

Georgia State University

ScholarWorks @ Georgia State University

Philosophy Theses

Department of Philosophy

5-14-2021

Moral Attitudes: Why Attitude Matters in Characterizing The Psychology of Moral Judgements

Tate Anders Krogstad
Georgia State University

Follow this and additional works at: https://scholarworks.gsu.edu/philosophy_theses

Recommended Citation

Krogstad, Tate Anders, "Moral Attitudes: Why Attitude Matters in Characterizing The Psychology of Moral Judgements." Thesis, Georgia State University, 2021.
https://scholarworks.gsu.edu/philosophy_theses/294

This Thesis is brought to you for free and open access by the Department of Philosophy at ScholarWorks @ Georgia State University. It has been accepted for inclusion in Philosophy Theses by an authorized administrator of ScholarWorks @ Georgia State University. For more information, please contact scholarworks@gsu.edu.

MORAL ATTITUDES: WHY ATTITUDE MATTERS IN CHARACTERIZING THE
PSYCHOLOGY OF MORAL JUDGEMENTS

by

TATE KROGSTAD

Under the Direction of Neil Van Leeuwen, PhD

A Thesis Submitted in Partial Fulfillment of the Requirements for the Degree of

Master of Arts

in the College of Arts and Sciences

Georgia State University

2021

ABSTRACT

Stephen Stich has recently claimed that moral judgements are not a distinct psychological natural kind from conventional judgements. He argues that our two best methods in philosophy and psychology fail to provide a characterization of moral judgements that is immune to counter arguments, and therefore, it is likely there is only the psychology of conventional judgements. I argue that Stich's arguments implicitly assume a mistaken assumption, that moral judgements are psychologically distinguished from conventional judgements only if they concern different kinds of content. This assumption is mistaken because moral judgements can be distinguished from conventional judgements in virtue of attitude type. I provide a theory of a moral attitude type to demonstrate the plausibility that moral judgements can be psychologically distinguished from conventional judgements in such a way. Furthermore, a large history of scholarship in philosophy and psychology on moral properties and content might have understandably lead Stich's argument astray.

INDEX WORDS: Norm, Attitude, Domain, Morality, Judgement, Conventionality

Copyright by
Tate Anders Krogstad
2021

MORAL ATTITUDES: WHY ATTITUDE MATTERS IN CHARACTERIZING THE
PSYCHOLOGY OF MORAL JUDGEMENTS

by

TATE KROGSTAD

Committee Chair: Neil Van Leeuwen

Committee: Andrea Scarantino

Eric Wilson

Electronic Version Approved:

Office of Graduate Services

College of Arts and Sciences

Georgia State University

May 2021

DEDICATION

This thesis is dedicated to my mother, Carol Lynn Krogstad, and my father, Kenneth Byron Krogstad. To my mother for her love and support, and my father for leaving me with his philosophical curiosity.

ACKNOWLEDGEMENTS

I would like to thank my advisor, Neil Van Leeuwen, and my committee members, Eric Wilson, and Andrea Scarantino for all their help throughout this process. They have honed my philosophical thinking and writing and been incredibly patient.

TABLE OF CONTENTS

DEDICATION		IV
ACKNOWLEDGEMENTS		V
1 INTRODUCTION		1
2 STICH AND THE MORAL DOMAIN		3
2.1 The Philosopher’s Strategy		4
2.2 The Psychologist’s Strategy		5
2.3 Stich’s Conclusion		8
3 STICH’S IMPLICIT ASSUMPTION AND ATTITUDE TYPE		9
4 A MORAL ATTITUDE TYPE		12
4.1 Moral Attitude Architecture		13
4.2 The Four Functional Characteristics		14
<i>4.2.1 Non-reversibility</i>		<i>14</i>
<i>4.2.2 Universalization</i>		<i>16</i>
<i>4.2.3 Externalization</i>		<i>17</i>
<i>4.2.4 Characteristic Emotional Activation</i>		<i>18</i>
5 METHODOLOGICAL IRONY AND MORAL CONTENT		20
5.1 A Correct Definition of The Concept of Moral Judgement		22
5.2 Stich and Moral Intuitions		25
5.3 Stich and A History of Scholarship on Moral Properties and Content		27

6 CONCLUSION.....30

WORKS CITED.....31

1 INTRODUCTION

Are moral judgements psychologically distinguishable from conventional judgements? According to Stephen Stich, no. In his 2018 work “The Moral Domain” and elsewhere, Stich argues that our best methods in philosophy and psychology fail to accurately characterize the psychology of moral judgements, and so it is likely that future attempts will also fail (2018, 2019a/b). To show this, Stich relies on studies that demonstrate putative moral norms and putative conventional norms do not systematically differ in terms of content, that there is an extensive overlap in the domains these judgements range over. Stich then infers that moral judgements do not have a unique psychological basis; their psychology is indistinguishable from conventional judgements. I am convinced by Stich that moral and conventional judgements have extensively overlapping domains of content, but I do not believe, as Stich does, that this entails there is no psychological distinction between moral and conventional *judgements*. I argue that Stich’s argument relies on an implicit, mistaken assumption: moral judgements are psychologically distinguished from conventional judgements only if they concern different kinds of content. This assumption is false because these judgements can differ in terms of attitude type.

Attitudes, such as imagining that P and believing that P, are particular manners of processing information, and they are distinguished by their functional characteristics. Content does not differentiate imagination and belief; one can imagine that ice cream is at the pier or believe that ice cream is at the pier. Instead, the dimension of variation between imagination and belief is captured by their characteristic manners of processing information — attitude type. It is possible that moral judgements are psychologically distinguished from conventional judgements in virtue of a moral attitude type. An irony from Stich’s work is that the empirical studies he references tacitly presuppose a moral attitude type. In what follows, in order to make the case

that a moral attitude exists, I will offer a theory of a moral attitude through the characterization of its functional features. Furthermore, if my argument is successful, an exciting, possible implication can be drawn from Stich's work. His argumentative mistakes might be symptomatic of a general tendency in philosophy and psychology that needs to be corrected, namely, to explain moral phenomena in terms of content differences while ignoring attitude type.

Before going on, I would like to address why it is worthwhile for philosophers and psychologists to investigate whether moral judgements are psychologically distinct from conventional normative judgements. Such a project has general epistemic value for human knowledge, but furthermore, if there is a true morality (however characterized), it is an important question to what extent humans have a chance at cognizing it. If humans could not track moral distinctions in their minds, then sharp and troublesome questions arise about the relevance of morality to humans. If it were true that moral judgements are not psychologically distinct from conventional judgements, then human psychology of norms might drastically differ from true moral norms. Morality, then, would lose much of its behavioral and social importance for humans. The enterprise of moral philosophy has also long been conducting research under the presumption that it is going to discover some kinds of truths that are not conventional in nature. And if it were true that there was no general distinction between moral and conventional judgments, then moral academia would be a very striking anthropological phenomenon and occasion serious reflection on the discipline as a whole. Beyond these issues, if there is not a true morality that humans can or cannot track, then there remains a pertinent question of whether morality is a human, psychological phenomenon. And such a question relies squarely on whether moral judgements are psychologically distinct from conventional judgements. While more hangs

on this question, the few points brought up here help demonstrate why we should engage in this investigation.

2 STICH AND THE MORAL DOMAIN

Stich writes in “The Moral Domain” that philosophers and psychologists are engaged in a project to provide a correct definition of the concept of moral judgement (2018). Ideally, such a definition would specify the target of moral theories, i.e., the phenomena these theories attempt to explain, and thereby differentiate the moral domain from the conventional domain. But to achieve this, Stich claims, a theorist must correctly distinguish the set of all moral claims from nonmoral normative claims (548). Philosophers take a strategy called the *method of cases* where they use necessary and sufficient conditions to capture the set of moral claims, like how a bachelor is defined in terms of an unmarried man. Psychologists take a strategy called the *Kornblith-Devitt method* where they attempt to provide a nomological cluster of properties, like how gravity is defined in terms of a correlating increase or decrease in attraction with mass. According to Stich, these methods are the most plausible and dominant strategies for providing a definition of morality (554). As Stich argues, however, both methods fail to provide a characterization that is immune to counterarguments (549- 551, 553). He states, “Neither method has enabled us to give a nonstipulative empirically supported definition of the moral domain... The conclusion that I am inclined to draw at this point is that the quest is doomed to failure.” In other words, Stich asserts, “There *is* no correct definition of morality. *There is no moral domain*” (554, original emphasis). While Stich’s conclusion might be taken to entail that there are no

moral restrictions, Stich remains silent on such a further claim. Stich only explicitly claims that moral judgements are not a distinctive psychological kind.¹

2.1 The Philosopher's Strategy

The method of cases uses our intuitions to identify moral judgements and distinguish them from other kinds of judgements, collects those judgements into a set, and then uses necessary and sufficient conditions to classify that set as moral (549).² Stich brings two problems to the table for this method: 1) it relies on the implausible classical theory of concepts, and 2) a growing body of evidence suggests that philosophical intuitions vary across different demographic groups (550). The classical theory of concepts holds that concepts can be captured via a set of necessary and sufficient conditions. Recent studies, however, show that this theory does not apply to most concepts. So, it is unlikely it applies to the concept of moral judgements (550; Smith & Medin, 1981; Laurence & Margolis, 1999). While other theories of concepts are available, Stich believes the formidable issue stems from his second worry.³ The method assumes that the correct definition of the concept of moral judgment is implicitly stored in the mental information that “guides people’s intuitions of the application of the term *‘moral judgement’*” (Stich, 2018, 551, original emphasis). But a 2017 study by Levine and colleagues, Machery and Stich suggests that some religious groups have varying intuitions over their classification of judgements as moral or conventional (Levine, Rottman, David, Stich & Machery, 2017). Moreover, work by Emma Buchtel and colleagues in 2015 found that Eastern

¹ If Stich is correct, then what exactly it entails for ethics is a significant question, particularly in terms of moral obligation and permissibility. My goal in this paper is to defend the distinctiveness of moral judgements as an attitude type and enable people to ask coherently what Stich’s argument entails.

² It should be noted Frank Jackson in 1998 provided an influential reconceptualization of the method of cases that is seriously different than what I have described here (Jackson, 1998). Stich does not address Jackson’s theory and instead opts for this simplified version of the method of cases.

³ Since Stich ignores Jackson’s compelling theory it is possible that Stich’s first horn is a strawman argument.

and Western participants varied over which acts they were more likely to classify as immoral (Buchtel et al., 2015). While Stich recognizes that further work needs to be done, he explains that the preliminary results strongly indicate that moral intuitions are not uniform across demographic groups (Stich, 2018, 551). Japanese citizens, for example, often believe that being rude to a host at a dinner party is immoral, U.S. citizens, on the other hand, often consider this as conventionally wrong.⁴ This suggests, Stich thinks, that the correct definition of morality cannot be implicitly stored in the mental information that guides people's moral intuitions.

At best, Stich writes, the method of cases can only provide us with a single demographic group's conception of moral judgement (551). However, "Without some reason to think that one of these cultural variants succeeds in picking out what really are the essential properties of morality while the others miss the mark, it looks like this traditional philosophical approach to characterizing the moral domain should be abandoned" (552).

2.2 The Psychologist's Strategy

Hilary Kornblith and Michael Devitt prominently characterized scientific analysis as what is now called the Kornblith-Devitt method (1998; 1996).⁵ They highlighted it as a contrast point to traditional methods of philosophical inquiry, and Stich interprets psychologists as using it when examining moral judgements. The method uses scientific investigation to discover a set of characteristic cooccurring properties of a natural kind, like oxygen and silicon cooccurring in quartz. These "essential features," as Stich calls them, are features that are typically found in the

⁴ Buchtel et al.'s study which Stich cites here presented Eastern and Western participants with the same behaviors and found that they differed in whether they judged the behavior to be immoral or conventionally wrong. The study suggests that eastern immoral behaviors are strongly related to incivility, as opposed to western immoral behaviors which are strongly related to serious harm. Stich's study with Levine and colleagues replicates the basic finding of Buchtel et al. that demographic groups vary in whether they consider the same behaviors to be moral or conventional. Although these studies are promising, further research should be done to strengthen this empirical claim. For sake of argument, I will assume Stich is right, and show that it still does not support Stich's wider claims.

⁵ The method was influenced by Hilary Putnam, who argued that it is the role of empirical science to establish the essential features of a natural kind, not analysis through and of ordinary concepts (1975).

natural kind and the conjunction of them are used to constitute a correct definition of the kind in question (Stich, 2018, 551). Concerning the project at hand, Stich interprets psychologists as taking moral judgements to be a psychological natural kind. “If that’s right,” he explains, “then the conjunction of [the features they discover] constitute an empirically supported definition of the moral domain” (552). First, Stich elaborates, the method gathers intuitive moral judgements into a representative list and then uses scientific investigation to determine a nomological cluster of properties, i.e., a set of properties that almost always occur together. The nomological cluster of properties are considered the essential features of the kind and thereby constitute a definition of it. As a consequence, any intuitively classified moral judgement that lacks these features is not considered part of the kind. Although intuitions are used to compile a set of paradigmatic moral judgements, those intuitions are not used to determine the essential features. Hence, Stich writes, intuitions play no significant role in determining the correct definition of morality and so the Kornblith-Devitt method does not fall prey to the criticisms of the method of cases. “When things work well,” Stich explains, “the method enables us to offer an empirically supported account of the essential features of a kind” (552).⁶

Stich interprets a prominent psychologist, Elliot Turiel, as using the Kornblith-Devitt method to establish the essential features of the moral domain and the correct definition of morality (552; Turiel, 1978, 1983; Turiel, Killen, & Helwig, 1987). Turiel argues that there are four features of our moral judgements that form a nomological cluster of properties: 1) reference to harm, 2) authority independence, 3) geographical universalizability and 4) temporal

⁶ A bellyache from Stich’s exposition here is that it is unclear how the essential features (the nomological cluster of properties) differ from being necessary and sufficient conditions. Stich does not spell out this difference and instead assumes that they are in some way relevantly distinguishable. Unfortunately, Stich’s silence complicates our understanding of how this method is sufficiently different from the method of cases, save using scientific investigation.

universalizability. According to Turiel, moral judgements are typically (1) justified by appeal to harm, (2) not influenced to change based on an authority figure's judgement, (3) applied to all persons across geographical location and (4) applied to all persons across temporal position (such as when wrongness of a behavior is judged to apply to a person's actions in another time and place).⁷ While subsequent research initially seemed to support Turiel's findings, there is now mounting evidence that his nomological cluster falls apart, that these features are not present in several instances of moral judgements (Stich, 2018, 552).

Stich presents work that shows these features are not found in all moral judgements. A 1993 study by Jonathan Haidt and colleagues reported that most of their participants judged, without appealing to harm, that having sex with a dead chicken was immoral (Haidt, Koller, & Dias, 1993). Earlier in 1987, Mordechai Nisan demonstrated that Israeli children also made moral judgements of wrongness without appeal to harm. However, their judgements did exhibit authority independence and geographical and temporal universalizability (Nisan, 1987; Stich, 2018, 552). In 2007 Kelly, Stich, Haley, Eng, and Fessler investigated moral judgements concerning serious cases of harm such as slavery and the punishment of whipping. Stich writes, "They found that, for many participants, judgements about these clearly harmful cases were *not* authority independent and *did not* generalize to other times and places" (553, original emphasis; Kelly et al., 2007). Most of these studies used participants from large-scale societies that comprise or closely resemble WEIRD people.⁸ A further study in 2015 by Fessler et al., though, shows that Turiel's nomological cluster also falls apart in small scale societies (Fessler et al.

⁷ It is important to note that the last three features can be viewed as attitudinal, rather than as part of the content of moral judgements.

⁸ WEIRD is an acronym applied to people from, or those who resemble people from, western, educated, industrialized, rich, and democratic societies. The acronym comes from a 2010 study by Henrich, Heine, and Norenzayan that showed most participants in psychological studies come from this kind of society that make up only a minority of the human population (2010).

2015). Stich explains that these studies demonstrate that moral judgements do not almost always involve harm, authority independence, and geographical and temporal universalizability at the same time. Hence, Stich concludes, Turiel's theory does not amount to a nomological cluster of properties of moral judgements and therefore cannot be used to determine the essential features of the moral domain or the correct definition of morality (2018, 552-553).

2.3 Stich's Conclusion

Stich believes that since our best attempts to specify the moral domain and establish the correct definition of the concept of moral judgment fail, then future attempts will also likely fail. What does this mean for morality? Stich argues that there is nothing on the table right now which provides us with good reason to believe that moral judgements are distinct from merely conventional judgements. So, he writes, "... if I am right about both the philosophers' conceptual analytic project and the psychological project spearheaded by Turiel, then there is no reason to think there is a special kind of normative judgement of the sort that philosophers and psychologists imagine" (2019/b, 6). Rather, Stich continues, "...the conviction that there must be a natural or well-motivated way of dividing normative judgements into those that are *moral* and those that are *nonmoral* is, I think, an illusion..." (2018, 554, original emphasis). According to Stich, we have a false belief that normative judgements are parsed into moral and conventional categories. This belief is false because there is no substantive psychological distinction that enables us to distinguish these judgements. Stich does not deny that "normative judgements about purity, reciprocity, authority, and kinship [etc.]" exist, rather, he denies that the psychology of these putative moral judgements is legitimately different and isolable from conventional judgements (554). Stich, here, is drawing an inference to the best explanation and believes that the most parsimonious explanation of the evidence he provides is that morality and

moral judgements are not a distinguishable normative category from conventionality. In other words, an undifferentiated normative judgement type, conventionality, would best explain and predict the relevant data, and so we should not extraneously posit that moral judgements are psychologically distinct. Furthermore, while it is possible for morality to exist independently of our psychology, Stich implies a skeptical view that if there are only conventional normative judgments, then morality and the moral domain must not exist.

3 STICH'S IMPLICIT ASSUMPTION AND ATTITUDE TYPE

I am convinced by Stich that there is no nonstipulative characterization of the moral domain, but I do not think his arguments achieve what they ultimately aim at. Stich's evidence largely refutes the two strategies by demonstrating that the *content* of moral judgements does not systematically differ from conventional judgements. But in order for Stich to infer from this that moral judgements are not psychologically distinct from conventional judgements he needs a further assumption. He needs to assume that moral judgements are psychologically distinguished from conventional judgements only if they concern different kinds of content. Without this assumption his inferences do not follow.

To clarify, recall the works of Emma Buchtel, Johnathan Haidt, Sydney Levine, and Kelly et al. Stich uses Buchtel et al.'s study and Levine and colleagues' study to show that intuitions over moral and conventional judgements vary and that a single norm can be the content of a moral or conventional judgment. Stich uses Haidt et al.'s study to show that one of Turiel's features, harm, is not found in all moral judgements. Moreover, Stich uses his own study with Kelly et al. to show that not all moral judgements involve a norm that is authority independent and generalized across time and space. Stich interprets Turiel's work as a

nomological cluster of *content* properties of moral judgements and argues that they are not almost always found in all moral judgements. Stich references these studies and others to refute the characterizations of the moral domain by philosophers and psychologists, but these studies do so by illustrating that there is an extensive overlap in the domains moral and conventional judgements range over. So, if Stich is to infer from these studies that moral and conventional judgements are not psychologically distinct, then his argument must implicitly rely on an assumption that moral judgements are psychologically distinguished from conventional judgements only if they concern different kinds of content. Given this assumption, Stich could reasonably infer that moral judgements are not psychologically distinguished from conventional judgements. However, without this assumption Stich's evidence only shows that moral and conventional judgements do not concern different kinds of content and that philosopher's and psychologist's content-based characterizations are incorrect.

However, Stich's argument is not warranted in its implicit assumption because it is possible that *attitude type* psychologically distinguishes moral and conventional judgements. Attitudes like imagination and belief can concern the same content, such as *that there is a cat on the sofa*, and remain psychologically distinct in virtue of processing that content differently. Hence, even if moral and conventional judgements do not range over different kinds of content, they can remain psychologically distinct through their differing manners of processing information, i.e., their functional characteristics. To sharpen this point, an easy analogy can be drawn with the attitudes of belief and desire.

Judgements of belief and judgements of desire can range over the same content, such as ice cream, but remain psychologically distinct. Consider,

- a) Adam *believes* ice cream is in front of him.
- b) Erica *wants* ice cream to be in front of her.
- c) Greg *believes* ice cream is sold at the pier.
- d) Tracy *wants* ice cream to be sold at the pier.

The content of these mental states does not determine the kind of mental state they are or the substantive psychological distinction between them. It is the particular attitude type (belief vs. desire) and corresponding functional characteristics that does this work. Since moral and conventional judgments *are* mental states, it is possible that attitude type does this work for them as well.

As Stich has shown, a single norm can be the content of a moral or conventional judgement. Consider the norms of spanking one's children and eating meat,

- e) Adam *morally judges* that spanking one's children is wrong.
- f) Erica *conventionally judges* that spanking one's children is wrong.
- g) Greg *morally judges* that eating meat is wrong.
- h) Tracy *conventionally judges* that eating meat is wrong.

Just like belief and desire, it could be attitude type that distinguishes these judgements. Stich's evidence can be interpreted as supporting this possibility given that we intuitively believe moral and conventional judgements are distinct. Since Stich's evidence shows that the content of moral and conventional judgements does not systematically differ, it is likely that the content of these judgements does not determine the kind of judgement being made or the substantive psychological distinction between them. Instead, this work might be done by attitude type just like mental states of belief and desire. Since a moral attitude type has not been ruled out, it is false that moral judgements are psychologically distinguished from conventional judgements *only if* they concern different kinds of content. Hence, Stich's argument has not established yet that there is no psychological distinction between moral and conventional judgements because his inference relies on a mistaken assumption. In fact, in an ironic twist, the methodology of

Stich's experiments tacitly presupposes that there is indeed a moral judgement attitude type. I will return to this point in section 5, but first let me characterize what I think that attitude is.

4 A MORAL ATTITUDE TYPE

The following notions, which Stich and I share, are presupposed throughout this discussion: a norm, rule, intrinsic motivation of norms, permissibility, impermissibly and obligation. Stich explains that, "a norm is a rule or principle that specifies actions which are required, permissible or forbidden independently of any legal or social institution" (Sripada & Stich 2005, 2). Moreover, Stich writes, "people are motivated to comply with norms as *ultimate ends*, rather than as means to other ends," such that people "exhibit a characteristic style of motivation in which the individual intrinsically values compliance with... rules even when there is no possibility of sanction from an external source" (8, original emphasis). In other words, the notion of a norm involves rules that specify permissible, impermissible or obligatory behaviors (or thoughts) that motivate people to act according to them independent of any legal or instrumental reason. Stich and I take it that norms in this sense are found universally in human cultures. Stich, however, takes it that moral norms are distinguished exclusively by their content. But I believe that moral judgements, as a type of judgement *about* norms, are distinguished by the way the mind functionally processes them, in particular through a moral attitude.

For the rest of this section, I will characterize and argue for four typically cooccurring functional features of moral judgments. While the specific details of my theory are my own, Adam Smith, Elliot Turiel, Jesse Prinz, and Kyle Stanford deserve credit as the forerunners of the features I present. In particular I interpret Elliot Turiel's work as positing attitudinal features of moral judgements. I also build on Adam Smith and Jesse Prinz's work that moral emotions

function differently than nonmoral emotions. Lastly, with the help of Kyle Stanford's compelling insights, I attempt to capture and analyze the objective phenomenology of moral judgements.⁹

4.1 Moral Attitude Architecture

My hypothesis is that if a moral attitude is a psychological natural kind, then it is functionally distinguishable from a conventional attitude in virtue of the following four features:

- i. Non-reversibility
- ii. Universalization
- iii. Externalization
- iv. Characteristic Emotional Activation

These four features can be tested for by investigating whether and to what degree an agent's mind characteristically processes a norm in these ways. My prediction is that when an agent makes a moral judgement, these four functional features will almost always be present. The conjunction of these functional features constitutes the moral attitude type as a natural kind, but it is not necessary for these four features to *always* be present in moral judgements. Further studies should be conducted to test for outlier cases and assess factors that contribute to their occurrence. But these cases will be unique, and I claim that similar to other natural kinds, we will find that these four features will be typical cooccurring properties of moral judgements. So, when an agent moralizes a norm *do not eat meat*, we can test whether and to what degree their mind processes that norm as (i) being insensitive to change, (ii) applying to all persons across time and space, (iii) objective, such that the agent experiences a distinctive phenomenology of objectivity of that norm and (iv) involving a pattern of emotional responses caused by a perception of violation or conformity to the rule. Essentially, our moral attitude is a propositional attitude

⁹ Prinz, 2007; Smith, A., 1759/1976; Stanford, 2018; Turiel, 1978, 1983, 1987.

concerning a rule, and the rules can vary in complexity such as *it is wrong to kill to eat meat* or *it is wrong to eat pork on religious holidays*. But when an agent moralizes a rule, acceptance of the rule only constitutes a moral judgement as a mental state if the agent typically processes the rule via these four features. Furthermore, because our moral attitude is just a distinctive manner of processing information, an agent can moralize any rule or norm. An important qualification of each functional feature, though, is that there is an open-ended metric for overriding factors that allows for degree and change. Testing for this open-ended metric requires that we identify it first, but that is a pursuit for further research. The open-ended metric is necessary because it will account for further psychological complexities of our moral judgements, particularly moral excuse.

4.2 The Four Functional Characteristics

4.2.1 Non-reversibility

The first feature is non-reversibility, and it is when an agent processes a norm as highly resistant to change. The core of this feature is motivated by Elliot Turiel's notion of authority independence. Turiel found that children judged a behavior of hitting another student as impermissible regardless of whether or not a teacher or other authority figure determined it was OK (Turiel, 1978, 1983; Turiel, Killen, & Helwig, 1987). The notion of authority independence is that moral judgements, as opposed to conventional judgements, are insensitive to change based on an authority figure's decree. While subsequent research by Kelly et al. showed that not all moral judgements exhibit this feature, such that some moral judgements were sensitive to an authority figure's decree, the notion can be broadened and amended to capture Kelly's findings (Kelly et al., 2007).

Non-reversibility takes the essential function of authority independence, broadens the influencer condition and assumes a metric for overriding factors to allow for degree. An influencer condition is anything that attempts to challenge the rule, such as an authority figure, circumstance, emotional appeal, rational argumentation, etc. The mental resistance or insensitivity to the influencer condition will always be strong such that an agent will be reluctant to admit of exceptions. However, depending on the influencer condition(s) and the rule, an agent could make an exception. If someone moralizes the rule *do not eat meat*, but finds themselves stranded on a desert island, they might admit of an exception to eat the land critters and sea life to stay alive. For the agent, the influencer condition has enough weight against the rule to admit of an exception, and so the agent morally excuses eating the land critters and sea life. Kelly et al.'s work can be interpreted as showing that for some people, authority decrees as an influencer condition is weighted enough to admit of exceptions. The weight of the influencer condition(s) is determined by an open-ended metric. The metric is an essential feature of non-reversibility because it allows for a degree of resistance and insensitivity that partly explains moral excuse.¹⁰ While conventional norms can be insensitive to change, agent's rarely process them in such a way; for example, an American who conventionalizes *men should wear pants as formal dress* might suddenly come to believe that this does not apply to men in Scotland.

¹⁰ It should be noted that moral development can also be captured by reference to an agent's open-ended metric. When a white nationalist drops their racist views, it is usually over a long period of time. A culmination of several influencer conditions weighing against a white nationalist's moralized racist rules could lead to a systematic change in the way that agent processes those rules. In the case of non-reversibility, this would mean that the agent ceased to process those rules as strongly insensitive to change. Similarly, when an agent ceases to moralize a rule all together, that agent no longer processes the rule via these four functional features. How and whether or not moral development happens is dependent on an agent's open-ended metric that allows for degree in these features.

4.2.2 *Universalization*

The second feature is universalization, and it is when an agent extends the applicability of a rule to all persons across time and space.¹¹ Several classic moral norms almost always exemplify this function as a feature of the attitude taken toward them: *do not murder*, *do not rape*, etc. When an agent universalizes the norm *do not murder*, the agent applies the obligation to adhere to it to everyone regardless of geographical location and temporal position. The agent then perceives any violation of the rule as impermissible. Conventional norms do not exhibit this function. When an agent conventionalizes a norm, the applicability of the rule is only extended to a certain demographic group(s). Universalization is motivated by Turiel's feature of universalizability, but it is also amended to involve an open-ended metric to allow for degree. This is because universalized moral judgements sometimes admit of narrow exceptions for localized groups.

To clarify, consider our lone survivor case again. The lone survivor morally excuses eating land critters and sea life to stay alive but, let us stipulate now, also universalizes the wrongness of eating meat. This narrow exception can be explained by the complexity of the norm in question, namely, *it is wrong to eat meat, unless it is necessary for survival*. This norm is more complex than *it is wrong to eat meat*, but anyone can still have a universalizing attitude towards this norm. The lone survivor adopts a universalizing moral attitude towards this complex norm given their circumstances, and then admits of a legitimate exception to eat meat. The lone survivor's open-ended metric helps capture why the lone survivor adopts the complex norm and

¹¹ Some universalized rules can have highly specific antecedents that restrict the universalization to only certain groups. Honor based systems are an example of having a universalizing attitude to a highly complex rule, such as, for all X, if X is of high status then it is shameful to do Y; in such cases, the group that the universalization applies to is relatively small. Universalization can happen to a greater or lesser extent depending on the complexity of the rule.

makes an exception. The open-ended metric could determine a threshold for when an agent includes a complex norm to be processed via their moral attitude, and hence, universalizes that new norm. A strict vegan, for example, might not consider survival, the added condition of the complex norm in question, to take precedent over their moralized general norm *it is wrong to eat meat*. The open-ended metric offers explanatory power to differentiate our lone survivor and strict vegan by determining unique psychological thresholds for moral processing inclusion.¹²

4.2.3 Externalization

The third feature is externalization, and it is when an agent ascribes objectivity to a norm and experiences that norm as objective. Kyle Stanford argues that moral norms are experienced as objective in a way that conventional norms are not. He explains, “moral norms exhibit a [strong] form of phenomenological externalization or objectification, one revealed most clearly in the unwillingness of subjects to tolerate disagreement without error concerning those norms and/or their implications” (2018, R5, para. 1). We can take from this that when an agent externalizes a norm, they process that norm as requiring a higher standard of evidence to be disregarded; that standard indicates the kind of objectivity the agent experiences and ascribes to moralized norms. Stanford notes that the standard is more modest than that of empirical fact, which means agents ascribe and experience a kind of objectivity that is also more modest. To clarify, Stanford writes that the standard is unlike that of “judgements of taste and preference,” where no objectivity is ascribed or experienced. Rather, he continues, the degree of objectivity

¹² Stich cited Kelly et al.’s study as direct evidence against Turiel’s feature of universalizability (Stich, 2018, 553; Kelly et al., 2007). The study found that some of their participants did not universalize the wrongness of whipping slaves to American slave owners 300 years ago and Greco-Roman societies. These participants judged that the norm was OK when considering various temporal and geographical conditions. But Kelly et al.’s study just shows that not *all* participants had the attitude of moral judgement that whipping slaves was wrong; some of the participants *morally judged* the norm to be wrong, and (shockingly) other participants only *conventionally judged* it to be wrong.

lays somewhere in between that of empirical fact and conventional norms (R5, para. 2). The point here is that when an agent moralizes a norm, they ascribe a certain level of objectivity to that norm that involves a corresponding phenomenological experience similar that of empirical fact. Stanford states that there is “reliable variation among human beings in whether they externalize or objectify even prototypically moralized judgments,” but that most “recognize such selective externalization of some but not all norms as part of their own moral phenomenology” (R2, last para.). We can interpret this as implying that the ascription of objectivity and resulting phenomenology comes in degrees depending on a particular person’s psychological metric.

4.2.4 Characteristic Emotional Activation

The fourth and last feature is characteristic emotional activation, and it is when an agent perceives a violation or conformity to a rule and reacts with a certain pattern of emotional responses. Call the pattern of emotional responses *moral emotions*; they are general emotions directed at social behaviors that are relevant to the activation of at least one of the three preceding attitudinal features.¹³ The core idea is motivated by Jesse Prinz. He writes, “I will say that moral emotions are emotions that arise in the context of morally relevant conduct. More specifically, moral emotions promote or detect conduct that violates or conforms to a moral rule” (2007, 68). In order to avoid circularity in the notion of “morally relevant conduct,” I will say that morally relevant conduct is specified by a rule that is typically processed in a given agent’s mind via the three preceding attitudinal features. Following Prinz, I claim there might be several moral emotions that we do not yet have words for, but commonly known ones are anger, contempt, disgust, guilt, shame, admiration, gratitude, gratification, and dignity (79, 81). While it

¹³ I am referring here to a cluster of emotional states that are often identified in the literature as cooccurring with moral judgments. Further work must be done to avoid circularity but developing a concept of moral judgement involves developing certain reactive emotional patterns. It is important to note that these emotions should not be defined *by* morality, but rather characterized as moral in virtue of their cooccurrence with moral attitude processing.

is possible to experience each of these emotions in a nonmoral context, they are functionally different in a moral context. Hence, there is *anger* and *moral anger*, *disgust* and *moral disgust*, etc.

Consider moral disgust; it is a reoccurring emotional response to a perception of a violation of a moralized rule, such as *do not eat meat*. When someone moralizes this rule and perceives someone eating meat, they do not respond with simple disgust like that towards rotten vegetables. Rather, they characteristically respond with an emotion, disgust (i.e., moral disgust) directed at the social behavior of the perpetrating agent. Moral disgust plays a different function here than the disgust towards rotten food. It suffices to say that there are many different functions of our moral emotions that differentiate them from their nonmoral counterparts. The point at stake is that the moral emotions are those that characteristically and systematically arise with moral judgements; these emotions do not arise with conventional judgements. Hence, when someone moralizes the rule *do not eat meat* and perceives a person eating meat, they will characteristically respond with an emotion directed at a particular social behavior, such as moral disgust, whereas someone who conventionalizes this rule will not.

Adam Smith's conception of sympathy can partly explain the evaluative complexity of moral emotions and moral judgements. Sympathy, according to Smith, "...denote[s] our fellow-feeling with any passion whatever" (1759, 1976, I.i.1.5). Smith claims that sympathy is a psychological mechanism that allows us to relate to another person's mental states. The information that is taken up by this mechanism can be sensory information of another person's behavior, expressions, or circumstances. It can also be the product of inferences about this information or the imaginations of being in the other person's situation. To apply this notion, let us consider moral disgust again. Smith writes, "The true cause of the peculiar disgust which we

conceive for the appetites of the body when we see them in other men, is that we cannot enter into them” (1759, 1976, I.ii.1.3). When an agent moralizes the rule *do not eat meat* and responds with moral disgust, they 1) perceive a violation of the rule, and 2) are unable to relate to the mental states of the perpetrating agent. Sympathy allows for thorough evaluation of a violation or conformity to a rule. The agent evaluates the violation in such a way as to capture circumstantial complexity that thereby enables the agent to respond with a fitting moral judgement and moral emotion. The particular psychology of the agent will determine which moral emotion is elicited. Hence, not everyone will respond with moral disgust at the sight of someone eating meat, they might respond with another moral emotion, such as moral anger.¹⁴

5 METHODOLOGICAL IRONY AND MORAL CONTENT

I noted in section 3 that Stich’s empirical work tacitly presupposes a moral attitude, let me now return to this point. Recall Stich’s study with Levine and colleagues; they investigated intuitive classifications of moral judgements among different religious groups. Their experiment presented participants with the statements, “Many people in Honduras believe that people should not kill others for no reason,” and “Many people in Italy believe that adults should not eat pasta with their fingers,” etc. They then asked their participants whether they and the people in their community agreed. Following that question they asked, “Now consider the judgement you just made. Is *that* a moral judgement or some other kind of judgement?” (Stich, 2018, 550, my emphasis; Levine, Rottman, David, Stich & Machery, 2017). Their experiment demonstrated that

¹⁴ A brief and possible counter argument that Stich might make is that a unique moral attitude type does not exist; rather, moral judgements could be the product of a hodgepodge of different attitudes like belief and imagination. Externalization could be a product of imagination ascribing objectivity to a rule, and moral norms being held to a higher standard of evidence could be a result of belief. Stich might be right that a moral attitude type does not exist. However, this rebuttal means he has granted the central point of my argument, that there is a substantive psychological distinction between moral and conventional judgements, and it is found in attitude type.

different religious groups differed in whether their judgement was moral. But here is the methodological irony: their experiment only makes sense if their participants have an intelligent grasp of different kinds of judgements and their corresponding attitude types. The last question asks participants to consider whether their judgement was moral, conventional, *or* some other kind of judgement, such as hypothetical, expectational, imaginative, desirous, etc. *The question tacitly presupposes that moral and conventional judgements are distinguished by attitude type in the same way that other kinds of judgements are, and moreover, that participants comprehend the difference.* Hence, for participants to accurately address the question, they must implicitly understand that attitude type distinguishes moral and conventional judgments. It follows that Stich's very argument, through its reliance on this empirical work presupposes the very thing it would deny: a moral attitude. So, how did Stich's argument miss its mark?

There are a few salient points to draw here about Stich's argument. A first relates to Stich's methodological irony; Stich's argument overlooks inherent ambiguity in the notion of a correct definition of the concept of moral judgement, and this seems to contribute to its mistaken assumption that moral judgements are distinguished from conventional judgements in virtue of concerning a certain kind of content. Stich's argument also makes a further general mistake in its interpretation of the empirical evidence concerning moral intuitions that broadly undermines its criticism against the philosopher's project. Stich's argument mistakenly infers that if moral intuitions differ across demographic groups, then there must not be a similar underlying concept of moral judgement across these groups. These mistakes seem to be symptomatic of a history of scholarship in philosophy and psychology that focuses on moral properties and the content of moral judgements, which suggests how Stich's argument could have been influenced and miss its mark.

5.1 A Correct Definition of The Concept of Moral Judgement

To recap, Stich explains early in his argument that the philosopher's project and the psychologist's project are aimed at constructing a *correct* definition of the concept of moral judgement (2018, 549). To this point, Stich references Alasdair MacIntyre and his influence on philosophy's movement towards analyzing moral utterances. Stich writes that, of MacIntyre and the philosophers engaged in analyzing moral utterances, they "believed that the concept of moral utterance or moral judgement specifies some (or perhaps all) of the essential properties of morality and, thus, that a definition that made this concept explicit would tell us what the essential features of morality are" (549). Stich follows this thought throughout his argument. If philosophers and psychologists are able to accurately characterize and capture the concept of moral utterances (judgements), then the definition they develop should incorporate the relevant information to distinguish moral judgements from all other nonmoral normative judgements. In other words, it should be able to distinguish the domain of moral claims from the domain of nonmoral normative claims. Any *stipulative* definition, to reiterate, such as a "Mormon concept of moral judgement" or a "secular liberal concept of moral judgement" will not do because there is no independent reason, at least right now, "to think that one of these cultural variants succeeds in picking out what really are the essential properties of morality while the others miss the mark" (551). Stich is right to emphasize that the philosopher's project and the psychologist's project should aim at constructing a *correct* definition of the concept of moral judgement, but unfortunately, the notion of a correct definition of the concept of moral judgement that Stich's argument works with here is inherently ambiguous, and this partly leads his argument astray.

Stich's argument does not specify *what* should be involved in a correct definition of the concept of moral judgement, and indeed, part of what the philosopher's and psychologist's

projects are doing is to determine what *is* involved. However, because the notion is unclear, a correct definition of the concept of moral judgement could involve any number of things, such as specifying a certain kind of content, attitude type, or both, etc. This ambiguity makes it difficult for Stich to have airtight interpretations of his evidence against the philosopher's and psychologist's projects, and I believe this contributes to Stich's argument assuming an implicit mistaken assumption about moral judgements and content. For example, as noted above, Stich's argument uses his study with Levine and colleagues and another by Emma Buchtel et al. to argue against the philosopher's method of cases, in particular its central idea that the correct definition of the concept of moral judgement is implicitly stored in the "information that guides people's intuitions about the application of the term "moral judgement" (551). The studies show that different demographic groups have varying intuitions over whether certain instances of their judgements are moral. Stich's argument attempts to infer from this that there must not be any unifying information across demographic groups that *does* or *could* constitute a correct definition of the concept of moral judgement. However, since the correct definition of the concept of moral judgement is ambiguous, it seems that Stich's argument cannot convincingly make this inference without assuming what kind of property or properties are involved in the correct definition. We can illustrate this criticism more concretely by taking closer look at Stich's study with Levine and colleagues.

In the study, the questions presented to participants involved widely different forms of content, i.e., killing people for no reason and eating pasta with fingers (550). If we assume that the correct definition of the concept of moral judgement involves specifying a certain kind of content, as Stich's argument seems to implicitly do, and people's intuitions actually vary over whether their judgements concerning these kinds of contents are moral, then Levine and

colleagues' study would provide good evidence against the philosopher's project. Specifically, it would show that the method of cases' central assumption is incorrect because there would be no relevantly similar or unifying information concerning *content* implicit in different demographic group's concepts of moral judgement that guides their intuitive moral judgments. But this inference depends on assuming what the correct definition of the concept of moral judgement involves, in particular, by specifying a certain kind of content. It is not unusual for philosophers and psychologists to get fixated on moral properties and moral content. G. E. Moore, for example, argued that ethical knowledge "rest[s] on a capacity for an intuitive grasp of fundamental ethical truths..." of which those fundamental ethical truths have value in virtue of an irreducible non-natural (moral) property that is the result of a specific intrinsic nature (Baldwin, 2010, Section 3, para 4; Moore, 1899, 1922, 1903/1993).¹⁵ The inherent ambiguity of the notion of the correct definition of the concept of moral judgment seems to have led Stich's argumentative interpretation of his evidence astray and implicitly rely on an assumption that content was the defining factor of moral judgements. If we assume the correct definition involves something different, though, like an attitude type, we can interpret and contrast Levine and colleagues' study in another way.

Different cultures and demographic groups could still have the same concept of moral judgement, i.e., a concept dependent on a moral attitude type, and simply apply it to different kinds of content. Their intuitions differ over which judgements are moral judgments not because they have different concepts of moral judgement, but rather because their intuitions are varying over which kind of content their moral judgements apply to. Hence, under this interpretation,

¹⁵ I'll argue in section 5.3 that such a history in scholarship in philosophy and psychology possibly influenced Stich's overall argument.

Levine and colleagues' study does not provide evidence against the philosopher's project. The study would not show that there is not a consistent set of information about the correct definition of the concept of moral judgement implicit in the information that guides different demographic groups' intuition about moral judgements. This point is important because it illustrates that Stich's argument makes a more general mistake concerning moral intuitions that further undermines his criticism against the method of cases.

5.2 Stich and Moral Intuitions

While Stich acknowledges that more studies beyond Levine and colleagues' and Buchtel et al.'s need to be conducted to confirm that philosophically important intuitions like moral intuitions *do vary* across different demographic groups, he assumes for present purposes that the research is on the right track (Stich, 2018, 551). Stich's argument takes this assumption to largely refute the method of cases. Stich writes,

Let's explore the implications of this assumption for the project of providing a definition of morality... The goal of the project of defining morality is to distinguish moral claims from nonmoral normative claims and to do it correctly, not stipulatively. But if our assumption is correct, it looks as though the best that the method of cases can give us is a characterization of the secular liberal concept of moral judgments, the orthodox Jewish concept of moral judgement... the Beijing Chinese concept of moral judgement, and so forth. (551).

Stich infers that, without independent reason to believe one of these variants picks out the correct definition of the concept of moral judgement, the method of cases should be considered inadequate to yield its target definition (551). In other words, Stich believes, provided his argumentative assumption that moral intuitions do vary across demographic groups, there must not be an underlying correct concept of moral judgement across these groups. And so, Stich's argument concludes, the method of cases "should be abandoned" (551). But Stich's primary inference against the method of cases is a non sequitur. Stich's inference does not follow because

he mistakenly believes that *if intuitive judgements of a certain kind differ, there must not be an underlying similar concept of those judgements.*

People can have the same underlying concept of a certain kind of judgement even though their specific intuitive application of those judgements differ. For example, take judgements of direction. Suppose I have lost a cat with a foreign exchange student from China, Cheng; Cheng intuitively judges that the cat went north, and I intuitively judge that the cat went south. Our intuitive application of directional judgement differs, notably here in terms of content (north vs south), but they arguably have and are guided by the same concept of directional judgement. We can push this analogy forwards towards philosophically important intuitions like moral judgements. In Buchtel et al.'s study, for example, "seemingly mild misbehaviors such as spitting, cursing, and littering were much more likely to be called immoral by Beijing than Western participants, while serious behaviors such as killing, stealing, and hurting others were much more likely to be called immoral by Western participants" (Buchtel et al., 2015). Suppose that Cheng intuitively judges that *cursing* is immoral, and I intuitively judge that *cursing* is not immoral. Just like our lost cat example, Cheng and I could still have the same underlying concept of moral judgement that guides our application but differ in terms of which content we intuitively apply that judgement to. In other words, Beijing and Western participants can vary in their intuitive application of moral judgement, but still have the same underlying concept of moral judgement that guides their application.¹⁶ This could also hold regardless of whether or not the correct concept of moral judgement involves a unique moral attitude type. Hence, Stich's criticism against the method of cases is more generally undermined by an unacknowledged

¹⁶ There could be several reasons why intuitive application would differ while the concept stays the same. For example, environmental factors could force a certain culture to adapt and acquire new moral norms in order to survive while a different culture was insulated from those factors and had no need to acquire such norms. Intuitions could simply be picking up on different cultural norms.

feature of intuitions. This means that the method of cases is still a possible avenue to discover the correct definition of the concept of moral judgement; whether it will remain an open question.

5.3 Stich and A History of Scholarship on Moral Properties and Content

Stich's argument makes a few serious mistakes that characteristically relate to content. I believe that a large history of scholarship in philosophy and psychology that focuses on moral properties and moral content makes these implicit mistakes more understandable. There is no way to tell if Stich was explicitly influenced by this history of scholarship, but the legacy and massive academic environment involved in this history suggests that Stich's argument might have been symptomatic of it. Though there are several philosophers and psychologists that focus on moral properties and moral content, I will note just a few to illustrate this history.

G. E. Moore popularized contemporary metaethics by arguing for moral realism in the 19th and 20th centuries. He argued that moral properties are nonnatural through his open-ended argument and held that ethical knowledge depended on an intuitive grasp of these moral properties. G. E. Moore's arguments sparked swaths of debate, but his central belief that moral value depends on a certain kind of nonnatural moral property has been argued for by many successors. For example, during and after Moore's lifetime, W. D. Ross followed Moore and argued that moral significance, broadly construed, is grounded in an objective nonnatural *relational* property. Though Ross disagrees with Moore's consequentialism, he writes, "...if we ask ourselves whether it is qua, the packing and posting of a book, or qua the securing of my friend's getting what I have promised to return to him, that my action is right, it is clear that it is in the second capacity that it is [only] right... it is right by its own nature and not because of its consequences" (1930, 44). Furthermore, Ross agrees with Moore that intuitions that pick out this nonnatural moral property are important, he states, "The existing body of moral convictions of

the best people is the cumulative product of the moral reflection of many generations, which has developed an extremely delicate power of appreciation of moral distinctions” (41). Moore’s influence on moral realism reverberates today, as Geoffrey Sayer-McCord argues that distinctive moral kinds are the basis of moral terms. He writes, “On the moral kinds approach, two people will be seen as using moral terms with the same meaning if, but only if, the use of the terms in question are appropriately causally regulated by the same moral, not natural, kind” (1997). We can ignore the subtle intricacies of Moore’s, Ross’, and Sayer-McCord’s claims and recognize that, while there has been a lot of push back on moral realism, there has also been a consistent and growing espouse of moral properties as the basis of moral value. Crucially, an overarching goal of these theories is to establish that morality is an objective phenomenon, and that correct apprehension of moral properties is the basis of accurate moral judgements. Recent work in psychology also demonstrates a similar emphasis on distinguishing accurate moral judgements in virtue of a specific kind of content.

Psychologists such as Kurt Gray et al. and E. B. Royzman et al. have stressed an importance of explaining and distinguishing morality through content, particularly in terms of reference to harm (Gray et al., 2012, 2014; Royzman et al., 2009). Gray et al. writes, “not only are moral judgments sensitive to perceived agency and experience, but all moral transgressions are fundamentally understood as agency plus experienced suffering—that is, interpersonal harm—even ostensibly harmless acts such as purity violations” (2014). In other words, Gray et al. propose that morality (and moral judgement) is rooted in the perception of at least two agents and interpersonal harm; they explicitly argue that “distinct moral domains can be understood through the lens of intention and suffering,” and that “The presence or absence of harm...distinguishes moral transgressions from conventional transgressions...” (2014). Royzman

et al. likewise argue that the moral domain can be distinguished in virtue of content specifying harm. They write, “the findings reviewed and reported here are largely consistent with Turiel’s basic insight that the process of selective moralization is effected by a system oriented towards a particular rule content and that this content is largely defined by acts or dispositions deemed intrinsically harmful to others” (2009). Emma Buchtel and colleagues cite these studies and base their research question on whether all moral judgements *do* reference harm. They explain that their research shows, “Chinese lay moral cognition may be closely tied to judgments of civility and politeness,” and they interpret this to be “in contrast to both classic and recent psychology theories of morality, which argue that moral judgment is universally based on perceptions of harm and suffering, justice, and fairness” (2015). These studies by Kurt Gray et al., E. B. Royzman et al. and Buchtel et al. illustrate a contemporary environment in psychology to analyze morality and moral judgement in terms of a certain kind of content.¹⁷

A long history of emphasis in moral philosophy on content and a similar, periodic surge in moral psychology might understandably lead Stich’s argument to assume that if there is no distinguishing content to moral judgements, then they must not be psychologically distinguishable from conventional judgements. And while we cannot know if Stich’s argument was explicitly influenced by this large history of scholarship in philosophy and psychology, its implicit argumentative mistakes seem symptomatic of this historical movement. For this reason, I believe moral philosophy and psychology should expand its explanatory net and begins to look at attitudinal features of human psychology.

¹⁷ David Velleman and Walter Sinnott-Armstrong have recently pushed for moral philosophy and psychology to steer clear of content explanations; but they seem to be exceptions (Velleman, 1992; Sinnott-Armstrong, 2016).

6 CONCLUSION

In response to arguments by Stephen Stich, a more precise characterization of the psychological states of moral judgements needs to be developed. Stich argues that the content of moral norms does not differ from conventional norms; he plausibly demonstrates that there is extensive overlap between the content of putative moral norms and putative conventional norms. Stich infers from this that there is no psychological distinction between moral and conventional judgements. But his inference rests on an implicit, mistaken assumption that moral judgements are psychologically distinguished from conventional judgements only if they concern different kinds of content. I have argued that this assumption is false because moral judgements can be psychologically distinguished in virtue of attitude type. I presented a theory of a moral attitude through the characterization of its functional features to demonstrate the plausibility that such a moral attitude type exists and distinguishes the psychology of moral judgements from conventional judgements. What I have offered is a first pass theory, and it is possible that there are features of our moral attitude that I have not addressed here. Nonetheless, when philosophers and psychologists investigate the distinction between moral and conventional judgements, I believe they should turn their attention towards this promising but underdeveloped area of research.

WORKS CITED

- Baldwin, Tom, "George Edward Moore", *The Stanford Encyclopedia of Philosophy* (Summer 2010 Edition), Edward N. Zalta (ed.)
- Buchtel, E., Guan, T., Peng, Q., Su Y., Sang, B., Chen, S., & Bond, M. (2015). Immorality East and West: Are prototypically immoral behaviors especially harmful, or especially uncultured? *Personality and Social Psychology Bulletin*, 41(10), 1382-1394.
- Devitt, M. (1996). *Coming to our senses*. Cambridge, UK: Cambridge University Press.
- Fessler, D. M. T., Barrett, H. C., Kanovsky, M., Stich, S., Holbrook, C., Henrich, J., ... Laurence, S. (2015), Moral parochialism and contextual contingency across seven disparate societies. *Proceedings of the Royal Society B*, 282, 1823
- Gray, K., Schein, C., & Ward, A. F. (2014). The myth of harmless wrongs in moral cognition: Automatic dyadic completion from sin to suffering. *Journal of Experimental Psychology: General*, 143, 1600-1615.
- Gray, K., Young, L., & Waytz, A. (2012). Mind perception is the essence of morality. *Psychological Inquiry*, 23, 101-124. doi: 10.1080/1047840X.2012.651387
- Haidt, J., Koller, S., & Dias, M. (1993). Affect, culture and morality, or is it wrong to eat your dog? *Journal of Personality and Social Psychology*, 65, 613-628.
- Henrich, J., Heine, S., & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and Brain Sciences*, 33(2-3), 61-83. doi:10.1017/S0140525X0999152X
- Jackson, F. (1998). *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford University Press.
- Kelly, D., Stich, S. P., Haley, K., Eng, S., & Fessler, D. (2007). Harm, affect and the moral/conventional distinction. *Mind and Language*, 22, 117-131).

- Kornblith, H. (1998). The role of intuition in philosophical inquiry: An account with no unnatural ingredients. In M. DePaul & W. Ramsey (Eds.), *Rethinking intuition* (pp. 129-143). Lanham, MD: Rowman & Littlefield.
- Laurence, S., & Margolis, E. (1999). Concepts and cognitive science. In E. Margolis & S. Laurence (Eds.) *Concepts: Core readings*. Cambridge, MA: MIT Press.
- Levine, S., Rottman, J., Davis, T., Stich, S., & Machery, E. (2017). *Religion and the scope of the moral domain*. Manuscript in preparation.
- Moore, G. E. (1903/1993). *Principia Ethica* Cambridge University Press, Cambridge: 1903.
*Revised edition with 'Preface to the second edition' and other papers, ed. T. Baldwin, Cambridge University Press, Cambridge: 1903
- Moore, G. E. (1922). "The Conception of Intrinsic Value" in *Philosophical Studies*. Reprinted in Baldwin's *Revised edition of *Principia Ethica*, 280-98.
- Moore, G. E. (1899). "The Nature of Judgment" *Mind* 8 176-93. Reprinted in **G. E. Moore: Selected Writings* 1-19.
- Nisan, M. (1987). Moral norms and social conventions: A cross-cultural comparison. *Developmental Psychology*, 23, 719-725.
- Prinz, J. J. (2007). *The emotional construction of morals*. Oxford: Oxford University Press.
- Putnam, H. (1975). The meaning of "meaning." In K. Gunderson (Ed.), *Minnesota studies in the philosophy of science: Vol. 7. Language, mind and knowledge*. Minneapolis: University of Minnesota Press.
- Ross, W. D. (1930). *The Right and the Good*. Oxford: Clarendon Press
- Royzman, E. B., Leeman, R. F., & Baron, J. (2009). Unsentimental ethics: Towards a content-specific account of the moral-conventional distinction. *Cognition*, 112(1), 159–174.

- Sayre-McCord, G. (1997). "'Good' on Twin Earth" *Philosophical Issues*, vol. 8, pp. 267-292.
- Sinnott-Armstrong, W. (2016). "The Disunity of Morality", in Liao 2016: 331–354.
- Smith, A. (1759, 1976) *The Theory of Moral Sentiments*, ed. D.D. Raphael and A.L. Macfie (Liberty Fund, 1976). ISBN: 978-0-86597-012-0 (paperback)
- Smith, E., & Medin, D. (1981). *Categories and concepts*. Cambridge, MA: Harvard University Press.
- Sripada, C., & Stich, S. (2005). A Framework for the Psychology of Norms. *The Innate Mind: Culture and Cognition*. DOI: 10.1093/acprof:oso/9780199733477.003.0012.
- Stanford, P., K. (2018). Moral externalization and normativity: The errors of our ways. *Behavioral and Brain Sciences*, 41, E119. doi:10.1017/S0140525X17002254
- Stich, S. (2019)/a "The Quest for the Boundaries of Morality," in Aaron Zimmerman, Karen Jones & Mark Timmons, eds., *The Routledge Handbook of Moral Epistemology*, (New York: Routledge). Pp. 15-37.
- Stich, S. (2019)/b Did religion play a role in the evolution of morality?, *Religion, Brain & Behavior*, DOI: 10.1080/2153599X.2019.1678511
- Stich, S. (2018) "The Moral Domain," in *The Atlas of Moral Psychology* ed. by Kurt Gray & Jesse Graham (New York & London: Guilford Press). Pp. 547-555.
- Turiel, E. (1978). Distinct conceptual and developmental domains: Social convention and morality. In H. Howe & c. Keasey (Eds.), *Nebraska Symposium on Motivation, 1977: Vol. 25. Social cognitive development*. Lincoln: University of Nebraska Press.
- Turiel, E. (1983). *The development of social knowledge*. Cambridge, UK: Cambridge University Press.

Turiel, E., Killen, M., & Helwig, C. (1987). Morality: Its structure, functions, and vagaries. In J. Kagan & S. Lamb (Eds.), *The emergence of morality in young children* (pp. 155-243).

Chicago: University of Chicago Press.

Velleman, J. D. (1992). The Guise of the Good. *Nous*, 26(1), 3-26. doi:10.2307/2215684