

**Sentimental Rules:
On the Natural Foundations
of Moral Judgment**

SHAUN NICHOLS

OXFORD UNIVERSITY PRESS

Sentimental Rules

This page intentionally left blank

Sentimental Rules

*On the Natural Foundations
of Moral Judgment*

SHAUN NICHOLS

OXFORD
UNIVERSITY PRESS

2004

OXFORD

UNIVERSITY PRESS

Oxford New York

Auckland Bangkok Buenos Aires Cape Town Chennai

Dar es Salaam Delhi Hong Kong Istanbul Karachi Kolkata

Kuala Lumpur Madrid Melbourne Mexico City Mumbai Nairobi

São Paulo Shanghai Taipei Tokyo Toronto

Copyright © 2004 by Oxford University Press, Inc.

Published by Oxford University Press, Inc.

198 Madison Avenue, New York, New York 10016

www.oup.com

Oxford is a registered trademark of Oxford University Press

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior permission of Oxford University Press.

Library of Congress Cataloging-in-Publication Data

Nichols, Shaun.

Sentimental rules : on the natural foundations of moral judgment / Shaun Nichols.

p. cm.

Includes bibliographical references and index.

ISBN 0-19-516934-4

I. Emotivism. I. Title

BJ1473 .N53 2004

1704—dc21 2003056382

9 8 7 6 5 4 3 2 1

Printed in the United States of America

on acid-free paper

For my parents, Gib and Sally Nichols

This page intentionally left blank

Preface

A few basic questions about the nature of morality command the attention even of those not steeped in philosophical training. Is morality grounded in rationality or does it depend crucially on the emotions? How did our moral system evolve to its present shape and character? Is morality objective? Would the rejection of moral objectivity have dire consequences for our normative lives? These are the kinds of questions that grab the novice to moral philosophy, and they are the questions that guided the explorations charted in this volume.

My pursuit of these fundamental metaethical issues has been thoroughly interdisciplinary, exploiting resources from philosophy, cognitive psychopathology, developmental psychology, cognitive anthropology, and social history. This interdisciplinary approach was not driven by any principled ecumenism. Rather, in trying to get a purchase on philosophical issues about, say, the role of sentiment in moral judgment or the genealogy of morals, I found myself driven to other disciplines to obtain the kind of information that seemed relevant to the issues. It turns out that there is a great deal of extant empirical work that is philosophically instructive, as I hope to display in this volume. In some instances, the relevant empirical information simply hadn't been sought, though it

seemed easy enough to obtain. Initially, I suggested to friends in psychology that they might want to run this or that experiment. Oddly, I found that many psychologists do not particularly like being told which experiments to run. So in some cases I was reduced to running experiments by myself.

These wide-ranging investigations led to an account of the nature of morality in the Humean tradition on which the emotions play a vital role in morality. On the theory I develop, emotions play a critical role in both the psychological and the cultural underpinnings of basic moral judgment. Norms prohibiting harming others are, I argue, associated with a fundamental emotional response, and this connection invests such norms and other “Sentimental Rules” with a distinctive status. Furthermore, I argue that such Sentimental Rules enjoy an advantage in cultural evolution, which partly explains the cultural success and historical development of certain moral norms. The account that emerges, I maintain, has broad ramifications for philosophical ethics.

Given the breadth of the issues taken up, one casualty has been scholarly depth. I have neglected many important works in contemporary metaethics. The situation is rather worse for the other disciplines I plunder. The ground I cover in this slim book is immodestly vast. I am trying to address basic issues that were plumbed by Hume, Kant, and Nietzsche, to name a few of the more intimidating figures. My hope is that naturalistic methods really will shed light on these basic issues. But in attempting to contribute to these hoary debates, this hope has been accompanied by an abiding sentiment of humility.

Acknowledgments

Given the ambitiously interdisciplinary nature of this work, it is perhaps not surprising that my intellectual debts run wide. Many philosophers and psychologists have provided advice and comments on various parts of the book, including James Blair, Deborah Boyle, Bill Casebeer, Jonathan Cohen, Justin D'Arms, Ronald de Sousa, John Doris, Susan Dwyer, Luc Faucher, Owen Flanagan, Trisha Folds-Bennett, Tim German, Todd Grantham, Dan Guevara, Paul Harris, Dan Haybron, Ned Hettinger, Daniel Jacobson, Todd Jones, Christopher Knapp, Larry Krasnoff, Glenn Lesses, Edouard Machery, Ron Mallon, Elizabeth Meny, Aaron Meskin, Adam Morton, Richard Nunan, Marty Perlmutter, Chandra Sripada, Steve Stich, Kyle Swan, and Hugh Wilder. Many thanks to all of these friends and colleagues for their help. Thanks also to my editor, Peter Ohlin, for his support and his speedy work.

Michael Gill and I have discussed the issues and arguments in this book frequently over the last several years. He also provided extensive comments on the entire manuscript. The book is considerably better than it would have been without his input, and I am most grateful for his help.

My greatest intellectual debt is unquestionably to Steve Stich. Before I ever met Steve he was one of my intellectual heroes for his knack at

extracting philosophically important conclusions from scientific work. His work has been a model of how to do naturalistic philosophy. Over the last dozen years, Steve has been a frequent collaborator and valued friend. In addition to providing characteristically insightful feedback on the work here, Steve has been a constant source of encouragement and inspiration. I would not have known how to begin, much less finish, this project were it not for Steve.

Finally I am deeply thankful for the love and support of my family, Heather, Sarah, and Julia.

Some portions of the book are drawn from previously published work. I am grateful to the publishers for permission to use material from the following articles: Mindreading and the cognitive architecture underlying altruistic motivation, *Mind and Language* 16 (2001): 425–55, with permission from Blackwell; Is it irrational to be amoral? How psychopaths threaten moral rationalism, *The Monist: An International Quarterly Journal of General Philosophical Inquiry* 85 (2002, Peru, Illinois, 61354): 285–304, reprinted by permission; Norms with feeling: Towards a psychological account of moral judgment, *Cognition* 84 (2002): 221–36, with permission from Elsevier; On the genealogy of norms: A case for the role of emotion in cultural evolution, *Philosophy of Science* 69 (2002): 234–55, with permission from the University of Chicago Press.

Contents

1. Norms with Feeling: Toward a Psychological Account of Moral Judgment, 3
 2. Sparks of Benevolence: The Varied Emotional Responses to Suffering in Others, 30
 3. Is It Irrational to Be Amoral? How Psychopaths Threaten Moral Rationalism, 65
 4. Philosophical Sentimentalism, 83
 5. Sentiment, Reason, and Motivation, 97
 6. A Fragment of the Genealogy of Norms, 118
 7. Moral Evolution, 141
 8. Commonsense Objectivism and the Persistence of Moral Judgment, 166
- References, 199
- Index, 219

This page intentionally left blank

Sentimental Rules

This page intentionally left blank

I

Norms with Feeling

Toward a Psychological Account of Moral Judgment

When I am at a loss to know the effects of one body upon another in any situation, I need only put them in that situation, and observe what results from it. But should I endeavour to clear up after the same manner any doubt in moral philosophy, by placing myself in the same case with that which I consider, 'tis evident this reflection and premeditation would so disturb the operation of my natural principles, as must render it impossible to form any just conclusion from the phenomenon. We must, therefore, glean up our experiments in this science from a cautious observation of human life, and take them as they appear in the common course of the world, by men's behaviour in company, in affairs, and in their pleasures. Where experiments of this kind are judiciously collected and compared, we may hope to establish on them a science which will not be inferior in certainty, and will be much superior in utility, to any other of human comprehension.

—David Hume, *A Treatise on Human Nature*

1. INTRODUCTION

Hume concludes his introduction to the *Treatise* with the above call for a research program. While naturalistic philosophers of every stripe applaud Hume's empirical turn, historians of philosophy note that Hume does not mean "experiment" the way we do today.¹ Rather, for Hume, cautious

1. A second qualification concerns his term "moral philosophy" which, in Hume's usage is not restricted to issues about morality, but rather, refers to the study of humans more generally, or the "science of human nature" (Hume, *Enquiry Concerning Human Understanding*, 5).

observation of human life primarily consisted in a kind of human ethology. One quietly observes people's behavior in their ordinary environments. Nonetheless, there is no reason to think that Hume would have shrunk from applying the standard experimental method to the study of moral judgment. On the contrary, it is likely that Hume would have championed the use of controlled experiments in moral psychology, had the methodology been readily available in his time. Hume's interests in the nature of moral judgment ranged widely, including issues about the role of sentiment in moral judgment, the role of reason in moral judgment, and the origin of moral judgment. To address these issues adequately requires attention to empirical details. In this volume, I will draw on recent work that attempts to supply such empirical details. Although the bulk of this research has been done by social scientists rather than philosophers, I suspect that were Hume alive today, this research would be high on his reading list.

Many of the deepest issues concerning the nature of morality would be illuminated if we had an adequate account of the nature of moral judgment. So it is scarcely surprising that researchers before and after Hume have invested enormous effort in trying to produce an account of moral judgment. The exploration of moral judgment in psychology stretches back for a century, through Kohlberg and Piaget. The philosophical lineage is much longer and enjoys an even more distinguished cast, including Kant, Hume, and Aristotle.

Throughout the twentieth century, philosophical work in metaethics largely ignored the psychological literature on moral judgment. Part of the explanation for this, I suspect, is simply that much of the best known psychological work seems not to intersect directly with the issues philosophers care about. Much of the psychological work on the nature of morality, for instance, charts developmental changes, gender differences, and cultural variations in moral cognition. But this kind of work seems not to address the core of morality—it does not tell us what is at the heart of moral judgment. As a result, it does not tell us what we want to know in metaethics. Over the last twenty years, a tradition in moral psychology has developed that really does, I will maintain, help us understand the nature of moral judgment. The research explores the ability to appreciate the distinctive status of morality, as reflected by the capacity to distinguish moral from conventional transgressions. In this chapter, I will suggest that this method plausibly reveals a capacity for a kind of core moral judgment. I will also present the prevailing accounts that might explain core moral judgment. There is a long history of perspective-taking accounts of moral capacities. More recently, there is an important proposal that ties moral judgment fairly directly to an affective response. I will

maintain that the evidence on core moral judgment indicates that both of these approaches are inadequate. I will argue that the capacity for core moral judgment depends on both a body of information about which actions are prohibited (a “normative theory”) and an affective mechanism that confers a special status on the norms.

2. CORE MORAL JUDGMENT

In the psychological literature, the capacity for moral judgment has perhaps been most directly approached empirically by exploring the basic capacity to distinguish moral violations (e.g., pulling another person’s hair) from conventional violations (e.g., chewing gum in class). This tradition in psychology began with the work of Elliott Turiel and has flourished over the last two decades (Turiel, Killen, and Helwig 1987; Dunn and Munn 1987; Smetana and Braeges 1990; Nucci 1986; Blair 1993). Turiel explicitly draws on the writings of several philosophers, including Searle, Brandt, and Rawls to draw the moral/conventional distinction. The distinction is characterized as follows: “Conventions are part of constitutive systems and are shared behaviors (uniformities, rules) whose meanings are defined by the constituted system in which they are embedded” (Turiel, Killen, and Helwig 1987, 169). Moral rules, on the other hand, are “unconditionally obligatory, generalizable, and impersonal insofar as they stem from concepts of welfare, justice, and rights” (Turiel, Killen, and Helwig 1987, 169–70). Although Turiel adverts to philosophical precedent for this distinction, the attempt to draw an analytic distinction between morality and convention is fraught with controversy. Fortunately, it is a controversy we can ignore. For we do not need to supply an analysis to see the significance of the data. The research program generated by Turiel’s work indicates that people distinguish moral violations from conventional violations along several dimensions.

Rather than embark on an attempt to define the moral and conventional domains, the easiest way to see the import of the data on moral judgment is to consider how subjects distinguish canonical examples of moral violations from canonical examples of conventional violations. Hitting another person is a canonical example of a moral violation used in these studies. Other frequently used examples of moral violations are pulling hair, stealing, and pushing another child. The examples of conventional violations that have been studied are much more varied. Some of the examples are violations of school rules, such as not paying attention during story time or talking out of turn. Some of the examples are violations of etiquette, such as drinking soup out of a bowl. Other examples

are violations of family rules, such as not clearing one's dishes. What is striking about this literature is that, from a young age, children distinguish the moral violations from the conventional violations on a number of dimensions. For instance, children tend to think that moral transgressions are generally less permissible and more serious than conventional transgressions. Children are also more likely to maintain that the moral violations are "generalizably" wrong, for example, that pulling hair is wrong in other countries too. And the explanations for why moral transgressions are wrong are given in terms of fairness and harm to victims. For example, children will say that pulling hair is wrong because it hurts the person. By contrast, the explanation for why conventional transgressions are wrong is given in terms of social acceptability—talking out of turn is wrong because it is rude or impolite, or because "you're not supposed to." Further, conventional rules, unlike moral rules, are viewed as dependent on authority. For instance, if at another school the teacher has no rule against chewing gum, children will judge that it is not wrong to chew gum at that school; but even if the teacher at another school has no rule against hitting, children claim that it is still wrong to hit. Indeed, a fascinating study on Amish teenagers indicates that moral judgments are not even regarded as dependent on God's authority. Nucci (1986) found that 100 percent of a group of Amish teenagers said that if God had made no rule against working on Sunday, it would not be wrong to work on Sunday. However, more than 80 percent of these subjects said that even if God had made no rule about hitting, it would still be wrong to hit.

These findings on the moral/conventional distinction are neither fragile nor superficial. On the contrary, the findings are quite robust. They have been replicated numerous times using a wide variety of stimuli (see Smetana 1993 and Tisak 1995 for reviews). Furthermore, the research apparently plumbs a fairly deep feature of moral judgment. For moral 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. It is found in young and old alike. Thus, it seems that the capacity for drawing the moral/conventional distinction is part of basic moral psychology.

Most of the above research on the moral/conventional distinction has focused on moral violations that involve harming others, and that will be my main focus as well. However, it is clear that harm-centered violations do not exhaust the moral domain. To take one obvious example, we think

it is wrong to cheat on one's taxes, but this has little direct bearing on harm. Furthermore, recent evidence indicates that the moral domain may not even be stable across cultures (e.g., Miller, Bersoff, and Harwood 1990; Haidt, Koller, and Diaz 1993). In a clever study by Jonathan Haidt and colleagues, they found that low socioeconomic status (SES) subjects were more likely than high SES subjects to maintain that people engaging in offensive or disrespectful actions (e.g., having sex with a dead chicken or cleaning the toilet with the national flag) should be stopped or punished (Haidt, Koller, and Diaz 1993). Haidt and colleagues conclude that it is parochial to think that harm is central to drawing the moral/conventional distinction (e.g., Haidt, Koller, and Diaz 1993, 625). Although there may be some relativity in the moral domain, the cross-cultural work also indicates that in all cultures, canonical examples of moral violations involve harming others (see, e.g., Hollos, Leis, and Turiel 1986; Nucci, Turiel, and Encarnacion-Gawrych 1983; Song, Smetana, and Kim 1987). Indeed, even Haidt and colleagues found that the subjects in different cultures and different SES groups made similar judgments about violations involving harm—for example, in all groups subjects tended to say that a girl who pushes a boy off a swing should be punished or stopped.

Thus, even though the moral domain is hardly exhausted by harm-based violations, it is plausible that judgments about harm-based violations constitute an important core of moral judgment. For the appreciation of harm-based violations shows up early ontogenetically (as we will see in section 3), and it seems to be cross-culturally universal. Brian Scholl and Alan Leslie make a related point about “theory of mind,” the capacity to understand other minds (Scholl and Leslie 1999). They note that, although there are cross-cultural differences in theory of mind, all cultures seem to share a core theory of mind that emerges early ontogenetically (140). Something similar might be said about the findings on moral judgment—despite the cross-cultural differences in moral judgment, the evidence indicates that all cultures share an important basic capacity, what I will call “core moral judgment.” The capacity to recognize that harm-based violations have a special status (as compared to conventional violations) is an important indicator of the capacity for core moral judgment. As a first approximation, the capacity for core moral judgment can be thought of as the capacity to recognize that harm-based violations are very serious, authority independent, generalizable and that the actions are wrong because of welfare considerations.

In the remainder of this chapter, I will develop an account of core moral judgment that draws on evidence from developmental psychology

and cognitive psychopathology.² In contemporary cognitive science, evidence from development and evidence from psychopathologies play particularly central roles in guiding theorizing about the psychological architecture underlying a capacity. For these sources give us a glimpse into which capacities might be independent from one another and which capacities seem to be inextricably linked. In chapters 3 and 4, I will return to the evidence from development and psychopathologies, and I will argue that the evidence poses problems for philosophical accounts of moral judgment from both the rationalist and the sentimentalist traditions. But before we can set out the philosophical ramifications, we would do well to get clear on the psychology. This will require looking closely at the available psychological accounts of the underpinnings of core moral judgment. There are two prominent approaches to the psychological architecture underlying moral judgment that we must consider—the perspective-taking account and Blair’s VIM-account. Both accounts, I will argue, are not adequate to the data.

3. PERSPECTIVE: TAKING ACCOUNTS OF MORAL JUDGMENT

The Piagetian tradition in moral psychology ties moral understanding to the capacity for perspective taking, or imagining oneself to have the mental states of another (e.g., Piaget 1932; Kohlberg 1984; Selman 1980; Damon 1977; see also Rawls 1971, chapter 8). The capacity for perspective taking itself has received enormous attention over the last twenty years, in the research on the capacity for understanding other minds, or “mindreading.” Researchers in this tradition have explored in detail the capacity for attributing mental states to others and predicting others’ behavior (e.g., Baron-Cohen, Leslie, and Frith 1985; Bartsch and Wellman 1995; Goldman 1989; Gopnik and Wellman 1994; Gordon 1986; Harris 1992;

2. For the philosophers, I should acknowledge that I am being deliberately ingenuous about some long-standing disputes in metaethics. In particular, I will not engage the noncognitivist view that, despite appearances, moral utterances (e.g., “it is wrong to murder”) do not express judgments or beliefs of the speaker but only noncognitive states (like disapproval). Simon Blackburn and other noncognitivists have developed tremendously clever devices for accommodating the apparently cognitivist surface features of moral discourse (e.g., Blackburn 1984, 1985). These noncognitivist pyrotechnics will not be addressed here. For the goal in this volume is to begin by taking commonsense moral thought at face value. On this approach we want to avoid, at least initially, invoking the kinds of subtle reinterpretations of moral discourse that are offered by noncognitivists. In addition, I am independently suspicious of some of the claims that undergird noncognitivism. For instance, noncognitivist accounts typically depend on the view that it is a conceptual truth that moral judgments are motivating, and this claim runs up against serious obstacles (see chapter 5, section 4).

Leslie 1994; Nichols and Stich 2003; Stich and Nichols 1992). This work on mindreading has been carried out largely in isolation from work in moral psychology, but the research on mindreading affords us the opportunity to evaluate perspective-taking accounts of moral judgment much more effectively.

Although the recent mindreading literature boasts no detailed perspective-taking account of core moral judgment, several recent theorists seem to suggest that perspective taking might be required for moral judgment. For instance, Alvin Goldman cites Schopenhauer and Rousseau as advocates of the view that empathy is “the source of moral principles” (1993, 355). And Goldman defines empathy in terms of perspective taking: “Paradigm cases of empathy . . . consist first of taking the perspective of another person, that is, imaginatively assuming one or more of the other person’s mental states. . . . The initial “pretend” states are then operated upon (automatically) by psychological processes, which generate further states that (in favorable cases) are similar to, or homologous to, the target person’s states. In central cases of empathy the output states are affective or emotional states” (1993, 351). In a similar vein, John Deigh claims that in order to grasp right and wrong in the deeper sense, one needs mature empathy, which involves *inter alia*, “taking this other person’s perspective and imagining the feelings of frustration or anger” (Deigh 1995, 758). Robert Gordon offers a more explicit and sophisticated perspective-taking account of how we determine whether an action is wrong, suggesting that we “imagine being in X’s situation, once with the further adjustments required to imagine being X in that X’s situation and once without these adjustments. If your response is the same in each case, approve X’s conduct; if not, disapprove” (Gordon 1995, 741).

Surely people sometimes use perspective taking in making moral evaluations. And the above authors are not sufficiently precise about which kinds of moral judgments depend on perspective taking to allow us to determine whether they would maintain that the basic capacity to draw the moral/conventional distinction depends on the capacity for perspective taking. There is certainly no systematic argument in the recent mindreading literature for the view that perspective taking is required for drawing the moral/conventional distinction. But the work on the moral/conventional distinction currently provides the clearest way to explore the basic capacity for moral judgment, so it will be of interest to see how a perspective taking account of this capacity fares against the evidence in any case. The evidence suggests that any attempt to defend that position will face some serious obstacles.

The first problem for perspective-taking accounts comes from developmental evidence. Children begin to appreciate features of the moral/

conventional distinction surprisingly early. Smetana and Braeges (1990) found that shortly before the third birthday, children were more likely to judge that moral violations generalized across contexts than conventional violations when they were asked, "At another school, is it OK (or not OK) to X?" (336). Further, according to Smetana and Braeges, after factoring in corrections for language, the results suggest that children generalize moral violations in this way shortly after the second birthday, and they recognize that conventional violations but not moral violations are contingent on authority at two years and ten months (Smetana and Braeges 1990, 342). More conservative estimates put the recognition that moral violations are not authority contingent shortly after the third birthday (Blair 1993). So, apparently young children can make these distinctions in controlled experimental settings. In addition, studies of children in their normal interactions suggest that from a young age, they respond differentially to moral violations and social violations (e.g., Dunn and Munn 1987; Smetana 1989).

The evidence of early success on the moral/conventional distinction sits alongside evidence of a somewhat later trajectory for perspective taking. Although at three years of age, children have some mindreading capacities, their perspective-taking abilities are still quite limited. Most famously, three-year-old children tend to fail the "false belief task." In the classic version of this task, Wimmer and Perner (1983) had children watch a puppet show in which a puppet, Maxi, put chocolate in a box and then went out to play. While Maxi was out, his puppet mother moved the chocolate to the cupboard. Children are asked where Maxi will look for the chocolate. Before the age of four, children tend to give the incorrect answer that Maxi will look in the cupboard, where the chocolate really is. This is another robust finding in developmental psychology. Dozens of studies have replicated the basic findings, and the typical result is that children do not pass the false belief task until after the fourth birthday (for a meta-analysis, see Wellman, Cross, and Watson 2001). As a result, there is reason to doubt that the young children who pass the moral/conventional task are proficient at determining the perspective of another person and then pretending to occupy that perspective.

The evidence on children poses an embarrassment for perspective-taking accounts, and the situation for these accounts is made worse by recent findings on children with autism. Children with autism have serious deficits to their capacity for understanding other minds (see e.g., Baron-Cohen 1995). Perhaps the best known finding is that autistic children fail the false belief task long after their mental age peers pass the task (Baron-Cohen, Leslie, and Frith 1985). They also have difficulty in other tasks that require taking the perspective of others. For instance, they

perform poorly when asked to determine which present is appropriate for another person (Dawson and Fernald 1987). In a series of studies on psychopathologies and moral judgment, R. James Blair presented autistic children with the standard moral/conventional task. Despite their deficiencies in perspective taking, Blair found that autistic children were able to make the moral/conventional distinction. Like normally developing children, autistic children treat moral transgressions as more serious and less authority contingent than conventional transgressions. As Blair (1993) points out, this finding poses a significant problem for perspective-taking accounts of moral judgment. For there is no doubt that autistic children have deficits in perspective taking and other sophisticated mindreading capacities. Hence, Blair's data on autistic children suggest that sophisticated mindreading abilities are not required to draw the moral/conventional distinction.

The data from development and psychopathology pose an obstacle for the perspective-taking account of core moral judgment. Arguments in this area are rarely decisive, and the above arguments against perspective-taking accounts are no exception. Nonetheless, the problems facing the perspective-taking account are certainly serious enough to warrant exploring alternative accounts of core moral judgment.

4. BLAIR'S VIM ACCOUNT

Armed with a dazzling series of experiments, Blair has developed the most detailed alternative to the perspective-taking account of moral judgment. Blair maintains that the capacity to draw the moral/conventional distinction derives from the activation of a Violence Inhibition Mechanism (VIM). The idea for VIM comes from Konrad Lorenz's (1966) suggestion that social animals like canines have evolved mechanisms to inhibit intraspecies aggression. When a conspecific displays submission cues, the attacker stops. Blair suggests that there's something analogous in our cognitive systems, the VIM, and that this mechanism is the basis for our capacity to distinguish moral from conventional violations.

Unfortunately, it is not entirely clear how VIM is supposed to produce the moral/conventional distinction, but we do get a broad outline from Blair (1995). It is useful to divide Blair's theory into two parts. The first part of the theory proposes that VIM generates a sense of aversion. VIM is activated by distress cues, but VIM-activation initially simply produces a withdrawal response. This VIM-activation becomes aversive through "meaning analysis": "the withdrawal response following the activation of