

*Unpublished draft: Please do not cite without permission from the author.*

## **Meta-Ethical Pluralism: Exploring the Evidence**

Jennifer Cole Wright

College of Charleston

Correspondence concerning this manuscript should be sent to: Jennifer Cole Wright, Department of Psychology, College of Charleston, 57 Coming Street, Charleston, SC 29424, email: [wrightjj1@cofc.edu](mailto:wrightjj1@cofc.edu).

Key words: Tolerance, Meta-ethics, Objectivism, Non-Objectivism, Pluralism, Variability, Incoherence

Word-count: 9,821

## Abstract

At the heart of meta-ethics is the debate about whether or not people’s moral claims assume the existence of moral facts—and if they do, whether or not those facts are non-relative and/or mind-independent. Often underlying this debate is the assumption that moral discourse is relatively uniform. But recently philosophers have challenged this assumption, arguing that moral discourse is actually quite variable (Gill, 2009)—even “incoherent” (Loeb, 2008)—displaying commitments to objectivism *and* non-objectivism, realism *and* anti-realism. In this paper, I provide empirical support for the existence of meta-ethical variability across several different assessments of people’s moral discourse. Further, I argue for the view that this “pluralism”—while perhaps *metaphysically* untenable—reflects a coherent (if implicit) pragmatic strategy meant to facilitate our navigation through ever-changing normative space by modulating the level of personal choice and social debate about particular moral issues that is deemed acceptable, both within and between socio-cultural groups.

## Meta-Ethical Pluralism: Exploring the Evidence

Jennifer Cole Wright

At the heart of meta-ethics<sup>1</sup> is the debate about whether people’s moral discourse (e.g., statements such as “Abortion is wrong!”) assumes the existence of moral facts—and if so, whether those facts are non-relative and/or mind-independent. This debate not only involves the question about what is in fact the case about moral claims, but also the question of what ordinary, competent language users (the “folk”) *take* to be the case when they utter moral statements. That is, when people say things like “Abortion is wrong!”, do they take themselves to be expressing their positive/negative feelings towards abortion, or other “pro/con” attitudes or affective affiliations with community norms—or do they take themselves to be conveying beliefs about matters of fact? And, if the latter, do they consider those facts to be objective or non-objective in nature?

Philosophers have employed a variety of intuitive, theoretical, and observational strategies to argue for both sides of the debate, and they remain strongly entrenched on either side—some arguing that the folk’s moral claims reflect an underlying commitment to moral realism/objectivism (Brink 1989; Darwall 1998; Harman 1996; Mackie 1977; Shafer-Landau 2003; Smith 1994), others arguing the contrary (Ayer 1952; Blackburn 1984; Dreier 1999; Gibbard 1990; Rachels & Rachels 2009; Wong 1984, 2006). Recently, however, Loeb (2008) has argued that the perpetuation of this long-standing debate between philosophers “...is evidence that inconsistent elements—in particular, commitments both to and against

---

<sup>1</sup> Or at least, as Gill (2009) calls it, “descriptive meta-ethics”, which involves providing the best analysis of our ordinary moral discourse.

objectivity—may be part of any accurate understanding of the central moral terms...” (p. 358).

In other words, his supposition is that people may be *both* objectivists and non-objectivists—that they “...use moral words *both* to make [objective] factual assertions *and* to do something incompatible with the making of such assertions...” (p. 363).

Similarly, Gill (2009) challenges what he calls the “Uniformity-Determinacy” (UD) thesis underlying most of 20<sup>th</sup> century meta-ethics, which holds that “our ordinary [moral] discourse is uniform and determinate enough to vindicate one side or the other” of the meta-ethical debate. As an alternative, he proposes both that (at least some parts of) our ordinary moral discourse failed to support either meta-ethical position (the Indeterminacy Thesis) and/or that sometimes provide support for one meta-ethical position, while yet other times providing support for the opposing position (the Variability Thesis).<sup>2</sup>

In this paper, I will be examining whether or not there is evidence for the latter of these theses, the possibility that people’s moral discourse will sometimes reflect cognitivist/objectivist commitments, while sometimes reflecting non-cognitivist/non-objectivist commitments. If such variability in our moral discourse exists, then we may be able to uncover instances of it using empirical methods. Considering that the project at hand is descriptive meta-ethics, such an endeavor seems well motivated—indeed, Loeb (2008) emphasized the importance of an empirical investigation into this issue, holding that the discussion of people’s meta-ethical commitments must at some point involve an examination of their actual “linguistic dispositions” (p. 356; i.e., their intuitions, patterns of thinking and speaking, semantic commitments, and other internal states—both conscious and not). Gill (2009) agrees with this sentiment, writing that “If

---

<sup>2</sup> Of the two, Loeb’s (2008) incoherentist view is more extreme than Gill’s (2009) variabilist view insofar as he maintains that the realist/anti-realist commitments reflected in people’s ordinary moral discourse cannot be disentangled from one another into discrete uses that reflect one kind of commitment, on the one hand, and the other kind of commitment, on the other (as the Variability Thesis suggests). See Sinnott-Armstrong (2009) for a discussion of these and other “mixed” meta-ethical positions.

we leave open the possibility that our uses of moral terms are more variable and indeterminate than previously assumed, descriptive meta-ethics will have to involve much more empirical investigation...we will have to give much more time and effort than most 20th century meta-ethicists did to the gathering of data” (p. 232)—one recommended approach being the testing of people’s reactions to moral statements under “experimentally responsible conditions” (such as those employed by social scientists).

Happily, though the status of people’s meta-ethical commitments has been a subject that has heretofore failed to attract attention from social scientists, a recent spark of interest has generated a handful of empirical inquiries, the results of which are intriguing, if preliminary. Of particular interest is the fact that (though not always acknowledged by the researchers themselves) much of data collected provides at least a *prima facie* challenge to the UD thesis, revealing what appears to be a mixture of *both* objective and non-objective meta-ethical commitments on the part of ordinary folk, both between and within persons.

This variability shows up in (at least) two forms. First, there is some evidence for *interpersonal* differences in people’s meta-ethical commitments. Nichols (2004), for example, found that while some of the people he interviewed gave objectivist responses to questions about moral disagreement—stating that the wrongness of a particular action was grounded by objective facts (and that, therefore, if two people disagreed about it, one would be mistaken)—others did not, stating instead that there were no objective facts of the matter and that people making contrary claims could be correct. And Sarkissian, Park, Tien, Wright, & Knobe (2012) found that while people gave objectivist responses to moral disagreement that occurred between members of the same culture, they gave more non-objectivist responses when it occurred between members of different cultures.

There is also some evidence for *intrapersonal* variability as well. Goodwin and Darley (2008, 2010, 2012) found that while people gave objectivist responses more frequently for moral issues than for other sorts of (non-moral) issues, they were nonetheless internally inconsistent in their meta-ethical stance, giving objectivist responses for some of the moral issues, but not others. Following up on this, Wright, Grandjean, and McWhite (2012) and Wright, McWhite, and Grandjean (2014) replicated (and expanded upon) this pattern of variability, verifying its presence even under more carefully constructed empirical conditions.

Of course, these early studies have only begun to scratch the surface of people's ordinary moral discourse and the underlying meta-ethical commitments it reveals. For one, thus far the questions employed have been both limited in scope and (as many concerned philosophers have pointed out)<sup>3</sup> insufficiently nuanced, at best only partially capturing the features that distinguish an objectivist from a non-objectivist stance. In addition, they have failed to adequately tease apart the major non-objectivist positions—e.g., non-cognitivism, subjectivism, and relativism. The goal of this paper, therefore, is to report on our recent attempts to advance this investigation further. To do this, we employed more finely-tuned assessments of meta-ethical commitments that were designed to better disentangle the different non-objectivist stances.

### **Meta-Ethical Commitments, Take 1: Objectivism vs. Subjectivism/Relativism**

Our first attempt was to explore people's meta-ethical commitments along the *objectivism/non-objectivism* dimension. This required distinguishing between two related, but distinct, forms of non-objectivism: namely, *subjectivism* (i.e., the truth of a moral statement is determined by the acting person's beliefs, attitudes, or feelings about the action/type of action or

---

<sup>3</sup> While I'm sure this criticism has made it into print by now, my main exposure to it has been through personal conversations and other forms of communication with philosophers interested in the research and its implications.

by the beliefs, attitudes, and feelings had by the culture to which the person belongs) and *relativism* (i.e., the truth of a moral statement is relative to the individual making it and/or culture to which the individual belongs).

The first step in the process was to successfully identify moral issues.<sup>4</sup> To this end, we asked people<sup>5</sup> to consider 20 different issues (e.g., “Smoking cigarettes in an enclosed public space”; “Engaging in prostitution”; “Selling children on the internet”; see Table 1 for a complete list) and identify what *kind of issue* they believed each to be: *moral* or *non-moral* (the “non-moral” choices being *personal preference/choice* or *social convention/norm*). This step allowed us to individualize our analysis of meta-ethical commitments to those issues each person actually considered to be moral.

Of the 20 issues presented, none were unanimously classified as moral (or, for that matter, non-moral), though 10 received a dominant<sup>6</sup> moral classification. These included issues such as child trafficking (selling children on the internet), stealing, racial/gender discrimination, incest, euthanasia, and prostitution (see Table 1 for the breakdown).

### ***Objectivism vs. Subjectivism***

Once this step was completed, people were interviewed about their meta-ethical commitments, employing those issues they had previously classified as moral. In order to examine their commitments to *subjectivism* specifically, participants were first asked whether it to be morally acceptable or unacceptable for them to *x*—the “*x*” being filled in with each moral issue under consideration (e.g., engaging in prostitution). After this, they were asked to consider

---

<sup>4</sup> While a seemingly straightforward task, previous research suggests that there is actually substantial disagreement, both between and within age groups, about which issues should be considered “moral” (Wright 2011; Wright et al. 2008, 2012), even though there is remarkable agreement about what features (harm/unfairness) make something a moral, rather than a non-moral, issue.

<sup>5</sup> 164 participants respond to the online survey we posted on Amazon Mechanical Turk (70% female; 89% Caucasian, 5% African-American, 1% Asian-American, 3% Hispanic, and 2% other).

<sup>6</sup> From here on, by “dominant” I mean above the level expected by mere chance, which is a percentage (in this case, at least 39%) that is generated by the number of choices a person can make and the number of people making them.

for each moral issue someone making the (sincerely endorsed) statement reflecting the opposite stance—i.e., if they had reported  $x$  being morally unacceptable, the statement they were asked to consider was “It is morally acceptable to  $x$ ” and vice versa). For each issue, they were then asked to report what they believed to be the case, their choices begin:

- *It would be morally (un)acceptable for that person to  $x$ .* The person would be correct because the rightness/wrongness of  $x$  is determined individually, by each person's beliefs, attitudes, or feelings about the act of  $x$ -ing or type of action that  $x$ -ing is.
- *It may or may not be morally (un)acceptable for that person to  $x$ .* Whether or not the person was correct would depend on the community in which that person lives. The rightness/wrongness of  $x$  is determined by a community's collective beliefs, attitudes, or feelings about the act of  $x$ -ing or type of action that  $x$ -ing is.
- *It would not be morally (un)acceptable for that person to  $x$ .* The person would be mistaken (as would anyone else who made this claim). The rightness/wrongness of  $x$  is determined by the action or type of action it is, independently of the person's or his/her community's beliefs, attitudes, or feelings about it. That is, there is something about  $x$ -ing or the type of action  $x$ -ing is that makes it right/wrong regardless of what that person or other people think or feel about it.

The first two options were coded as *subjective*; the last as *objective*.

The results showed clear variability in the meta-ethical responses people gave for each moral issue. For example, of the 10 issues dominantly classified as moral, none received a consistent meta-ethical response—7 were dominantly classified as objective, while the remaining 3 were dominantly classified as subjective (Table 2). Importantly, those issues dominantly classified as non-moral (either personal or social) were also dominantly classified as subjective.

There was also substantial *intrapersonal* variability: 78% of the people were pluralists, giving objective responses to some of the moral statements and subjective responses to others, while only 19% of the participants were consistent objectivists and 3% consistent subjectivists (Table 3).

## *Non-Relativism vs. Relativism*

In order to examine their commitments to *relativism*, participants were first given a short paragraph introducing them to the idea of relative vs. non-relative terms.

Consider the difference between the term “triangular” vs. the term “tall”. The first of these terms is a *non-relative* term, meaning that the context in which it is uttered does not influence its truth value—e.g., the statement “That shape is triangular [i.e., it is a shape with three sides and three corners]” is either true or false *of the shape being talked about* no matter who says it, when it is said, or what frame of reference is being used. If it is true that the shape being referred to is triangular in one context, then (barring something happening to change the shape) it will always be true that it is triangular, regardless of the person making the statement and/or the time, place, situation in which it is uttered.

On the other hand, “tall” is a *relative* term, and, therefore, the statement “Naomi is tall” could be true or false, depending on the context/the frame of reference under which it is uttered—e.g., whether we are comparing Naomi, who stands 5’6”, to a group of women from a Black Hmong village in Vietnam (who, at their tallest, stand about 5”) or to a group of NBA players (who, on average, stand about 6’7”). It would also be the case that we’d consider the statement “Naomi is tall” to be true if uttered by a Black Hmong woman, but not true if uttered by an NBA player. In other words, for relative terms, the person making the statement and/or the time, place, situation in which it is uttered makes a difference. Frame of reference is important for determining truth-values.

Please keep this distinction between *relative* and *non-relative* terms in mind as you participate in the next exercise.

After reading this short paragraph, people were asked to consider for each of their moral issues two people who read about someone engaging in that behavior and then making two different (and opposing) statements about it—one stating that “It was wrong for that person to *x*” and the other stating that “It was not wrong for that person to *x*”.<sup>7</sup> Given these opposing statements, participants were asked to choose which they thought would be the case, between:

- ONLY ONE of these statements would be correct—either it is true that it was wrong for that person to *x* or it is true that it was not wrong for that person to *x*, regardless of who is making the statement or the contexts in which it is being made. Both statements cannot be correct.

---

<sup>7</sup> Strictly speaking, what we are testing here is *appraiser*, rather than *agent*, relativism. While agent relativism holds that the appropriate frame of reference is the moral framework of the person who performs the act, or of the cultural group to which the person belongs, appraiser relativism holds that the appropriate frame of reference is the moral framework of the person who makes the moral judgment, or of the cultural group to which the person belongs.

- BOTH of these statements could be correct—whether it is true that it was wrong for that person to  $x$  or not wrong for that person to  $x$  depends on who is making the statement and/or the context in which it is made.<sup>8</sup>

The first option was coded as *non-relative*; the second as *relative*.

Once again, the results showed clear variability in whether people considered the truth of the moral statements to be relativized. For the same 10 issues as above, the statements about 5 of them were considered dominantly non-relative, whereas the statements about 3 of them were considered dominantly relative—the remaining 2 were split (Table 2). Again, statements about the issues classified as non-moral were also considered dominantly relative.

There also continued to be intrapersonal variability: 85% of the participants were pluralists, with only 8% of the participants consistently endorsing non-relativism and 7% consistently endorsing relativism (Table 3).

### ***Categorical vs. Non-Categorical Imperatives***

In his presentation of his “error theory”, Mackie argued that people’s moral claims made reference to objective values—values that were “categorically imperative” in nature (1977, p. 29) and, therefore, providing people with a reason to do/not do the action *independently from* (and even *in spite of*) any actual desires, inclinations, beliefs (etc.) that they might have to do/not do it. Therefore, in addition to the above questions, we decided to also ask participants to consider, for each issue  $x$ , the following question:

If there were people who did not believe that there was anything wrong with doing  $x$  (or  $x$ -ing)—and, indeed, they wanted to do it— would there be any reason for that person to nonetheless refrain from doing it?

For each issue, they were asked to pick the best response from the following:

- *There would be no reason for them not to  $x$ .* They should feel free to  $x$  if they so desired.

---

<sup>8</sup> Thank you to Edward Jarvis and others for help in formulating the subjectivism and relativism questions.

- *There still might be a reason for them not to x.* People in their family/community might disapprove of *x*-ing or type of action that *x*-ing is.
- *There still might be a reason for them not x.* It is against the law and they could get in trouble for *x*-ing or for engaging in the type of action that *x*-ing is.
- *There is still a strong reason for them not to x.* It would be bad<sup>9</sup> for them to *x*, even if they don't think so and they wanted to do it (and even if no one else would disapprove or punish them for doing so).

The first three options were coded as *non-categorical*; the last as *categorical*.

Not surprisingly, people also displayed variability in their responses here—4 of the 10 moral issues were dominantly considered categorical, while 4 were dominantly considered non-categorical, with the remaining 2 being split (Table 2). The non-moral issues were all dominantly classified as non-categorical. And, 82% percent of the participants were pluralists, with 12% consistently choosing the categorical response and 6% the non-categorical response (Table 3).

In sum, our investigation revealed a high and consistent degree of variability (or, what I call “pluralism”) in the way people think about moral issues and evaluate moral discourse. Across several distinct lines of questioning, people were both objectivists—and non-objectivists. They viewed the status of some moral actions as objectively determined, moral claims about those actions non-relative, and reasons for/against them categorical, while at the same time viewing the status of other moral actions as subjectively determined, moral claims about them relative, and reasons for/against them non-categorical. And they disagreed—for some moral issues more than others—about whether any particular moral issue should be treated objectively or not.

---

<sup>9</sup> We used the term “bad” here because elsewhere we’d used “unacceptable” and “wrong” and we wanted to demonstrate that variability in people’s meta-ethical responses is present across multiple moral terms.

## Meta-Ethical Commitments, Take 2: Cognitivism vs. Non-Cognitivism

Our next project was to explore people's meta-ethical commitments along the *cognitivism/non-cognitivism* dimension, examining whether people view moral statements as being the sorts of things that are either true or false, like beliefs (*cognitivism*) or the sorts of things that are neither true nor false, like feelings or attitudes (*non-cognitivism*).

In order to examine this, we had to first distinguish between two distinct non-cognitivist claims (Joyce, 2009). The first is the claim of *semantic nonfactualism*, which is the denial that moral statements express propositions or have truth conditions (i.e., that they are “truth-apt”). The second is the claim of *psychological non-cognitivism*, which denies that the mental states moral statements are conventionally intended to convey are beliefs (or other related cognitive mental states). While typically non-cognitivists accept both of these negative claims, they can nonetheless come apart, so we decided to test each separately.

We anticipated that the difference between the cognitivist and non-cognitivist positions would not be obvious to people, so we carefully constructed Introductory Exercises that people had to complete before they could move forward. Only those people who did so successfully were invited to complete the rest of the survey.

### *Semantic Non-Factualism*

In order to properly attune people<sup>10</sup> to the difference between statements that are “truth-apt” and those that are not, they were instructed to read through the following very carefully and then answer some questions at the end:

---

<sup>10</sup> We had 122 participants respond to an online survey, either as a part of a college study or posted on Amazon Mechanical Turk (67% female; 88% Caucasian, 7% African-American, 2% Asian-American, and 3% Hispanic).

Some statements assert propositions that are what we call “truth-apt” – that is, they are meant to reflect *matters of fact* about the world (though sometimes they may fail to do so), which means they will be either true or false. For example, if I said to someone that “Boston, MA is north of Miami, FL” I would be stating something that is truth-apt – it is either true or false. In this case, we can easily establish whether my statement is true or false (e.g., by looking at a map). And, as it turns out, it is true. If, however, I had stated that “Boston, MA is south of Miami, FL”, it would have been false. Either way, the important thing is that there is a *fact of the matter* (in this case, the geographical relationship between Boston and Miami) that my statement was meant to assert.

Determining the truth/falsity of statements like the above is relatively easy. But sometimes it isn’t easy. Consider, for example, the statement that “The earth is the only planet in our galaxy with life on it”. We simply don’t know at this point (and, indeed, we may never know) whether this statement accurately reflects a matter of fact (that is – whether it accurately reflects how many planets in our galaxy actually currently support life). So, we have no way of establishing whether the statement it is true or false – but, nonetheless, it is still truth-apt. It is either true or false – i.e., either the earth is the only planet in our galaxy with life on it or it isn’t. So, if one person said “The earth is the only planet in our galaxy with life on it” and another person said “Earth is not the only planet in our galaxy with life on it”, one of these people would be correct and the other one mistaken (even if we can’t say at this point which one is which).

Consider, on the other hand, claims like “Peanut butter ice cream is delicious” or “Jazz music is the best form of music ever invented” or “Riding on the roller coaster at Eliches is awesome!” Unlike the statements considered above, these statements aren’t truth-apt. They are neither true, nor false – there isn’t a *fact of the matter* about the world that they are intended to reflect. In other words, there isn’t an actual fact of the matter about whether peanut butter ice cream tastes delicious or riding the roller coaster at Eliches is awesome. Some people enjoy the taste of peanut butter ice cream, others don’t; Some people have a great time riding the roller coaster at Eliches, others don’t. So, if one person said “Riding roller coasters is awesome!” and another person said “Riding roller coasters is absolutely terrifying!” it wouldn’t make sense to say that one of the two was correct and the other mistaken. This is because neither of these statements are intended to accurately reflect some fact about roller coaster riding – rather, they are expressions of people’s liking/disliking of or approval/disapproval for something (in this case, riding roller coasters). In other words, statements like “Riding on roller coasters is exciting” or “Peanut butter ice cream is delicious” are not truth-apt – they are neither true, nor false. Instead, they are expressions of what we call people’s “pro/con attitudes” (i.e., their positive/negative feelings, likes/dislikes, approval/disapproval, etc.).

It is important to recognize that truth-apt statements about ice cream and roller coaster riding can be made – for example, “Meredith hates peanut butter ice cream” or “I really love riding the roller coaster at Eliches” are both statements that are either true or false (either Meredith hates peanut butter ice cream or she doesn’t, etc.). To illustrate further: Imagine that Meredith said “I hate peanut butter ice cream”. In this case, she’d be stating something that is truth-apt, since her statement asserts a fact of the matter about herself

(namely, that she hates peanut butter ice cream). But if instead she said “Peanut butter ice cream is disgusting”, she’d be stating something that is not truth-apt, since it is a statement intended to express her dislike of peanut butter ice cream.

For the questions that follow, please keep this distinction in mind, as you’ll be asked to identify which statements you think are “truth-apt” (i.e., asserting *matters of fact* that are either true or false) and which statements you think are not “truth-apt” (i.e., expressing *pro/con attitudes*, and so are neither true, nor false).

After reading this, people were then given 10 statements (5 of each type, see Table 4) and asked to identify them as either “truth-apt” or not. Only those participants who identified at least 9 of the 10 correctly (94 participants or 77%) were invited to continue with the study.

At this point, the remaining participants were given 20 different issues to consider (see Table 5 for a complete list). Once again, they were asked to identify what *kind of issue* they believed each issue to be, *moral* or *not-moral*. Of the 20 issues considered, 12 were dominantly classified as moral: including issues such as trafficking children (selling children on the internet), stealing, rape, infidelity, cheating, incest, and racial discrimination (Table 5).

Once this step was completed, participants were asked to consider for each issue the statement “It is wrong to *x*” (the “*x*” being filled in with the particular issue under consideration – e.g., selling children on the internet) and then identify whether they thought each statement was either:

- “truth-apt” (*assertions of matters of fact* that are either true or false)
- not “truth-apt” (*expressions of positive/negative feelings*, pro/con attitudes, etc. that are neither true, nor false)

Of the 12 issues dominantly identified as moral, 3 were dominantly identified as truth-apt, 4 as not truth-apt, and the remaining 5 were split between the two. The issues dominantly classified as non-moral were also dominantly identified as not truth-apt (Table 6). Thus, once again, we found meta-ethical variability, with people treating *some* moral statements as assertions of matters of fact, while treating *other* moral statements as expressions of pro/con (presumably con)

feelings/attitudes. And this variability was also once again intrapersonal: 76% of the people were pluralists, identifying some moral statements as truth-apt and others as not, while only 9% of the participants provided consistently cognitivist responses and 15% consistently non-cognitivist responses (Table 3).

### ***Psychological Non-Cognitivism***

Switching focus to the 2<sup>nd</sup> related, though distinct, claim of non-cognitivism (the denial that moral statements are meant to express beliefs or related “cognitive” mental states), we gave a separate group of people<sup>11</sup> the following Introductory Exercise:

People make different kinds of statements – some of which *assert beliefs*, others of which *express feelings*. Consider, for example, if I said to someone that “Boston, MA is north of Miami, FL”. What I am doing is expressing my belief that something is this case – namely, that there is a fact of the matter about the geographical relationship between Boston and Miami. My intention is to assert a belief, which in this case turns out to be true. But, there are also times when the beliefs we assert with our statements are false, like if I would have said “Boston, MA is south of Miami, FL” instead. But that does not change the fact that such statements assert a *belief* about something being the case. For our purposes, it doesn’t matter whether the beliefs being asserted are true or false – all that matters is that we sometimes make statements that are intended to assert beliefs about things that we take to be matters of fact about the world.

The same goes for statements that involve beliefs whose truth/falsity cannot be established. For example, I might state something like, “The earth is the only planet in our galaxy with life on it”. This isn’t the sort of belief that can currently be established as true or false – we don’t know at this point (and, indeed, we may never know) whether my belief accurately reflects a fact of the matter about life in the galaxy or not. But, nonetheless, my objective in making this statement is to assert a belief about something I take to be true, even if it can’t be established for sure whether or not I’m correct.

Consider, on the other hand, my statement that “Peanut butter ice cream is delicious!” or “Jazz music is the best form of music ever invented” or “The roller coaster at Eliches is terrifying!” Here, these statements are not intended to be assertions of beliefs about matters of fact – i.e., that peanut butter ice cream is the sort of thing that is, in fact, delicious or that riding the roller coaster at Eliches is the sort of activity that is terrifying). Rather, they are expressions of *positive and/or negative feelings* and *attitudes* that I have

---

<sup>11</sup> We had 116 participants respond to an online survey, either as a part of a college study or posted on Amazon Mechanical Turk (61% female; 89% Caucasian, 6% African-American, 1% Asian-American, and 4% Hispanic).

about the subject matter (in this case, really liking peanut butter ice cream and not liking the roller coaster at Eliches).

When I make these sorts of statements, I am fully aware that they aren't true or false (like the statements considered above). While it may be true that I like the taste of peanut butter ice cream and don't enjoy riding on the roller coaster at Eliches, there isn't actually a *fact of the matter* about peanut butter ice cream being delicious or the roller coaster being terrifying – after all, it would make perfect sense for someone to reasonably state the opposite and neither of us would be mistaken. In other words, the objective of statements like "Riding on roller coasters is terrifying" or "Peanut butter ice cream is delicious" is to express our positive/negative feelings (pro/con attitudes, liking/disliking, approval/disapproval) about something, not to assert beliefs about things that we take to be true.

Of course, I can believe (i.e., take it to be true) that I or someone else really likes peanut butter ice cream and doesn't like riding the roller coaster at Eliches and my statements can assert such beliefs – such as if, for example, I were to say "Meredith really loves peanut butter ice cream" or "Peanut butter ice cream is my favorite". These statements involve beliefs about Meredith and myself that are either true or false. But statements like "Peanut butter ice cream is disgusting!", on the other hand, are not.

To further illustrate, consider the following two statements:

- Larry loves Bon Jovi
- Bon Jovi rocks!

The first statement involves the *assertion of a belief* about Larry (namely, that he loves Bon Jovi – which could be true or false); the second, on the other hand, does not assert a belief (there is no fact of the matter about Bon Jovi "rocking" that can be established as true or false) but instead *expresses a person's positive attitude* (their appreciation, enjoyment, approval) towards Bon Jovi.

For the questions that follow, please keep this distinction in mind, as you'll be asked to identify those statements you think were intended to *assert beliefs about matters of fact*, those intended to *express positive/negative feelings, attitudes, etc.* about a topic, and those intended to do both.

After reading this, people were then given the same 10 statements (Table 4) and asked to identify them as statements that were intended to 1) express positive/negative feelings, attitudes, etc., 2) assert beliefs about matters of fact, or 3) both.<sup>12</sup> Only those participants who identified at least 9 of the 10 correctly (98 participants or 84%) were asked to continue with the study.

---

<sup>12</sup> Since often our utterances do more than one thing—e.g., convey a belief while at the same time expressing an emotional reaction—we felt that it was important to include this third option. But, for the purposes of analysis the

At this point, the remaining participants were given the same 20 issues to consider. As above, they were asked to identify what kind of issue they believed each to be and of the 20 issues considered, 11 were dominantly classified as “moral”—not surprisingly, the same as before, with the exception of 1<sup>st</sup> trimester abortions, which dropped below chance (Table 5).

Once this step was completed, they were asked to consider for each issue the statement “It is wrong to do *x*” (the “*x*” being the particular issue under consideration) and then identify whether the statement was intended to:

- assert *beliefs* about matters of fact
- express positive/negative *feelings, attitudes, etc.*
- both<sup>13</sup>

Of the 11 issues dominantly identified as moral, 3 were dominantly identified as only expressions of positive/negative feelings/attitudes, 5 were dominantly identified as asserting beliefs (either alone or accompanied by an expression of feelings/attitudes), and the remaining 3 were split. As before, all issues classified as non-moral were also dominantly identified as expressing positive/negative feelings/attitudes (Table 6).

Taken together, these results provide additional support for the existence of meta-ethical variability in people’s cognitivist/non-cognitivist commitments, viewing some moral statements as intended to only convey expressions of people’s feelings/attitudes, but others to convey beliefs about matters of fact. And again, there was substantial intrapersonal variability: 75% of people gave pluralist responses, while only 9% of the participants provided consistently cognitivist responses and 16% consistently non-cognitivist responses (Table 3).

---

last two options were collapsed together since the issue of primary importance was whether there was an assertion of a belief, regardless of whether or not it was accompanied by an expression of a feeling/attitude.

<sup>13</sup> As Sinnott-Armstrong (2009) reminds us, it is not inconsistent with cognitivism for feelings/attitudes to be uttered alongside beliefs in the same statement, “We can state facts and also express emotions or issue imperatives at the same time, such as when I say ‘There’s a spider on your leg’ in order to state a fact and also to express fear and to warn and alert you to danger (Sinnott-Armstrong 1993).

## Significance of Meta-Ethical Pluralism

All together, the “descriptive meta-ethical data” collected and reported here, alongside previous research, supports the presence of variability in the way people think and talk about moral issues across several distinct meta-ethical positions. And though it is certainly likely that the methods utilized here, while an improvement on previous research, have still failed to fully capture relevant meta-ethical distinctions and/or tap into people’s meta-ethical commitments in important ways, they nonetheless reveal that something very interesting is going on with people’s moral discourse.

So what is to be made of this variability—this “pluralism”—in people’s meta-ethical commitments? Gill (2009) suggests that we may simply be guilty of a degree of “internal incoherence” and the variability in our moral discourse evidence of mistakes in our thinking (and speaking) about moral issues—mistakes that, unfortunately, could result in the adoption of “morally disastrous views” (p. 233). And Loeb (2008) argues that “...ordinary people are at bottom widely and irremediably, if perhaps only implicitly, conflicted about questions of moral objectivity” (p. 363), rendering their moral discourse irremediably incoherent.

I am aware of no philosophical position that provides a cogent argument for a metaphysical (or epistemological) pluralism in meta-ethics, the absence of which is a *prima facie* reason to view any pluralism as the result of confusion (or incoherence) on the part of ordinary moral language users. Nonetheless, I would argue that there is a *pragmatic* “method in the madness” (likely implicit) to meta-ethical pluralism—namely, that it aids in the individual and collective navigation of normative space within an ever-changing social environment.

Elsewhere (Wright et al., 2012, 2014) I’ve argued for the view that meta-ethical pluralism serves an important “psycho-social function” by modulating the level of permissible choice and

dialogue about moral issues, both within and between socio-cultural groups. To illustrate: viewing a moral issue as objectively-grounded removes it from the realm of legitimate personal/social negotiation (i.e., individual and/or social attempts to condone and/or promote it are deemed unacceptable, censorship/prohibition supported). On the other hand, viewing a moral issue as non-objectively grounded allows people to acknowledge its *moral* significance (i.e., that it is not simply a personal/conventional matter), while at the same time maintaining room for open dialogue and debate—thus, social censorship/prohibition are viewed less favorably. While support for this view has already been discussed (Wright et al., 2012, 2014), we decided to examine it further here.

For the first set of studies (*Objectivism vs. Subjectivism/Relativism*) people's level of public and private tolerance for potentially divergent moral beliefs/values was assessed by asking them how willing they would be to privately and publicly engage in dialog with someone who believed that *x*-ing (the "*x*" being filled in with each moral issue) was morally acceptable, as well as to privately and publicly support and/or shun him/her for engaging in *x*-ing (3 = *very willing* to -3 = *very unwilling*). In addition, they were asked both to consider what would be true of the person *x*-ing (e.g., "They would probably not be that different from me"; "They would not have the same values as I do"; "They would be a bad person", etc.) and what would be true of themselves if were they to *x* as well (e.g., "I'd never do it!"; "I'd be happy/proud that I did it, even if other people didn't accept it"; "I wouldn't think it's that big of deal") from 3 = *very true* to -3 = *very not true*. These questions were then collapsed into four summary tolerance variables: *private tolerance* ( $\alpha = .85$ ), *public tolerance* ( $\alpha = .70$ ), *judgments about other* ( $\alpha = .89$ ), and *judgments about self* ( $\alpha = .68$ ).

Comparing the difference in people's tolerance for divergent moral actions that they viewed as objectively grounded to those they viewed as non-objectively grounded revealed that they were significantly less tolerant of divergent moral actions that they viewed to be objectively-grounded—both privately and publicly.<sup>14</sup> They also had more negative views of anyone who engaged in these issue-relevant actions and of themselves, were they to engage in them (Fig. 1).<sup>15</sup>

For the second set of studies (*Cognitivism vs. Non-Cognitivism*) we added questions about people's internal motivation. For each issue, they were asked what would be the case for someone who had honestly stated "It's wrong to *x*", and yet had done *x*. Specifically, they were asked to rate (3 = *strongly agree* to -3 = *strongly disagree*) the degree to which the following would be the case:

- They would feel guilty if they ended up doing it.
- They would feel disgusted by others that did it.
- They would feel motivated not to do it.
- They would feel outraged at others that did it.
- If they did it, they would actively try to avoid doing it again in the future.

People were also asked questions about how they would respond towards someone who engaged in each issue-relevant action, their response options including both negative (e.g., "I would try to convince him/her to stop", "I would be disgusted", "I would seek to have that person punished") and positive (e.g., "I would think that person was worth getting to know better", "I would understand, it's his/her choice") statements. These questions were then collapsed into two summary tolerance variables: *intolerance* ( $\alpha = .85$ ), *tolerance* ( $\alpha = .71$ ).

Comparing the difference in people's tolerance for divergent moral actions that they viewed as non-cognitive (in the sense that they had classified people's previous statements about

---

<sup>14</sup> Paired-sample t-tests,  $t_s(157) = 8.8$  to  $10.3$ ,  $ps < .001$

<sup>15</sup>  $t(157) = 10.1$ ,  $p < .001$  and  $t(157) = 4.6$ ,  $p < .001$

them as *only* expressing feelings/attitudes) to those they viewed as cognitive (previous statements having been classified as asserting beliefs, with or without feelings/attitudes accompanying) revealed that people attributed significantly stronger internal motivation to those who had engaged in the latter type of moral actions.<sup>16</sup> They also displayed significantly less positive tolerance—and more negative tolerance—towards anyone who engaged in them.<sup>17</sup> And finally, they reported that it was significantly less okay for either themselves or someone else to engage in them (Fig. 2).<sup>18</sup>

These results together continue to support the view previously proposed (Wright et al., 2012, 2014), that meta-ethical pluralism allows for divergent moral actions to be treated differently—some tolerated to a degree (though certainly not as much as divergent non-moral actions); others strongly prohibited and punished. And it allows for social discussion about some moral issues to remain open, while discussion about other moral issues is considered “closed”. And while all moral issues are expected to “pull hard on us” internally, motivating us to refrain from them—and making us feel ashamed/guilty/disgusted at ourselves and others when we fail—some moral issues are expected to have more of a pull than others.

## **Concluding Remarks**

An objectivist might argue that just as many of our current beliefs about the physical nature of the world will one day turn out to be false (and therefore need to be replaced) many of our current moral beliefs will likewise turn out to be wrong—and therefore need to be changed. But the socio-cultural (and personal) difficulty of rejecting previously held moral beliefs and adopting new ones can be much greater than that of rejecting previously held scientific beliefs in

---

<sup>16</sup>  $t(143) = 6.4, p < .001$

<sup>17</sup>  $ts(144) = 10.9$  and  $11.2, ps < .001$

<sup>18</sup>  $ts(144) = 3.5$  and  $2.7, ps = .001$  and  $.007$ , respectively

favor of new ones. After all, our moral beliefs occupy the very center of what we take ourselves to value and who we take ourselves to be. Thus this sort of transition—from one set of moral beliefs to another—can require a major social/political/cultural (and personal) transformation. Such transformations cannot happen overnight—indeed, if pushed to hard, they can result in significant psychological distress and socio-political conflict.

The view that I am arguing for here is that meta-ethical pluralism helps to facilitate these psycho-social transformations, providing a means by which issues can move through normative space. On the one hand, it allows people to acknowledge the moral significance of an action, event, or situation without yet requiring (or even being allowed to require) the censorship and prohibition that typically accompanies moral issues. It “holds open” the social space for people to carefully consider an issue, to engage in social debate and dialogue—to achieve a sort of reflective equilibrium—before granting the issue its full (objective) moral weight. On the other hand, it is also allows people to continue to hold onto their moral convictions about an issue, while at the same time creating room for discussion and alternative choices/practices—with the hope that at some point the issue will shift out of the moral domain altogether.

Consider, from our own past, the social/political/cultural transformation required to embrace the fact that the ownership of other human beings is not a legitimate means to economic gain (one that individuals can choose to engage in or not), nor merely a social custom—rather, it is an inhumane and morally intolerable practice. Living in the south (indeed, in the city where over half of the soon-to-be slaves arrived in the US to be sold), I have witnessed firsthand the historical dialogue that treated slavery as a topic for social debate and discussion, where people of influence openly and publically argued both sides.<sup>19</sup> But, this is no longer the case—indeed, it

---

<sup>19</sup> An example of which was brought vividly to life by Spielberg’s recent cinematic depiction of the adoption of the 13<sup>th</sup> amendment in *Lincoln*.

would be hard to imagine an open public forum on the topic being actually considered, much less allowed (unless only as a historical lens). So, we might argue that this issue has successfully moved through normative space, being first recognized and discussed as having important moral significance—and yet, something that people had the freedom to choose, based on the pangs of their own moral consciences—to finally (via a combination of both reflective dialogue, rhetoric, and force) being granted its full moral weight and, thereby, forbidden.<sup>20</sup>

On the flip side, consider the transformation required to acknowledge that a behavior previously prohibited as morally repugnant and punishable—e.g., mixed-race relationships—is in fact a behavior of no more moral significance than any of the other relationship choices we make. Today, mixed-race relationships are a common (and widely accepted) sight on the streets of the city in which I live—something that once was considered by many to be an outrage, even a “lynch-able” offense.

Of course, even if this is correct, unless a reasonable philosophical defense of pluralism can be given, objectivists such as myself must acknowledge that an “error theory” of some sort (albeit, quite different from Mackie’s 1977) may be warranted after all—namely, that people’s moral discourse erroneously treats certain moral issues as non-objectively grounded and/or non-cognitive. But, as mentioned earlier, my account provides at least a *pragmatic* defense to this practice, something akin to a form of “exculpatory pretense”, such as was introduced to explain other linguistic practices (Turri 2012). That is, perhaps people merely talk (not necessarily consciously) *as if* a moral issue is non-objectively grounded so as to create/maintain space for choice, discussion, and debate—long enough, at least, for agreement between reasonable parties

---

<sup>20</sup> This is not to say that it is still not globally practiced—sadly, perhaps to an even greater degree than ever before.

to be reached, at which point the issue's actual objective status is acknowledged or the issue's moral status is abandoned.<sup>21</sup>

We must not neglect the non-objectivists/non-cognitivists, however. After all, the presence of meta-ethical pluralism can be accounted for within those frameworks as well, perhaps more easily so. It seems less problematic that people would speak non-objectively/non-cognitively about some moral issues, while speaking objectively/cognitively about others. After all, a number of the non-objectivist/non-cognitivist positions on offer (e.g., Blackburn 1984; Gibbard 1990; Timmons 1999) go a long way towards accounting for the cognitivist/objectivist-sounding language people use to discuss moral issues.

Either way, one thing is clear—the “descriptive” story about people's meta-ethical commitments is not nearly as uniform as was once supposed. And if people's linguistic dispositions in any way reflect deeper meta-ethical truths, then this complexity presents both philosophers and social scientists with important “food for thought”.

---

<sup>21</sup> This idea receives some support from the relationship between perceived consensus and objectivity found by Wright, et al. (2014) and Goodwin and Darley (2010, 2012). The former found that perceived consensus (both within one's own and across different cultures) mediated the relationship between people's meta-ethical commitments and their tolerance for divergent moral beliefs/values, while the latter found that manipulations of perceived consensus resulted in corresponding changes in objectivity—the greater the consensus, the greater the objectivity reported. These findings suggest a tight (and most likely bi-directional) relationship between the two variables.

## References

- Blackburn S (1984) *Spreading the Word: Groundings in the Philosophy of Language*. New York: Oxford University Press
- Brink D O (1989) *Moral Realism and the Foundations of Ethics*. New York: Cambridge University Press
- Darwall S (1998) *Philosophical Ethics*. Boulder, CO: Westview Press
- Dreier J (1999) Transforming expressivism, *Noûs*, 33, 558-572
- Gibbard A (1990) *Wise choices, apt feelings: A theory of normative judgment*. Cambridge: Harvard University Press
- Gill M (2008) Indeterminacy and Variability in Meta-ethics, *Philosophical Studies*, 215-234.
- Goodwin G, Darley J (2008) The psychology of meta-ethics: Exploring objectivism. *Cognition*, 106, 1339-1366
- Goodwin G, Darley J (2010) The perceived objectivity of ethical beliefs: Psychological findings and implications for public policy. *Review of Philosophy and Psychology*, 1, 1-28
- Goodwin G, Darley J (2012) Why are some moral beliefs perceived to be more objective than others? *Journal of Experimental Social Psychology*, 48, 250-256
- Haidt J (2012) The righteous mind: Why good people are divided by politics and religion
- Harman G (1975) Moral relativism defended. *The Philosophical Review*, 84, 3-22
- Harman G (1996) Moral relativism. In G. Harman & J. J. Thomson (Eds.), *Moral relativism and moral objectivity* (pp. 3-19). Oxford: Blackwell
- Joyce R (2009) Moral Anti-Realism, *The Stanford Encyclopedia of Philosophy* (Summer 2009 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/sum2009/entries/moral-anti-realism>.

- Loeb D (2008) Moral incoherentism: How to pull a metaphysical rabbit out of a semantic hat. In W Sinnott-Armstrong (Ed.) *Moral psychology: The cognitive science of morality: intuition and diversity* (Vol. 2; pp. 355-385). Cambridge: Bradford Books
- Mackie J L (1977) *Ethics: Inventing Right and Wrong*. New York: Penguin
- Nichols S (2004) After objectivity: An empirical study of moral judgment. *Philosophical Psychology*, 17, 3-26
- Rachels J, Rachels S (2009) *The Elements of Moral Philosophy*. New York: McGraw Hill
- Sarkissian H, Park J, Tien D, Wright J C, Knobe J (2011) Folk moral relativism. *Mind & Language*, 26:4, 482-505
- Shafer-Landau R (2003) *Moral Realism: A Defence*. New York: Oxford University Press
- Sinnott-Armstrong W (1993) Some Problems for Gibbard's Norm-Expressivism, *Philosophical Studies* 69, 297–313.
- Sinnott-Armstrong W (2009) Mixed-up meta-ethics. *Philosophical Issues*, 19, 235-256.
- Smith M (1994) *The Moral Problem*. Oxford: Blackwell
- Timmons M (1999) *Morality without foundations*. New York: Oxford University Press
- Turri (2012) The test of truth: An experimental investigation of the norm of assertion, manuscript under review
- Wong D B (1984) *Moral Relativity*. Berkeley: University of California Press
- Wong D B (2006) *Natural Moralities: A Defence of Pluralistic Relativism*. New York: Oxford University Press
- Wright J C (2011) Children's and adolescents' tolerance for divergent beliefs: Exploring the cognitive and affective dimensions of moral conviction in our youth. *British Journal of Developmental Psychology*, online version, 1-18. DOI:10.1111/j.2044-835X.2011.02058.x

- Wright J C, Cullum J, Schwab N (2008) The cognitive and affective dimensions of moral conviction: Implications for tolerance and interpersonal behaviors. *Personality and Social Psychology Bulletin*, 34:11, 1461-1476
- Wright J C, Grandjean P T, McWhite C (2012) The meta-ethical grounding of our moral beliefs: Evidence for meta-ethical pluralism. *Philosophical Psychology*, 336-361.
- Wright, J C, McWhite C, Grandjean P (2014) The cognitive mechanisms of intolerance: Do our meta-ethical commitments matter? In T Lombrozo, S Nichols, J Knobe (Eds). *Oxford Studies in Experimental Philosophy* (Vol. 1)

## Tables

	Personal	Social	Moral
<b>selling children on internet</b>	7%	8%	<b>84%</b>
<b>eating part of another human being</b>	10%	13%	<b>77%</b>
<b>steal money and/or supplies from the large company where you work</b>	20%	9%	<b>71%</b>
<b>conscious discrimination on basis of race/gender</b>	10%	20%	<b>70%</b>
<b>having sex with someone other than spouse</b>	25%	6%	<b>68%</b>
<b>refusing to provide help to those who need it</b>	<b>40%</b>	8%	<b>52%</b>
<b>helping terminally ill patients</b>	<b>43%</b>	7%	<b>50%</b>
<b>engaging in prostitution</b>	<b>39%</b>	17%	<b>44%</b>
<b>eating your pets (that died from an accident)</b>	31%	25%	<b>44%</b>
<b>terminating pregnancy</b>	<b>56%</b>	3%	<b>40%</b>
burning the American flag	20%	52%	28%
eating factory-farmed meat	76%	7%	17%
watching pornographic videos	80%	7%	14%
using recreational drugs	75%	13%	12%
smoke cigarettes in enclosed public space	37%	52%	11%
playing violent video games	73%	16%	11%
using swear words in public	50%	41%	9%
publicly criticize your government and engage in protest events	61%	33%	6%
wearing pajamas to business meeting	40%	56%	4%
getting tattoos/body piercings	92%	6%	2%

*Table 1.* Domain Classification

	Obj	Subj	Non-Rel	Rel	Cat	Non-Cat
<b>selling children on internet</b>	<b>92%</b>	8%	<b>86%</b>	14%	<b>75%</b>	25%
<b>eating part of another human being</b>	<b>79%</b>	21%	<b>68%</b>	32%	<b>69%</b>	31%
<b>steal money and/or supplies from the large company where you work</b>	<b>76%</b>	24%	<b>70%</b>	30%	45%	55%
<b>conscious discrimination on basis of race/gender</b>	<b>80%</b>	20%	<b>69%</b>	31%	<b>70%</b>	30%
<b>having sex with someone other than spouse</b>	<b>63%</b>	37%	56%	44%	<b>61%</b>	39%
<b>refusing to provide help to those who need it</b>	39%	<b>61%</b>	24%	<b>76%</b>	35%	<b>65%</b>
<b>helping terminally ill patients</b>	33%	<b>67%</b>	24%	<b>76%</b>	16%	<b>84%</b>
<b>engaging in prostitution</b>	51%	49%	37%	<b>63%</b>	43%	57%
<b>eating your pets (that died from an accident)</b>	53%	47%	37%	<b>63%</b>	45%	55%
<b>terminating pregnancy</b>	34%	<b>66%</b>	18%	<b>82%</b>	23%	<b>77%</b>
burning the American flag	50%	50%	42%	<b>58%</b>	38%	<b>62%</b>
eating factory-farmed meat	32%	<b>68%</b>	14%	<b>86%</b>	17%	<b>83%</b>
watching pornographic videos	29%	<b>71%</b>	18%	<b>82%</b>	12%	<b>88%</b>
using recreational drugs	36%	<b>64%</b>	27%	<b>73%</b>	23%	<b>77%</b>
smoke cigarettes in enclosed public space	28%	<b>72%</b>	21%	<b>79%</b>	30%	<b>70%</b>
playing violent video games	29%	<b>71%</b>	12%	<b>88%</b>	10%	<b>90%</b>
using swear words in public	19%	<b>81%</b>	11%	<b>89%</b>	5%	<b>95%</b>
publicly criticize your government and engage in protest events	27%	<b>73%</b>	16%	<b>84%</b>	10%	<b>90%</b>
wearing pajamas to business meeting	40%	<b>60%</b>	26%	<b>74%</b>	30%	<b>70%</b>
getting tattoos/body piercings	34%	<b>66%</b>	11%	<b>89%</b>	2%	<b>98%</b>

Table 2. Item-Level Variability

	Fully Non-Objective	Pluralist	Fully Objective
Subjectivity	3%	78%	19%
Reasons	6%	82%	12%
Relativity	7%	85%	8%

	Fully Non-Cognitivist	Pluralist	Fully Cognitivist
Semantic Non-Factualism	15%	76%	9%
Psychological Non-Cognitivism	16%	75%	9%

*Table 3. Intrapersonal Variability*

Not Truth-Apt/Expression of Feelings/Attitudes	% correct	
Golden Retrievers are better dogs than Chihuahuas.	89%	88%
Heavy Metal music sucks!	88%	93%
Strawberries are tastier than raspberries.	87%	91%
Abstract art is a waste of time and space.	89%	92%
Walking on the beach at sunset is relaxing.	80%	83%
Truth-Apt/Assertion of Beliefs	% correct	
Penguins are birds that can't fly.	97%	97%
Golden Retrievers are bigger dogs than Chihuahuas.	94%	98%
Water is H <sub>2</sub> O.	97%	100%
Triangles are sturdier for construction (hold more weight) than squares.	87%	86%
Benjamin Franklin was the third president of the United States.	94%	95%

*Table 4.* Introductory Exercise Examples

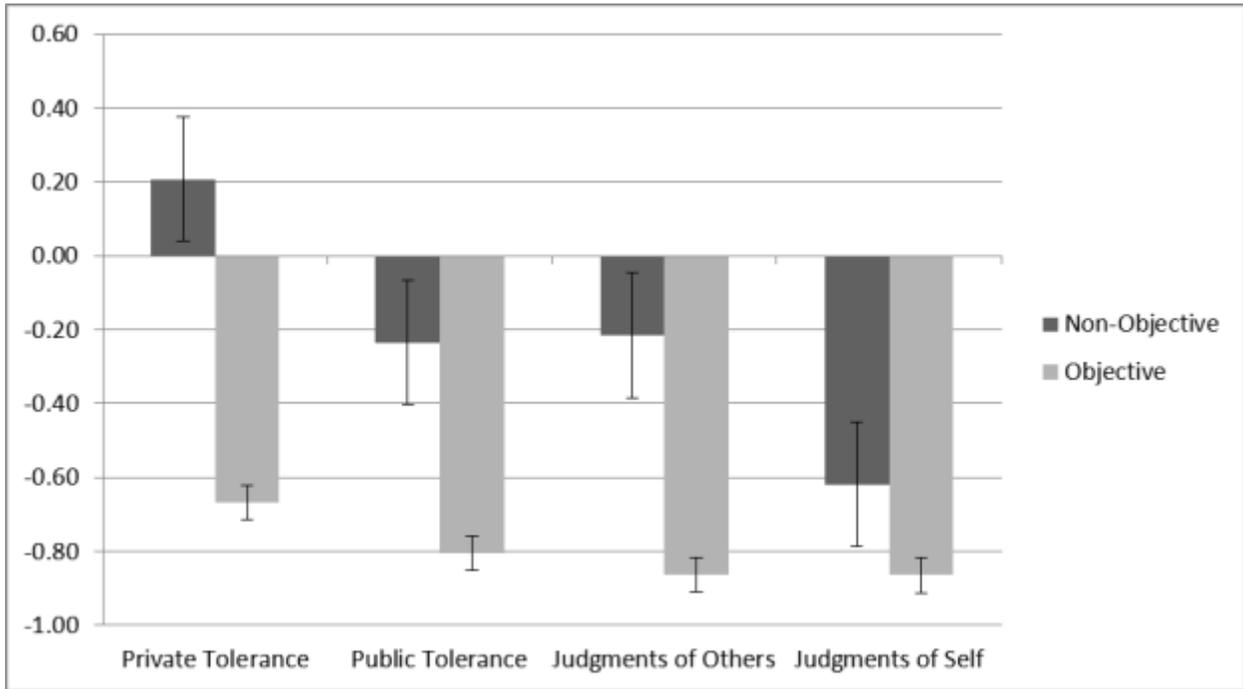
	Personal	Social	Moral	Personal	Social	Moral
Selling children on the internet.	3%	5%	<b>91%</b>	6%	6%	<b>89%</b>
Taking things that don't belong to you.	7%	10%	<b>83%</b>	12%	11%	<b>76%</b>
Forcing someone else to have sex.	11%	6%	<b>83%</b>	14%	10%	<b>76%</b>
Cheating on an exam.	14%	5%	<b>80%</b>	15%	10%	<b>75%</b>
Cheating on one's spouse.	11%	11%	<b>78%</b>	18%	9%	<b>74%</b>
Knowingly overcharging someone for a product or service in order to make more money.	20%	11%	<b>70%</b>	17%	13%	<b>70%</b>
Consciously discriminating against someone on the basis of race.	13%	18%	<b>68%</b>	13%	26%	<b>61%</b>
Having consensual sex with another close family member.	13%	20%	<b>67%</b>	22%	17%	<b>61%</b>
Doing nothing in the face of people in our communities that could use our help.	<b>48%</b>	5%	<b>47%</b>	<b>37%</b>	15%	<b>48%</b>
Buying products made overseas in sweat shops.	<b>43%</b>	12%	<b>46%</b>	<b>44%</b>	12%	<b>44%</b>
Engaging in vigilante justice (take the law into your own hands) when you or someone you love has suffered an injustice.	<b>46%</b>	18%	<b>36%</b>	<b>50%</b>	17%	<b>33%</b>
Getting a 1st trimester abortion.	<b>63%</b>	1%	<b>36%</b>	<b>70%</b>	4%	26%
Burning the American flag in protest of our governmental policies.	<b>51%</b>	24%	25%	<b>51%</b>	27%	22%
Watching pornographic videos.	<b>72%</b>	4%	24%	<b>76%</b>	7%	18%
Using illegal drugs.	<b>61%</b>	16%	23%	<b>66%</b>	16%	17%
Eating factory-farmed (instead of free range) cows, pigs, and chickens.	<b>80%</b>	6%	14%	<b>61%</b>	22%	17%
Smoke cigarettes openly in enclosed public places.	<b>57%</b>	30%	13%	<b>84%</b>	3%	13%
Recycling paper, plastic, cardboard, and metal.	<b>66%</b>	24%	11%	<b>52%</b>	36%	12%
Gay couples getting married.	<b>72%</b>	18%	10%	<b>59%</b>	32%	9%
Privately owning guns or other dangerous weapons.	<b>83%</b>	12%	5%	<b>82%</b>	11%	7%

Table 5. Domain Classification

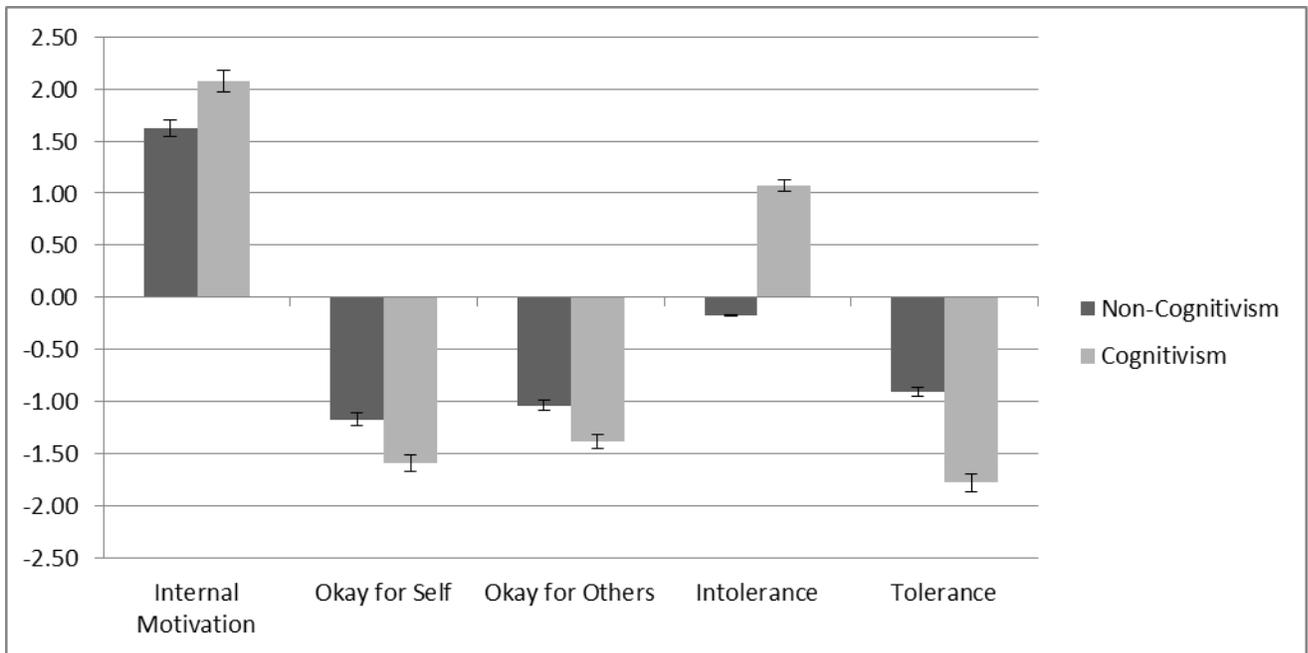
	Moral Only				Moral Only			
	Non-Objective	Objective	Non-Objective	Objective	Express Feeling/Attitude	Assert Belief/Both	Express Feeling/Attitude	Assert Belief/Both
Selling children on the internet.	27%	<b>73%</b>	27%	<b>73%</b>	5%	<b>95%</b>	6%	<b>94%</b>
Taking things that don't belong to you.	48%	52%	41%	59%	7%	<b>93%</b>	7%	<b>93%</b>
Forcing someone else to have sex.	26%	<b>74%</b>	23%	<b>77%</b>	22%	<b>78%</b>	16%	<b>84%</b>
Cheating on an exam.	48%	52%	49%	51%	<b>47%</b>	53%	<b>45%</b>	55%
Cheating on one's spouse.	58%	42%	57%	43%	38%	<b>62%</b>	36%	<b>64%</b>
Knowingly overcharging someone for a product or service in order to make more money.	59%	41%	50%	50%	35%	<b>65%</b>	32%	<b>68%</b>
Consciously discriminating against someone on the basis of race.	47%	53%	32%	<b>68%</b>	29%	<b>71%</b>	22%	<b>78%</b>
Having consensual sex with another close family member.	60%	40%	50%	50%	37%	<b>63%</b>	24%	<b>76%</b>
Doing nothing in the face of people in our communities that could use our help.	<b>89%</b>	11%	<b>80%</b>	20%	<b>62%</b>	38%	<b>56%</b>	44%
Buying products made overseas in sweat shops.	<b>78%</b>	22%	<b>67%</b>	33%	<b>57%</b>	43%	<b>44%</b>	56%
Engaging in vigilante justice (take the law into your own hands) when you or someone you love has suffered an injustice.	<b>78%</b>	22%	<b>79%</b>	21%	<b>53%</b>	47%	<b>50%</b>	50%
Getting a 1st trimester abortion.	<b>83%</b>	17%	<b>67%</b>	33%	<b>67%</b>	33%	<b>57%</b>	43%
Burning the American flag in protest of our governmental policies.	<b>74%</b>	26%	48%	52%	<b>58%</b>	42%	35%	<b>65%</b>
Watching pornographic videos.	<b>92%</b>	8%	<b>73%</b>	27%	<b>71%</b>	29%	31%	<b>69%</b>
Using illegal drugs.	<b>69%</b>	31%	52%	48%	<b>67%</b>	33%	38%	<b>63%</b>
Eating factory-farmed (instead of free range) cows, pigs, and chickens.	<b>93%</b>	7%	<b>77%</b>	23%	<b>45%</b>	55%	0%	<b>100%</b>
Smoke cigarettes openly in enclosed public places.	<b>80%</b>	20%	50%	50%	<b>82%</b>	18%	<b>55%</b>	45%
Recycling paper, plastic, cardboard, and metal.	<b>77%</b>	23%	40%	60%	<b>54%</b>	46%	36%	<b>64%</b>
Gay couples getting married.	<b>77%</b>	23%	56%	44%	<b>52%</b>	48%	25%	<b>75%</b>
Privately owning guns or other dangerous weapons.	<b>86%</b>	14%	<b>100%</b>	0%	<b>68%</b>	32%	33%	<b>67%</b>

Table 6. Item-Level Variability

## Figures



*Figure 1.* People's tolerance for divergent moral actions: Objective vs. Non-Objective



*Figure 2.* People's tolerance for divergent moral actions: Cognitivism vs. Non-Cognitivism