches many qualifiers to this hope) understand more of the truths of morality; we will see through our historical and personal contingencies and more moral truths will be apparent to us.

Each of the moral philosophers examined here use inferences to the best explanation of moral progress to support their respective ethical theories. The theories represented by this group are diverse and range from utilitarianism to dispositional theories of value. What they have in common is the argument structure: they attempt to explain whatever form of moral progress they pick out by appeal to a realist metaethics.

2.3 Railton's Social Unrest Analysis

Peter Railton in “Moral Realism” also posits an argument to the best explanation of the type that we have been considering. We must explain his argument separately, however, because it comes attached to a wider theoretical apparatus that will prove useful in later discussion. Railton's realism is based around the thesis “that moral norms reflect a certain kind of rationality, rationality not from the point of view of any particular individual, but from what might be called a social point of view” (Railton 190). Social rationality considers “the interests of all potentially affected individuals counted equally under circumstances of full and vivid information” (Railton 190). Thus, what is rational in Railton's sense here is that which takes into account the needs and wants of all affected and considers them without prejudice; it is idealized, because it takes them into account under conditions of maximal information. To inform this

thesis, Railton contends that a society that discounts the interests of some of its members will be subject to unrest and resentment from the disaffected members because the society has violated social rationality (191). He elaborates:

[T]he discontent produced by departures from social rationality may produce feedback that, at a social level, promotes the development of norms that better approximate social rationality. The potential for unrest that exists when the interests of a group are discounted is potential for pressure from that group—and its allies—to accord fuller recognition to their interests in social decision-making and in the socially-instilled norms that govern individual decision-making. It therefore is pressure to push the resolution of conflicts further in the direction required by social rationality, since it is pressure to give fuller weight to the interests of more of those affected. (Railton 193)

Social unrest can give way amongst disaffected groups to social pressure to change the moral norms that conflict with their interests. Since these interests are objectified,⁴ it is against social rationality to dismiss them. Pressure to undo or refine the norms in question pushes society further towards the prescriptions of social rationality. Therefore, any social group capable of applying effective pressure can elicit change. As a result, one would expect society to progress “toward the inclusion of the interests of (or interests represented by) social groups that are capable of some degree of mobilization” (Railton 194-5).

Railton argues that his theory will have explanatory and predictive value suggesting that “over time, and in some circumstances more than others, we should expect pressure to be exerted on behalf of practices that more adequately satisfy a criterion of rationality” (196-7). He then attempts to explain three historically observable trends in the evolution of moral norms. The first of these, and the most important for present purposes, is generality; moral communities have expanded over time to incorporate larger and larger social groups, starting with the family and progressing to the entire species (Railton 197). This trend occurs because social communities

⁴ See the earlier parts of Railton’s article for his notion of an objectified interest. To fully explicate the notion here is not entirely necessary, since it is not critical to the attack that Nichols launches on Railton’s view.

have expanded and incorporated more and more people, from family to tribe to state (Railton 197). Railton explains the mechanism behind this trend as a result of the inclusion of moral groups into the society:

Expanding social entities frequently subjugate those incorporated within their new boundaries, and the means by which those thus oppressed have secured greater recognition of their interests have been highly conflictual, and remain—perhaps, will always remain—incomplete. (Railton 197)

As a new group is assimilated into a society, the society subjugates and oppresses them. This oppression breeds resentment because the interests of the new group are not being recognized. The lack of recognition invites the potential for social unrest, and reforms are made in the system.

The second trend that Railton picks out is the “humanization” of morality. Many origins and sources of moral norms reside in institutions such as religion and codes of honor that lack reference to actual human concerns and needs (Railton 197-8). Railton explains:

While vestiges of these views survive in contemporary moral theory, it is typical of almost the entire range of such theory, and of much of contemporary moral discourse, to make some sort of intrinsic connection between normative principles and effects on human interests. Indeed, the very emergence of morality as a distinctive subject matter apart from religion is an instance of this pattern. (Railton 198)

While Railton makes no explicit argument showing how his social unrest analysis of morality explains the trend to humanization, it is not difficult, from the sketch above, to piece together the reasoning behind his appeal. The idea of humanization that he appeals to is one whereby our moral codes take into account our actual interests. The social rationality model Railton offers claims that an insufficient inclusion of human interest in the moral theory will result in resentment and a pressure for the code to change. In order to be responsive to our interests and have a binding affect on us, moral principles need to make some reference to our interests, or

else social pressure will be exerted upon them and they may be overthrown. The shift away from religious morality can be explained in this way, because religious morality fails to make reference to our actual interests, thus violating the sense of social rationality he has outlined. Therefore, pressure has been exerted upon it to change and moral theories have undergone a process of humanization.

A third explanatory advantage that Railton attributes to his theory is the predictive power it demonstrates regarding certain types of moral norms. Moral norms most closely attending to interests that everyone possesses and protecting those interests should demonstrate an increased stability and early onset in a society (Railton 198). Oppositely, where moral debate surrounds an issue that is less likely to be easily solved in a manner consistent with everyone's interests, the respective norms will be less stable and agreed upon much later (Railton 198-9). Railton explains:

Given a suitable characterization of the conditions that prevailed during the processes of normative evolution described by these patterns, the present theory claims not only that these changes could have been expected, but that an essential part of the explanation of their occurrence is a mechanism whereby individuals whose interests are denied are led to form common values and make common cause along lines of shared interests, thereby placing pressure on social practices to approximate more closely to social rationality. (Railton 199)

He claims that his theory not only predicts these trends in the history of moral evolution—it also explains why these trends occur. Such an explanatory and predictive power in this case and others is something that Nichols will single out as a virtue of Railton's argument.

Given Railton's argument, it could be difficult to see how exactly he fits into the definition of realism in use here. To some extent, for Railton, morality is dependent upon our interests. The following passage helps to illuminate exactly why Railton can be categorized with the other moral realists being studied here:

Just as explanations involving assessments of individual rationality were not always replaceable by explanations involving individual *beliefs about* what would be rational, so, too, explanations involving assessments of social rationality cannot be replaced by explanations involving *beliefs about* what would be morally right. For example, discontent may arise because a society departs from social rationality, but not as a result of a belief that this is the case. (Railton 192)

Our actual beliefs about what is morally right and wrong do not determine what counts as socially rational. Railton's theory, therefore, fulfills my fourth condition for a theory to be considered realist. But what about the first three? At the onset of his discussion, Railton outlines certain theses that his theory will support:

I will argue for a form of moral realism which holds that moral judgments can bear truth values in a fundamentally non-epistemic sense of truth; that moral properties are objective, though relational; that moral properties supervene upon natural properties, and may be reducible to them; that moral inquiry is of a piece with empirical inquiry; that it cannot be known *a priori* whether bivalence holds for moral judgments or how determinately such judgments can be assessed; that there is reason to think that current moralities are wrong in certain ways and could be wrong in quite general ways...and that, while there are perfectly general criteria of moral assessment, nonetheless, by the nature of these criteria no one kind of life is likely to be appropriate for all individuals and no one set of norms appropriate for all societies and all times. The position thus described might well be called 'stark, raving moral realism', but for the sake of syntax, I will colorlessly call it 'moral realism'. (Railton 165)

How much of a "stark, raving moral realist" is Railton? From this passage, we can see that Railton thinks moral discourse is truth-apt—thus fulfilling the first condition. Railton does not argue for an error-theory, so his theory seemingly fulfills the second condition. While the possibility of pluralism he admits to may seem problematic, given that he separates the moral fact of the matter, as defined by what is socially rational, from actual beliefs about what is morally right, this does not threaten his theory's fulfillment of the fourth condition. Problems arise, however, with the third condition.

It is not completely clear that Railton is arguing for a substantial enough sense of truth-aptness to fulfill the requirements that Denham sets up for the discourse, as he perceives it, to be

cognitive. Recall the earlier cited statement from Denham that “a more substantial truth-predicate may only be defined on a discourse in which convergence is explained by subjects’ ability to accurately represent states of the world, rather than the propensity to respond to its states in some other ways” (105). Railton’s theory seems to appeal to a strong enough notion of the truth-aptness of moral judgments to make the discourse cognitive: social rationality is defined as the state in which the interests of everyone are taken into account *under the condition of being maximally informed*. Since the interests are objectified and maximally informed in the social point of view he argues for, his account is not reduced to providing a minimal sense of truth for moral discourse because it makes use of something beyond our actual beliefs and interests. Therefore, Railton’s view seems to fulfill the third constraint and can count as realist for present purposes.

2.4 Conclusion

The arguments that moral realists have espoused in an attempt to explain moral progress are as varied as the types of realism they defend. Though we may have managed to oversimplify the playing field by lumping so diverse a group of philosophers together, we are now prepared to evaluate Nichols’ attack on the realists’ inferences to the best explanation and his sentimentalist account of moral evolution.

Chapter 3

Nichols and the Affective Resonance Account

3.1 Introduction

In his *Sentimental Rules: On the Natural Foundations of Moral Judgment*, Shaun Nichols explicates his Sentimental Rules account of moral judgment, which incorporates a sentimentalist metaethical theory. In this chapter, I want first to offer a brief account of Nichols' Sentimental Rules theory. Once this is in place, I will examine the anthropological research that Nichols examines and his proposed extensions to it. With his extension to "epidemiological" accounts of cultural transmission, Nichols hopes to offer a convincing analysis of the genealogy of our harm norms that bolsters his sentimentalist metaethics. Given this explanation, I will examine Nichols criticisms of the inferences to the best explanation of the moral realists.

3.2 Nichols and the Sentimental Rules Account

Nichols begins his explanation of the Sentimental Rules account of moral judgment by examining the nature of the moral/conventional distinction as it has been explored in experimental psychology. The moral/conventional distinction picks out the testable conceptual differences between convention and morality. Conventions are behaviors that fit into an overall system and are given meaning by their place in such a system (Nichols 5). The classical example of a conventional system is our system of manners, and the individual manners norms within it are given meaning only as part of the wider system of manners. Rules distinguished as moral are those rules that are unconditionally binding, capable of generalization, and stemming from impersonal concerns (Nichols 5). Nichols indicates that the evidence on this distinction is the best way to begin an account of moral judgment, and he summarizes the empirical results:

[M]oral violations are treated as distinctive along several different dimensions. Moral violations attract high ratings on seriousness, they are regarded as having wide applicability, they have a status of authority independence, and they invite different kinds of justifications from conventional violations. Finally, this turns out to be a persistent feature of moral judgment. (Nichols 6)

This list summarizes the differences between moral violations and conventional violations as revealed through empirical research, most notably using children as subjects. The research upon which he draws focuses primarily on harm-based transgressions as moral violations. It indicates that even children, in judging harm-based transgressions to be wrong, base their wrongness on the harm inflicted upon the victim—the wrongness is not based on the act violating the rules of some authority (Nichols 6). Another interesting finding on the moral/conventional distinction is that children in very different settings and cultures will judge violations of harm-based norms to be wrong, hence their wide applicability (Nichols 6). Conventional violations, on the other hand, are seen as violating a rule of some authority figure and such violations would not be violations given that the authority figure is not present: conventional norms lack wide applicability (Nichols 6). After rejecting different analyses of the data, Nichols explicates his own account.

Central to his account is the idea of an internal ‘normative theory’. He is very vague on what constitutes a normative theory, but he maintains that it must be internally represented (the agent must be able to apply the norms in question), that it is a set of rules, and these rules need not be connected by any substantial ethical theory:

“Normative theory” is not intended in any inflated sense. Rather, even a motley set of rules prohibiting certain behaviors will count as a normative theory. Internally represented rules concerning table manners, for instance, will count as a normative theory. The normative theory of central interest to us, however, is the normative theory prohibiting harmful actions. (Nichols 16)

Again, his focus is on harm-based norms. Incorporating research into the behavior and abilities of psychopaths, Nichols maintains that the ability to make moral judgments requires an affective

mechanism to be present in the agent. This affective mechanism responds to instances of suffering in others. Core moral judgments, for Nichols, are judgments about actions that violate our harm-based norms and elicit a strong negative response in the agent (18). The interaction between the affective mechanism and the normative theory is crucial for Nichols, since core moral judgment can only take place at the intersection of the two. Nichols summarizes:

Core moral judgment depends on two mechanisms, then, a normative theory prohibiting harming others, and some affective mechanism that is activated by suffering in others. Core moral judgment thus implicates what I'll call *Sentimental Rules*, rules prohibiting actions that are independently likely to elicit strong negative affect. The set of rules or normative theory prohibits actions of a certain type, and actions of that type generate strong affective response. (Nichols 18)

So the set of norms that constitute our normative theories hold sway over us because they are linked to strong emotional reactions as a result of our affective mechanisms regarding suffering in others. We judge an action that produces harm in others as being wrong because it violates our norms that influence us by causing a negative emotional response when we view suffering in others.

Nichols characterizes the affective mechanism that gives rise to altruistic motivation—the motivation that backs our harm-based norms—very minimally. He wants to posit something that will explain why young children seemed moved by suffering in others, so it cannot involve complex mindreading abilities, since he has already provided empirical evidence that young children lack the ability to attribute some beliefs and motivations to others. He defines the affective mechanism backing harm-based norms thus:

Altruistic motivation depends on a mechanism that takes as input representations that attribute distress, for example, “John is experiencing painful shock,” and produces as output affect that, inter alia, motivates altruistic behavior. Following the terminology introduced above, I'll call this system the Concern Mechanism. (Nichols 56-7)

The Concern Mechanism is the mechanism that takes indications of pain or distress in others and produces some sort of affect that can motivate altruistic behavior. So core moral judgments are those whereby an action is viewed to be in violation of a harm-based norm that is backed by the Concern Mechanism that provides motivational force for our harm-based norms.

In order to understand the story that Nichols will tell about the origin/evolution of norms, it is necessary to examine the relationship he outlines between the normative theory and the affective mechanism. This point is one on which Nichols' account is rather sketchy. He does not commit to any specific relationship. After outlining two possible views, he concludes:

Of course, one needn't take an extremist view at all. It's possible, for instance, that the affective system plays both a developmental and an on-line role in generating and preserving the nonconventional normative judgments. The point I wish to emphasize is simply that even if nonconventional normative judgment does involve both a normative theory and an affective mechanism, it remains to be seen exactly how those different mechanisms interact to enable the distinctive kinds of judgments subjects make about disgusting violations and moral violations. (Nichols 29)

So Nichols, in a very non-committal fashion, claims that neither the view that norms arise from the affective mechanism nor the view that they work perfectly side-by-side need be correct. It could be a little of both.

3.3 Nichols' Metaethical Commitments

The question naturally arises: what does Nichols' Sentimental Rules account of moral judgment obligate him to metaethically? I have been referring to Nichols as a sentimentalist but have not offered a substantive definition of sentimentalism. A sentimentalist moral theory is one that bases moral judgments on our emotional or sentimental reactions. Like realism, it is a label that incorporates a broad range of theories. Sentimentalism can also be characterized in terms of the four conditions laid out in the first chapter. Such sentimentalist theories will generally fulfill

the first and second conditions, that moral claims are truth-apt and that some count as true. What differentiates them from realist theories is their stance on conditions 3 and 4. Because sentimentalists tend to base moral facts upon the moral beliefs that we actually have (i.e., our reactions or approbation or disapprobation), they do not fulfill the fourth condition, and it is also not clear that they are committed to a substantial enough conception of truth to make the moral discourse cognitive. Since they base the truth-aptness of moral claims on our actual sentiments, there is an analogy between their views and the view of fashion that Denham argues fails to make the discourse cognitive. Basing the truth of moral claims on our actual sentiments seems to merely make judgments in the discourse minimally truth-apt, but fails to make the discourse cognitive.

Is Nichols a sentimentalist? He does directly compare his own view with standard accounts of sentimentalism (in Chapters 4 and 5 of *Sentimental Rules*). At one point, he does directly address some of these issues:

The Sentimental Rules account does not purport to give a conceptual analysis of moral terms...unlike philosophical sentimentalist views, the Sentimental Rules account is not committed to a conceptual connection between moral judgment and motivation...the Sentimental Rules account accommodates the central sentimental idea that moral judgment is typically attended by emotion. (Nichols 99)

Here, he explicitly aligns himself with sentimentalist accounts of moral motivation, picking out the common emphasis on emotion. Emotion gives rise to our moral beliefs, and our moral beliefs drive our norms. These norms are the moral facts of the matter for Nichols. So in this regard, Nichols is clearly in line with the sentimentalists.⁵ However, Nichols also appeals to the dissociability of the affective system (i.e., the concern mechanism) and the normative theory, and this distinctness does pose some problems for pigeon-holing Nichols straight into the

⁵ For a more substantive discussion of this argument, see section 4.2 of the present work.

sentimentalist camp (100). Further problems with categorizing Nichols arise because of his silence on the normative theory central to his account. He says at one point that “moral argument and moral discourse can largely be construed as reasoning over the content and implication of the internally represented normative theory” (Nichols, 107-8). Thus, the normative theory is independent of the affective system in a substantial enough sense to be invoked without invoking the affective system as well. Nichols is an externalist, and the point he wants to drive home with all this is that the content of moral judgments does not consist solely in the emotions that they invoke, so the dissociability of the normative theory from the affective system should not automatically exclude him from counting as a sentimentalist—indeed, it seems most sentimentalists would want to allow for dispassionate discussion of normative theories.

One distinct feature of the Sentimental Rules account of moral judgment (that will end up being important for present purposes) is its rejection of objectivity. After recounting research on the survival of moral judgments after the forsaking of objectivity, Nichols concludes:

For both the character and the cultural evolution of moral judgment depend on the emotions that we happen to have. Thus there is some reason to think that we should turn away from our commonsense commitment to moral objectivity. Such a move away from objectivity appears to be less disruptive than some have feared. The evidence recounted in the last section suggests that spurning objectivity by no means eradicates the power and authority of harm norms. This is hardly surprising if, as I have argued, harm norms acquire their distinctive status because of their connection with powerful emotional responses. Those emotional responses are not extirpated by coming to reject moral objectivity. The emotions that make moral judgment distinctive continue to burn even in people who abandon the lay metaphysics of morality. (Nichols 198)

The absence of objectivity does not reduce all moral violations to simple conventional violations. Even without objectivity, moral judgment retains its distinctiveness because of its close kinship

with strong emotional responses, such as our responses to suffering in others that make our harm norms culturally fit.

We have seen that Nichols' *Sentimental Rules* account of moral judgment commits him to something very similar to sentimentalism (with qualifications), and he rejects objectivity. These are the only two qualifications that I will need to make my later arguments (see section 4.2) sound.

3.4 An Excursion into Epidemiology

In Chapter Six of *Sentimental Rules*, Nichols sets out to sketch an account of the genealogy of norms and ends up endorsing what he calls an “epidemiological” approach. He sets his genealogical account apart from the herd by trying to develop an account of why certain norms persist—he is not attempting to give an *origin* story about harm-based norms. Indeed, Nichols will conclude the chapter by saying, “One promising way to investigate the genealogy of norms is by considering not the origin of norms, but rather, what makes certain norms more likely to prevail” (Nichols 139). He claims:

Instead of seeking an account of the origin of moral norms, we might try to determine which features make certain norms more likely to prevail than other norms. This might then help us explain why harm norms prevailed. To pursue this project, we need to turn to cognitive anthropology. (Nichols 121)

He wants to take an account of cultural transmission from cognitive anthropology, the epidemiological approach (Nichols 121). This approach focuses on mental representations and their fitness for survival given the psychological make-up of the creatures in question (Nichols 122).

Nichols pulls the epidemiological account from the work of Dan Sperber and his associates. Nichols wants to highlight two features of the epidemiological account:

First, on the epidemiological approach, we need to attend closely to the universal features of human psychology. Second, we can expect that cultural items that are more easily remembered will have greater cultural fitness. (Nichols 123)

Pascal Boyer, a follower of Sperber, has offered an account of religion using the epidemiological approach that he thinks shows that

[T]he ‘cultural mind’ typical of recent human evolution is not so much an unconstrained mind as a mind equipped with a host of complex, specialised [sic] capacities that make certain kinds of mental representations likely to succeed in cultural transmission. (Boyer 93)

Given the focus in Boyer’s paper on the things Nichols most wants to highlight, it will be worthwhile to take a brief look at Boyer’s account of religion in order to get a better grasp of what cognitive anthropologists are aiming for in providing epidemiological explanations of cultural phenomena.

Boyer ties religious concepts to “*intuitive ontology*, that is, a set of categories and inference-mechanisms that describe the broad categories of objects to be found in the world...and some causal properties of objects belonging to these categories” (100). For Boyer intuitive ontology, perhaps over-simplified, is a set of intuitive categories that break up the world. He provides an explicit, though not exhaustive, list that includes intuitive versions of biology, physics, and theory of mind (Boyer 100). These categories allow us predictive abilities for all things that fit into these categories. He finds a niche for religious concepts:

Religious concepts are constrained by intuitive ontology in two different ways: [1] they include explicit *violations* of intuitive expectations, and [2] they tacitly activate a *background* of non-violated ‘default’ expectations. (Boyer 100-1)

Religious concepts situate entities within existing ontological categories and then violate these categories. However, the situation of these entities within the ontological categories allows the believer to draw inferences against the background knowledge presupposed by the respective ontological category. For example, take the Christian tradition of the snake talking to Eve in the

Garden of Eden. The snake is situated into the category of animals and counterintuitively assigned the property of being able to speak a human language. The believer can, however, make further inferences as a result of the background information available to her (e.g., the snake had scales, it had fine, pointed teeth, etc.), that do not violate the intuitive aspects of that ontological category.⁶

Thus far, Boyer's account has highlighted one of the central features of epidemiological accounts that Nichols has wanted to accentuate and use for his own devices: universal features of human psychology. The universal feature being appealed to here is an intuitive ontology. So what about Nichols' second resource: an account of cultural fitness? Boyer explains:

From an anthropological viewpoint, it would seem that both explicit counterintuitive and default background elements are necessary for cultural fitness...In order to have minimal cultural fitness, descriptions of religious entities found in the cultural input must be *attention-grabbing*, otherwise they will be forgotten or discarded. This quality is provided by violations of intuitive expectations. Also, the concepts built by subjects on the basis of that input must have some *inferential potential*, and this is provided by default assumptions. (Boyer 104)

Boyer claims that the counterintuitive aspects of religious concepts are what make them easily remembered and their cultural fitness lies in part in this fact. He provides experimental data from research establishing the greater ease of recall for counterintuitive items (Boyer 105-6).

Nichols is quick to identify (but insufficiently explains) two virtues of an epidemiological approach. First, this methodology can account for cross-cultural variation of harm-based norms or, at the very least, is consistent with such variation, since it merely attempts to demonstrate the staying-power of certain harm-based norms (Nichols 123). Second, cultural transmission approaches demonstrate the power of the cumulative effects a harm-based norm may have on a culture, whereas origin stories have to demonstrate its effect at its initial endorsement (Nichols

⁶ For additional examples, see Boyer pages 100-3.

124). He hopes that his expansion of the epidemiological account can set up the necessary apparatus to support these virtues.

Nichols wants to augment the epidemiological account endorsed by Sperber and his followers (like Boyer), by expanding the cognitive mechanisms used in cultural transmission to include affective mechanisms, such as the one he has previously used to great effect throughout his Sentimental Rules account. He explains:

As Sperber and Boyer develop the epidemiological approach, they recommend that to understand cultural transmission, we attend to species-general information-based cognitive mechanisms like intuitive physics and intuitive psychology. However, there is a quite different class of basic mental mechanisms that are almost certainly crucial on an epidemiological approach: emotion systems. Sperber and Boyer devote little attention to the role of affective mechanisms as forces of cultural attraction. Yet a number of affective mechanisms are regarded as universal denizens of human psychology. And the idea that affective mechanisms partly determine which cultural items succeed is certainly consistent with the epidemiological approach. (Nichols 124)

Emotions can factor into cultural transmission for an obvious reason: we simply care more about those things that elicit a strong emotional response in us (Nichols 124). Nichols cites data from Ekman showing that certain emotional responses are common across cultures. These include disgust, fear, sadness, and anger (Nichols 125).

At this point, Nichols wants to show that the ways in which something that elicits a strong emotional response will be more culturally fit for transmission, and he lands upon retention as a key factor in this fitness (Nichols 125). After surveying some of the evidence regarding affect and retention, he draws two conclusions:

First, the crucial mnemonic dimension for cultural fitness is long-term retention, and that is exactly the dimension that affect most clearly facilitates. Second, the retention benefits cannot be attributed to a self-serving bias to remember things that are affectively pleasing. For the evidence indicates that retention benefits are generated when the stimuli elicit negative affect. (Nichols 125-6)

Given the data showing a correlation between emotional response and long-term retention, Nichols asks how affect can increase the likelihood of retention. He briefly discusses four possible explanations of what may serve as the mechanism linking the two, but he reaches no real conclusions on the issue. He summarizes the argument for the augmentation of the epidemiological account thus:

Affective systems provide a rich source of likely human universals, and there is considerable evidence that affective systems contribute to greater retention. We can combine these two facts to generate a fairly interesting prediction: cultural items that are likely to elicit a basic emotion will be more culturally fit than cultural items that are affectively neutral. This is significant because insofar as there are eliciting conditions for basic emotions that are broadly consistent across cultures, we can expect cultural items that have those features to be better remembered and hence have greater cultural fitness. (Nichols 127)

Given the association of emotions possessing cross-cultural elicitation patterns with norms enjoying greater retention, Nichols then sets out to show that the retention enabled by emotional response helps to explain the evolution and transmission of our norms.

Nichols argues that affect-backed harm-based norms will be more persistent because they “prohibit an action that is emotionally upsetting” (128). He summarizes with his “‘Affective Resonance’ Hypothesis”:

Normative prohibitions against action X will be more likely to survive if action X elicits (or is easily led to elicit) negative affect. (Nichols 129)

He is quick to add that he is not arguing that affect-backed norms are the only culturally fit norms; he merely wants to establish that affect-backed norms are more likely to be culturally fit than their non-affect-backed counterparts (Nichols 129).

All that being said, Nichols recognizes that he still needs to back up the possibility of affect-backed norms possessing greater cultural fitness with empirical research (129). Since many factors seem to convolute our moral norms, they may not be the paradigm cases to which

we should turn for evidence regarding the success of affect-backed norms (Nichols 129).

Nichols proposes a different line of inquiry:

Moral norms are contentious on multiple fronts. So if we want to understand how norms are transmitted in a culture, it will be best to explore first the transmission of less controversial norms. I will focus on norms governing manners in our culture. (Nichols 129-30)

To avoid the complicated terrain of moral norms, Nichols sets out to establish his expansion of epidemiological accounts by first examining manners norms. Since disgust is a basic emotion (and presumably, a universal one), Nichols wants to latch on to disgust-backed manners norms to argue that affect-backed norms will exhibit a greater cultural fitness than will non-affect-backed norms. Nichols appeals to the work of Norbert Elias to show that disgust-backed manners norms have a tendency to persist in a culture through time (130-3). Nichols rightly acknowledges that while Elias' work does provide substantial evidence that disgust-backed manners norms demonstrate significant staying power, it does not show that disgust-backed manners norms show greater cultural fitness and staying power than non-affect-backed manners norms (133). Nichols conducts his own research into establishing the second necessary part of his conclusion (133-7), and I am willing, for present purposes, to grant that he establishes this conclusion.⁷

After going through a number of cases regarding disgust and other affect-backed norms, Nichols concludes, "the emotions played a significant historical role in determining which norms survived into the present" (138-9). That being said, he continues:

The naturalistic genealogical approach I have taken in this chapter offers a radically different explanation of the striking fact that so many of our norms are closely coordinated with our emotions. At least part of the reason our norms are

⁷ My reasons for omitting an exegesis of his discussion are two-fold. First, I am merely taking into account simple space constraints. Second, nothing Nichols explains or establishes in this section of *Sentimental Rules* seems particularly contentious or controversial. Since the goal of the present work is simply to examine which ethical theory provides the best explanation for the available evidence of moral evolution, I want merely to examine the success of each theorist at doing just this—explaining moral evolution. With this in mind, I want to grant Nichols his anthropological conclusions, which will allow him to formulate an explanation of moral evolution.

connected with our emotions is that affect-backed norms have greater cultural fitness and hence are more likely to survive than norms that fail to engage our affective responses...Our emotions played a historical role in determining which norms survived into the present. (Nichols 139)

The connection, essential to Nichols' sentimentalist account, between the norms we have and our emotions is in part described by the effect of affect on the cultural fitness and thus, survival of our norms. So what Nichols has offered us in this chapter is an account of the persistence of our affect-backed norms through time.

3.5 Patterns of Moral Evolution

In Chapter 7 of *Sentimental Rules*, Nichols wants to utilize the anthropological evidence just set out (and his expansion thereof) to provide, using the Affective Resonance model, an explanation of three phenomena. The first is the cross-cultural ubiquity that harm norms exhibit. Harm-norms exist in every studied culture (with a few suspect exceptions) and this fact should be explained by an account of the genealogy of norms (Nichols, 142). Second, despite the ubiquity of harm norms, some explanation must be given for the variations in harm norms cross-culturally (Nichols 142-3). The third phenomenon he puts thus:

It has become commonplace in discussion of moral evolution that, in the long run, moral norms exhibit a characteristic pattern of development. The familiar account is that harm norms tend to evolve from being restricted to a small group of individuals to encompassing an increasingly larger group. That is, the moral community expands...As far as I know, there has not been a systematic cross-cultural study of normative evolution that confirms this. But I suspect that the pattern is real. It is manifestly the case in our own culture that the moral community has show this kind of development. (Nichols 143)

This phenomenon, though lacking a systematic confirmation, is intuitively plausible to Nichols, and he zeroes in on the social history of norms regarding cruelty to animals and norms regarding corporal punishment to establish it in Western culture.

As regards cruelty to animals, Nichols cites research detailing the emergence of legislation prohibiting various forms of inhumane treatment of animals (144-5). His attention is focused on the social history of Britain in the nineteenth century, but he claims that these findings can be extrapolated to Western history more generally. He claims, “One of the striking features of this history is that it makes clear that it is naïve and simplistic to assume that there is a single origin story for norms prohibiting cruelty to animals” (Nichols 145). He briefly examines the literature surrounding this issue, and finds numerous explanations of the norm regarding the humane treatment of animals. He generalizes:

This serves to reinforce worries about trying to find *the* origin of harm norms. For in the case of norms against harming animals, there is no single origin. We might then, be leery of undertaking to find *the* origin of norms against harming people.⁸ (Nichols 145)

Nichols reinforces the idea of providing an account of the historical development of norms and the changes they have undergone instead of attempting a historical or philosophical origin story.

The research that Nichols appeals to informs his discussion of corporal punishment cross-culturally. He traces norms regarding physical punishment and execution from the late Middle Ages through the present in Europe and the United States (Nichols 145-6). In this sphere, Nichols finds that harsh physical punishments have declined since the late Middle Ages, with things like amputation and torture altogether disappearing by the present day (146). He also acknowledges a general decline in capital punishments (Nichols 146).

After illuminating these “important historical trends in harm norms” (146-7), he concludes:

⁸ I would like to point out that just because no single origin story can be divined in the case of norms prohibiting cruelty to animals, there is no good reason to think that there is not a single origin story regarding the origin of norms prohibiting harm to humans. Indeed, I think it is much more plausible that there is a single origin story regarding harm norms toward humans than it is that there is one for harm norms toward animals. However, I do not wish to harp on this issue. I raise it merely to point it out and then dismiss it.

We want a genealogical account that can explain such similarities and historical trends as well as the cross-cultural differences in harm norms across cultures. That is, of course, a preposterously ambitious project. My goal here will be to try to provide a partial explanation of the similarities and changes in a way that is fully compatible with the differences. (Nichols 147)

Nichols wants to provide an account of the trends in moral evolution that he has pinpointed using his Sentimental Rules account of morality. This explanation will constitute his inference to the best explanation of moral evolution in favor of his own brand of sentimentalism.

3.6 The Best Explanation (According to Nichols)

We have already seen how the moral realists interpret the supposed expansion of the moral community over time. It is now time to allow Nichols a chance to explain the relevant patterns in moral evolution. Earlier in his book, Nichols examines empirical data investigating the affective responses of individuals perceiving suffering in others. From this data, Nichols concludes that human beings seem to be hardwired from a very young age to recognize and respond to the suffering of others (even some nonhuman animals) with distress (154-5). The concern mechanism he has posited earlier in the book seems to be a human universal as evidence for it is exhibited cross-culturally. Appealing to evidence for its being cross-cultural, Nichols argues that the concern mechanism and its typical reactions to suffering in others fit neatly into his epidemiological account (155). In his discussion of manners norms, he argued that those norms prohibiting an action that elicited negative affect will be more likely to persist than those that do not. He claims:

In keeping with this, I suggest that our emotional sensitivity to suffering in others helped to secure for harm norms the central role they occupy in our moral outlook. Suffering in others leads to serious negative affect, so harm norms would prohibit actions that are likely to elicit negative affect. Thus, if affect-backed norms are more culturally fit, the norms against harming others should have increased cultural fitness over norms that are not backed by affective

response. That is, harm norms, like norms against disgusting behavior, enjoy Affective Resonance, which enhances their cultural fitness. (Nichols 155)

Harm norms are affect-backed just like disgust-backed manners norms. Therefore, like disgust-backed manners norms, harm norms should have an increased cultural fitness since they will be more easily remembered and subject to long-term retention due to the emotional response violations of them will produce in an onlooker. Nichols puts it thus, “For as harm norms entered the culture, their emotional resonance would have contributed to their cultural cachet” (155).

The preceding amounts to Nichols’ argument for the ubiquity of harm norms. In virtue of our basic responses to what they prohibit, harm norms will more likely than not persist in a culture in which they emerge. This phenomenon explains their widespread existence and continuance over time. Nichols also claims that the account he has offered provides an explanation of the great variation of harm norms cross-culturally:

For the claim is simply that harm norms will have enhanced cultural fitness. This allows for considerable normative diversity, because it concedes that cultural processes play a vital role in the development of norms. Because cultural processes implicate a complex and variegated set of forces, it is hardly surprising, on this view, that there is so much diversity in the norms found in different cultures. (Nichols 155-6)

Since cultural change is influenced by a great many factors, it is unsurprising that harm norms would vary from culture to culture, since the different factors will cause the norms to be manifested differently in different cultures.

With this analysis, Nichols provides us with an explanation for two of the three phenomena he seeks to explain: the ubiquity of harm norms and the cross-cultural variation of harm norms. Now, he must explain the distinctive evolution of harm norms. However, he elaborates:

Two central characteristics of the evolution of harm norms might be teased apart. First, as we’ve seen, harm norms seem to become more inclusive, that is, cultures

seem to develop a more inclusive view of the set of individuals whose suffering matters. Second, harm norms come to apply to a wider range of harms among those who are already part of the moral community—that is, there is less tolerance of pain and suffering of others. (Nichols 156)

Conveniently enough, the two cases of the evolution of harm norms that Nichols has explained earlier fit these two characteristics perfectly. The prohibitions on cruelty towards animals are examples of the expansion of the moral community to include as relevant the suffering of those not previously recognized as part of the moral sphere, and this inclusion can be explained by the Affective Resonance account because it maintains that the response cues we demonstrate when confronted by the suffering of others occur not merely when confronted by humans, regardless of their race or gender (thus the expansion of the moral community to include all humans such as Singer cites), but also when we are confronted by suffering in animals as well (Nichols 156-7).

Nichols explains:

Consequently, so long as one remains within the confines of species that are likely to inspire reactive distress or concern (this will probably exclude lots of insects), harm norms that are more inclusive will have a survival advantage over other norms. Thus, the Affective Resonance account seems well suited to explain the evolution of norms against cruelty to animals. These norms plausibly resonated with a preexisting tendency to respond emotionally to the suffering of animals. (Nichols 157)

Exposure to the suffering of animals causes us distress, and this negative affect gives cultural fitness to those norms prohibiting cruelty to animals.

As regards corporal punishment⁹ and decreased tolerance of suffering amongst those already considered part of the moral community, Nichols says the following:

Witnessing or hearing about harsh corporal punishments is likely to trigger reactive distress or concern. Indeed, because our immediate responses to another's suffering is largely insulated from background knowledge...even if one is convinced of a convict's guilt, witnessing severe punishments on that individual

⁹ For a radically different interpretation of Western restrictions on corporal and capital punishment, see Michel Foucault's *Discipline and Punish: The Birth of the Prison*.

is likely to elicit reactive distress or concern. As a result, norms prohibiting those kinds of punishments will gain some fitness advantage. (Nichols 157)

Regardless of the contingent facts surrounding another person, his or her suffering is likely to elicit negative affect in an observer. As has been stated before, this affect will enhance the cultural viability of a norm prohibiting this suffering. Thus, Nichols' account can offer a substantial explanation of the three basic patterns of moral evolution that he outlined at the beginning of Chapter 7: the ubiquity of harm norms, cross-cultural variation of harm norms, and the expansion of the moral community.

Furthermore, Nichols' account can even be used to explain why some moral norms do not survive the test of time. He surveys what little historical evidence he can find regarding norms that are no longer in existence, such as restrictions on pride and eating for pleasure, and he concludes:

These norms that have fallen away, of course, also have the feature that they prohibit actions that are unlikely to elicit reactive distress or concern (or any other core emotion). Hence, as the Affective Resonance account would suggest, the norms that are connected to reactive distress and concern seem to survive well, whereas many norms that are not so connected have disappeared. (Nichols 158)

Naturally, those norms that have not survived are those not connected with any core emotional response or concern mechanism being activated in an onlooker. Thus, they have less potential for long-term retention and are more likely to be forgotten.

At this point, Nichols adds a point that may be well worth considering, since it clarifies how his account is supposed to work:

It bears emphasizing that the Affective Resonance explanation is not that individuals typically recognize that certain norms fit well with their emotions and accordingly decide to adopt those norms. The idea that people deliberately try to achieve some kind of equilibrium by bending their norms to fit their affect seems rather implausible. Rather, the Affective Resonance proposal approaches the phenomena from the broad vantage of cultural evolution. The central idea is

simply that on balance, affect-backed rules will be more attractive and this advantage will accumulate down the ages. (Nichols 159)

Nichols' account of the genealogy of morals does not work on an individual agent basis. The account is explanatory only in the long run. The culture as it changes over time is the basic unit in this account, and norms backed by affect will have a greater likelihood to survive in a culture as it changes through time.

3.7 The Attack on Realism

After providing the Affective Resonance explanation of the evolution of (at least some) moral norms outlined above, Nichols turns his attention to attacking the explanations that moral realists provide for moral evolution and showing that his explanation is a better one: as he says, "After all, the form of argument under investigation is inference to the best explanation" (160). Nichols proceeds by undermining the arguments from realists like Sturgeon (he calls their accounts "simple appeals to moral facts") first, and then he attacks Railton's argument separately because Railton's social unrest model possesses more explanatory power when providing an account of the evolution of moral norms.

Nichols singles out simple appeals to moral facts as being "thin" explanations because of their lack of predictive ability:

For without an independently established story about moral facts, the simple realist story does not generate specific predictions about the direction of moral evolution. What one needs to know antecedently is, what are the moral facts? The trouble is, of course, that there is no generally accepted story about moral facts. That's part of the reason, after all, that realists appeal to moral progress—to shore up their claim that there *are* moral facts. As a result, the simple realist story provides only the barest sketch of an explanation—if there were moral facts, then that would explain why there are robust trends. (Nichols 160-1)

The analysis provided by such realists is an inadequate explanation, because it lacks predictive specificity—it can only explain large-scale movement of moral beliefs towards the moral fact of the matter. Since realists have agreed upon no substantive conclusions regarding what exactly the moral facts are, no specific predictions can be offered as to evolution of moral norms. Thus, simple appeals to moral facts lack predictive specificity. In fact, Nichols, earlier in the text, has already provided us with an account of what the simple realists have to offer in terms of the explanatory scope of their accounts:

Certain core harm norms are virtually ubiquitous because everyone has figured out this much about morality. The norms evolve in a characteristic way because people tend to get closer to the truth. And there is considerable variation because the process is a difficult one, distorted by self-interest and cultural idiosyncrasies. (Nichols 153)

Such are the explanations that a moral realist has to offer. These explanations, however, are rather thin and tell us very little about the moral facts themselves and how our moral beliefs do change over time.

Conversely, the Affective Resonance account can deliver a much more robust explanation that can provide much more predictive specificity (Nichols 161). His account predicts that those norms that are affect-backed will be more likely to survive through time and be more culturally fit (Nichols 161). This specificity can be a burden of a theory if its predictions can be experimentally disconfirmed, but this is not the case with the predictions of the Affective Resonance account, which are in fact verified by the empirical research that Nichols has surveyed (161). Nichols argues:

It is a commonplace in philosophy of science that, all else being equal, a theory with greater predictive success and explanatory depth is to be preferred. On these grounds, the Affective Resonance is to be preferred to the realist account. For the realist account only explains why we find moral evolution, whereas the Affective Resonance account explains both why we find moral evolution, and why it takes the course it does. (Nichols 161)

On the Affective Resonance account, moral evolution occurs because we are sensitive to the suffering of others and this evolution proceeds the way it does because affect-backed norms are more culturally fit than non-affect-backed norms. This account possesses more explanatory value than the simple realist accounts.

The characterization that Nichols provides of the realist arguments here may be accused of committing a straw man fallacy against the realists that he has rather indiscriminately lumped together. I want to particularly single out Peter Singer's argument as one that Nichols insufficiently addresses. Singer is a utilitarian, so for him the moral facts are that pain is bad and causing pain is morally bad. Though Singer may be at a loss to explain the emergence of some harm norms, it seems he has an easy time explaining the emergence of norms prohibiting cruelty towards animals. Singer could claim that the moral evolution occurred because the erroneous and widespread belief that animals could not feel pain in the same way that humans do was corrected, and thus our moral beliefs about the standing of animals changed to be further in line with the moral fact of the matter—i.e., that causing animals pain is morally bad. Though this argument is only a sketch, I think it is sufficient to show that Nichols underestimates some of his enemies. That being said, I do think that Nichols' explanation from the Affective Resonance account is vastly stronger and more revealing in its explanatory power than even Singer's account. While Nichols may be underestimating some of the views he is attacking, his account provides a better inference to the best explanation than any of the accounts of this first group of realists.

Nichols treats Railton's account separately because it "manifestly does not suffer from a lack of predictive specificity" (161). He acknowledges:

Railton's proposal is intriguing. First, it does not suffer from the lack of predictive specificity that other realist accounts do. And there can be no doubt that social unrest has played a vital role in shaping the norms that we have. This is probably true for both moral and nonmoral norms. The proposal also helps to reinforce that, just as there are plausibly multiple different origins for moral norms, there are plausibly multiple factors influencing the evolution of moral norms. (Nichols 162)

All praise for Railton's account aside, Nichols pinpoints three explanatory failings on the part of Railton's theory. The first is a failing of Railton's account to explain the extent to which certain social reforms tend to go when they occur. Instead of merely doing what is required to maintain domestic peace, change is usually more pervasive (Railton 193). Nichols rightly points out that the social unrest account cannot make much sense out of this tendency (163).

The second failing of Railton's social unrest account emerges from the fact that many instances of the evolution of moral norms do not involve any real potential for social unrest. Railton cannot explain why the moral community should have expanded to include the entire species, since many groups will never have the potential to create unrest in most cultures (Nichols 163). Nichols' Affective Resonance account, as we have seen, can, however, explain this phenomenon. Again, conveniently, the examples of cruelty to animals and corporal punishment are perfect models for this failing as well. Under Railton's model, prohibitions on cruelty to animals would have to be explained in terms of the potential for some group to apply social pressure, but it is not clear who would constitute such a group and our norms only emerged after the ability of animals to create social unrest was effectively over (Nichols 163). Nichols' account, as previously illustrated, does offer a compelling account of how these norms may have emerged. As regards corporal punishment, it seems unlikely that anyone facing corporal punishment would be in a position to incite social unrest, but once more Nichols account has already provided us with an illustration of why corporal punishment has died out

(Nichols 164). I wish to point out that Nichols' response to Railton on this point is not as charitable to Railton as it should be. While Nichols focuses only on the possibility of the condemned exerting social pressure, it is possible that the friends, family, and followers of the condemned may be able to exert sufficient social pressure to change things. Instances of the deaths of martyrs strengthening social movements are abundant, and it is obvious that Nichols neglects this potential argument from Railton.

The third failing of Railton's account is its inability to explain the close correlation between our emotions and our norms—which Nichols' Affective Resonance account has no trouble illuminating (Nichols 164). Once again, Nichols may not be charitable enough to Railton's account because Railton explicitly ties social rationality to human interests—this is how it explains the phenomenon of the humanization of morality—even though those interests are objectified.

3.8 Conclusion

Nichols' Affective Resonance account of moral evolution seems plausible. It explains why those norms that prohibit actions that elicit a negative affect in us will be more culturally fit and, thus, more likely to persist over time. Furthermore, this account can be used to offer specific explanations of particular phenomena evident within moral evolution, such as the expansion of moral communities and the ubiquity of harm norms. In this regard, the Affective Resonance account offers a better inference to the best explanation than do any of the competitor moral realist theories. Should we therefore accept Nichols' sentimentalist ethics incorporating the Affective Resonance account as defeating the different varieties of realism examined in Chapter 2?

Chapter 4

Conclusion

4.1 Introduction: So Who Won?

If we grant that the arguments attempted by the moral realists and Nichols are legitimate in their goals, it is very clear to me that Nichols has triumphed over the various forms of realism described here. His account does exhibit wider explanatory and predictive ability, and thus should be accepted as providing the best inference to the best explanation.

However, I want now to raise an objection against the arguments that we have seen thus far. I will argue in the following sections that there are really two uses of the inference to the best explanation of moral progress at play in the arguments I have examined. The first use renders the arguments circular, so any attempt to refute an opponent using such an argument results in the offending theorist begging the question against his opponent. The second use renders the inference to the best explanation sound, but it does not leave us in a position to say much about moral facts, regardless of their nature.

4.2 The First Use and the Charge of Circularity

As I addressed in the first chapter, the moral realists and the sentimentalists use two different terms for the phenomena examined in this thesis. The moral realists describe moral evolution as a process of moral progress. It is named moral progress, because we are getting, progressing, closer and closer to the alignment of our moral beliefs with the moral facts. We advance toward the moral fact of the matter as we move forward through time. Nichols does not refer to these phenomena, which he explains with his genealogy, as moral progress. He describes it as moral evolution and identifies each phenomenon as a pattern thereof.

Why these terminological differences? It emerges because the two terms are used to describe two very different things given the respective theoretical apparatuses of the realists and the sentimentalists. For the realists, moral beliefs change because we are getting closer to the moral facts. The truth of certain moral facts and our increasing awareness of them cause us to reconcile our moral beliefs with them over time. Conversely, Nichols offers an account of moral evolution to show not merely the change in our moral beliefs, *but also changes in the moral facts themselves*. However, seeing this implication of Nichols' view takes some additional argument.

I said earlier that one of the features that will separate the moral realists from a sentimentalist like Nichols is that the moral realists are committed to a more substantial sense of truth for the truth of moral facts. For Nichols, the truth of any moral fact, in some sense that he does not clearly illuminate, depends upon our emotional responses to it. Recall Nichols' earlier anthropology of moral norms that have fallen away with time. The shift in moral norms is dependent upon the potential for affective arousal by violations of the norms. Those norms that are affect-backed will exhibit more cultural staying power. It seems ludicrous for Nichols to maintain that long forgotten moral norms still hold sway over us—indeed, it seems impossible for Nichols to say that adherence to those norms is even justified anymore. For Nichols, moral beliefs change over time as we adjust our norms to more fully represent our emotional responses. Emotions drive our beliefs about morality, and our moral beliefs drive our norms. The norms are the moral facts of the matter.

The moral facts are dependent upon our emotions. To believe that we feel a certain way in response to something is to make an incorrigible inference—one cannot doubt that he has a certain emotional response towards something when confronted with it. On some level, we cannot be wrong about our emotional responses to things. As our moral beliefs change, the

moral fact of the matter changes. This is surely a minimal conception of the truth predicate according to Denham's discussion, and it is not clear that moral discourse is truly cognitive under Nichols' conception of moral judgments. Furthermore, Nichols explicitly rejects objectivity, and it is not at all clear to me that he can even substantially avoid a charge of relativism.¹⁰

The realists want our moral beliefs to track the moral facts, which are, in some sense, independent of us. Thus, their interpretations of moral evolution will always proceed in a certain way with our beliefs progressing (hopefully) toward a convergence on the moral fact of the matter. Conversely, the sentimentalists claim that the moral facts are dependent upon our moral beliefs—i.e., moral facts are mind-dependent. As such, if they can merely show that our moral beliefs have changed over time, then they can successfully demonstrate moral evolution, since such an account of moral evolution need not appeal to anything beyond historical variation in our moral beliefs. This metaethical disagreement is a fundamental difference between the competing theories. Furthermore, the main point that the realists want to push is that their thick concept of moral progress with its dependence on moral facts that are independent of our actual moral beliefs is superior to the sentimentalists' conceptions of the same things. The realists use their inferences to the best explanation to attack other theories on just these points. They use moral progress to argue for moral facts, and this use gets them into trouble.

As I said earlier, I want argue that there are actually two uses of the inference to the best explanation at play in the section of the debate we have examined thus far. At this point, we

¹⁰ It seems to me that Nichols wants to avoid reaching a conclusion such as this. I think he can make a plausibly strong case for some degree of pseudo-objectivity in his theory (utilizing a notion of universal emotional responses). However, the minimal truth-aptness of moral facts implicit in his sentimentalist theory seems to make him susceptible to the charge of relativism. For a better discussion of the problem of objectivity for sensibility theorists, see David Wiggins "A Sensible Subjectivism?" For an empirically informed account that is unabashedly relativist, see Jesse Prinz's forthcoming book *The Emotional Construction of Morals*.

have hit on the first one. This first use does not tell us much metaethically because of the presuppositions that facilitate it. First, the theorist hypothesizes moral facts. Next, he or she makes observations of some phenomena (here, patterns of moral evolution). Finally, the theorist claims that his or her metaethics is the best explanation—for realists it is moral progress (in their thick sense) and for anti-realists (here, sentimentalists) it is facts changing with beliefs. In this first case, the inference to the best explanation is circular, because the descriptions of the respective phenomena presuppose the metaethical theory being defended by the argument. When the various metaethical camps attempt to use the inference to the best explanation in this manner to argue for a particular direction of fit between moral beliefs and moral facts, they are not arguing divergent interpretation of the data, they are elucidating or spelling out the consequences of their different basic assumptions underlying the data. Since the realist is presupposing a direction of fit for moral beliefs and moral facts, he or she begs the question when using these explanations against theorists assuming a different direction of fit.

An example is helpful at this point. Sturgeon's argument as presented above is an instance of this first use of the inference to the best explanation and, as such, it is subject to a charge of circularity. His account of the changes in moral beliefs regarding slavery depends upon the fact that slavery was *morally* more despicable in the eighteenth and nineteenth centuries than it had been previously. Moral beliefs changed as a result of this moral fact. Thus, moral evolution is explained as an alignment of moral beliefs with the moral facts of the matter. Sturgeon is trying to use this inference to argue for the existence of mind-independent moral facts. However, he is appealing to mind-independent moral facts regarding slavery in his argument—that it was *morally* worse at one time period than it was in others. He helps himself to mind-independent moral facts and wants the conclusion he derives from this assumption to

establish the existence of mind-independent moral facts. Thus, his inference to the best explanation is circular, and for him to levy this argument against Harman is for him to beg the question against Harman.

4.3 The Second Use and Lack of Content

Is this what Nichols is doing with his Sentimental Rules account? Are all such inferences to the best explanation question-begging? I will answer both of these questions in the negative.

In response to the first question, Nichols' argument to the best explanation seems to be of a different form than those of Sturgeon, Slote, Smith, Singer, and Nagel. Nichols is offering an empirical explanation of empirical phenomena: his claim is that an epidemiological framework provides a sufficiently powerful explanation of the empirical phenomena of changing moral beliefs. Such an empirical explanation will not logically entail a metaethical theory, nor will it casually entail a metaethical theory. What such an empirical explanation can do is render some features of other accounts explanatorily idle. This is the second use of inferences to the best explanation that we can identify. Such a use is non-circular, because the explanation given is purely empirical in nature and no suppositions about moral facts are built into the argument itself. Nichols has not helped himself to any substantive claim regarding moral facts (so far as his inference to the best explanation is concerned) nor has he done anything beyond offering a purely empirical account of empirical phenomena.

It is helpful to now return to the argument from Gilbert Harman that Sturgeon attacks when offering his argument to the best explanation. Harman argues that moral facts are explanatorily idle—they do not figure into any explanation of an event or of the moral judgments of an event:

The observation of an event can provide observational evidence for or against a scientific theory in the sense that the truth of that observation can be relevant to a reasonable explanation of why that observation was made. A moral observation does not seem, in the same sense, to be observational evidence for or against any moral theory, since the truth or falsity of the moral observation seems to be completely irrelevant to any reasonable explanation of why that observation was made. The fact that an observation of an event was made at the time it was made is evidence not only about the observer but also about the physical facts. The fact that you made a particular moral observation when you did does not seem to be evidence about moral facts, only evidence about you and your moral sensibility. (Harman 7)

Though observational evidence can count for a scientific theory, observation cannot provide evidence for a moral theory, because the moral fact of the matter cannot explain the observation and why the agent in question made it. That a particular moral observation is made seems only to reflect on the psychology of the agent and not the moral fact of the matter. Thus, mind-independent moral facts are explanatorily idle, since they do not help us understand why certain moral judgments are made. Later, Harman adds:

Observational evidence plays a part in science it does not appear to play in ethics, because scientific principles can be justified ultimately by their role in explaining observations...by their explanatory role. Apparently, moral principles cannot be justified in the same way. It appears to be true that there can be no explanatory chain between moral principles and particular observings in the way that there can be such a chain between scientific principles and particular observings. Conceived as an explanatory theory, morality, unlike science, seems to be cut off from observation. (Harman 9)

There is no explanatory link between moral principles and the particular moral observations that we make, and this causal break is the source of the contrast between science and ethics—observations are explained by scientific principles. Mind-independent moral facts are explanatorily idle as regards the moral beliefs that we actually have. What is missing from Harman's argument is an empirical explanation of the moral observations or beliefs that we actually have.

With his inference to the best explanation, Nichols is rendering mind-independent moral facts explanatorily idle by offering an empirical account of the moral beliefs that we actually possess. His explanation provides more explanatory power and predictive success than do the accounts of the realist, and his explanation does not make an appeal to mind-independent moral facts. Since the better explanation does not need to appeal to such metaphysical entities for it to succeed, it renders the mind-independent moral facts explanatorily idle and superfluous. Nichols' explanation is more elegant in that it does not require the existence of such metaphysical entities (mind-independent moral facts). Thus, mind-independent moral facts are rendered superfluous and explanatorily idle.

It is important to note that Nichols does not see the break between his argument and the arguments of the moral realists. He regards them as being of the same piece. I hope that I have established that the inference to the best explanation of moral evolution is being used in two distinct ways in the literature that has emerged on the topic. That Nichols fails to notice the two uses of the argument is a serious shortcoming on his part. His argument is of a different type than those he is trying to defeat.

One serious problem that arises for Nichols is that the success of his argument (as an *empirical* explanation of empirical phenomena) is dependent upon non-moral descriptions. He has built no content into our moral beliefs. If we do insert content into our moral beliefs, can we provide a purely empirical account of the kind that Nichols offers? Sturgeon and the other moral realists insert content into our moral beliefs, only to succumb to circularity when offering arguments to the best explanation. To consistently insert content into moral beliefs as they figure into an empirical explanation of empirical phenomena, it seems that we must be committed to a

naturalistic reduction of moral terms and moral content to naturalistic non-moral terms and non-moral content.

What are the prospects for such a project for someone who wants to maintain some commitment to moral realism? One realist figure examined here offers an account that may provide an instance of the second use of the inference to the best explanation. Recall this earlier cited passage from Railton:

I will argue for a form of moral realism which holds that moral judgments can bear truth values in a fundamentally non-epistemic sense of truth; that moral properties are objective, though relational; that moral properties supervene upon natural properties, and may be reducible to them; that moral inquiry is of a piece with empirical inquiry; that it cannot be known *a priori* whether bivalence holds for moral judgments or how determinately such judgments can be assessed; that there is reason to think that current moralities are wrong in certain ways and could be wrong in quite general ways...and that, while there are perfectly general criteria of moral assessment, nonetheless, by the nature of these criteria no one kind of life is likely to be appropriate for all individuals and no one set of norms appropriate for all societies and all times. The position thus described might well be called 'stark, raving moral realism', but for the sake of syntax, I will colorlessly call it 'moral realism'. (Railton 165)

In this summation of his views, Railton explicitly claims that moral properties at least supervene upon natural one and may be reducible to them. While he does not offer an actual naturalistic reduction of moral properties, he admits the possibility of providing one. Also, Railton, like Nichols, is providing a purely empirical account of empirical phenomena and does not build content into the beliefs he tracks. Railton's inference to the best explanation, therefore, seems to be of a piece with Nichols' inference to the best explanation.

Are we in a position to evaluate the success of Nichols' argument to the best explanation as it attempts to better the arguments of the various realists? There seems no point in comparing his argument with those of realists whose inferences to the best explanation employ such arguments in the first use outlined above. The arguments of Sturgeon and those arguments

similar to his are circular and thus not sound, so there seems no need to directly compare Nichols' account to their accounts—if his is not circular, it will be more successful than their accounts. What about Railton? Railton's argument is not obviously circular, and, furthermore, it seems to employ the inference to the best explanation in the same manner that Nichols uses it. Nichols' explanation does seem to have more explanatory power and predictive success than Railton's social rationality account. The only hope for Railton's account is track the conditions of social rationality and do a better job at confronting the empirical phenomena than does Nichols. Even though Nichols' attacks on Railton may not be incredibly charitable, his account is more thorough and elegant in the explanations it provides and more predictively useful than Railton's.

4.4 Directions for Future Research

As the previous section makes clear, inferences to the best explanation of moral evolution/progress may still have a place in moral philosophy, so long as the use of such arguments is restricted to the second use I outlined above. My personal metaethical leanings are towards realism, so that Nichols offers the superior inference to the best explanation does not sit well with me. With this in mind, some possible directions for future research include creating a more expansive account of Railton's social rationality that could be effectively used to accommodate all the examples that Nichols' Sentimental Rules account has no problems explaining. Also, though this topic has received extensive treatment, the prospects for a naturalistic reductivist moral realism need to be explored if a realist is to offer an argument of similar force to Nichols'. While I am skeptical about the reduction of moral terms to natural non-moral terms, I, like Railton, am willing to concede its possibility. However, any attempt at

naturalistic reduction of moral terms is usually countered with G. E. Moore's Open Question Argument and thus commits the naturalistic fallacy.¹¹ There are larger issues with attempts at naturalizing moral terms, and the theorists who attempt it have an uphill battle to fight.

4.5 Conclusion: What About Moral Evolution/Progress?

The arguments from the previous sections of this chapter invite us to despair at the lack of fruitfulness in arguing from moral evolution/progress. What can we legitimately infer from the abundant empirical research surrounding these patterns of moral evolution? I hope now that the relevance of my opening quote from Nagel is apparent. Without a better understanding of the nature of moral facts, we cannot hope to reasonably infer anything about the superiority of one or another theory of moral progress. To suppose that any of the ethical theories scrutinized here is absolutely correct is, as Nagel points out, a very vain position to take. What these theories succeed in doing is to open up the floor for debate and create intellectual room to work out the issues surrounding, and in some cases enshrouding, the very idea of moral facts. The belief that morality can progress is one of the driving forces in that progression: the hope that change is possible is one motivating factor for the change that occurs.

If inferences to the best explanation of moral evolution are still to have a place in moral philosophy, then they must be used in a much more restricted sense than the way in which they have been previously utilized. They must be employed in the second use I have outlined above if they are to avoid the charge of circularity and be at all informative. Even with the soundness of this style of argument established, it is unclear how informative such arguments may be if they cannot successfully provide us with a naturalistic reduction of moral terms.

¹¹ For a full discussion of the Open Question Argument and the naturalistic fallacy, see G. E. Moore's *Principia Ethica*.

Works Cited

- Boyer, Pascal. "Evolution of the Modern Mind and the Origins of Culture: Religious Concepts as a Limiting-Case." *Evolution and the Human Mind: Modularity, Language and Meta-Cognition*. Ed. Peter Carruthers and Andrew Chamberlain. Cambridge: Cambridge University Press, 2000. 93-112.
- Denham, A. E. *Metaphor and Moral Experience*. Oxford: Oxford University Press, 2000.
- Foucault, Michel. *Discipline and Punish: The Birth of the Prison*. Trans. Alan Sheridan. New York: Vintage Books-Random House, Inc., 1977.
- Harman, Gilbert. *The Nature of Morality: An Introduction to Ethics*. New York: Oxford University Press, 1977.
- Horwich, Paul. *Truth*. Oxford: Basil Blackwell, 1990.
- Moody-Adams, Michelle M. "The Idea of Moral Progress." *Metaphilosophy* 30 (1999): 168-85.
- Moore, G. E. *Principia Ethica*. Cambridge: Cambridge University Press, 1903.
- Nagel, Thomas. *The View from Nowhere*. Oxford: Oxford University Press, 1986.
- Nichols, Shaun. *Sentimental Rules: On the Natural Foundations of Moral Judgment*. Oxford: Oxford University Press, 2004.
- Prinz, Jesse. *The Emotional Construction of Morals*. Oxford: Oxford University Press, 2007. (forthcoming)
- Railton, Peter. "Moral Realism." *The Philosophical Review* 95 (1986): 163-207.
- Singer, Peter. *The Expanding Circle: Ethics and Sociobiology*. New York: Farrar, Straus and Giroux, 1981.
- Slote, Michael. "The Rationality of Aesthetic Value Judgments." *The Journal of Philosophy* 68 (1971): 821-839.
- . "Is Virtue Possible?" *Analysis* 42 (1982): 70-76.
- Smith, Michael. *The Moral Problem*. Oxford: Blackwell, 1994.
- Sturgeon, Nicholas. "Moral Explanations." *Morality, Reason, and Truth: New Essays on the Foundations of Ethics*. Ed. David Copp and David Zimmerman. Totowa, N.J.: Rowman and Allanheld, 1985. 49-78.

Wiggins, David. "A Sensible Subjectivism?" *Needs, Values, Truth: Essays in the Philosophy of Value*. 3rd ed. New York: Oxford University Press, 1998. 185-214.

Wright, Crispin. *Truth and Objectivity*. Oxford: Oxford University Press, 1994.

Vita

Franklin Worrell was born and raised in Mobile, Alabama. He received a Bachelor of Arts in philosophy with full college honors and a Bachelor of Arts in English, concentration creative writing from Louisiana State University in 2005. His honors thesis explored Kant's ethics and was directed by Professor Husain Sarkar. He continued his studies at Louisiana State University by pursuing the Master of Arts in philosophy.