

In the Thick of Moral Motivation

Wesley Buckwalter¹ · John Turri^{1,2}

© Springer Science+Business Media Dordrecht 2016

Abstract We accomplish three things in this paper. First, we provide evidence that the motivational internalism/externalism debate in moral psychology could be a false dichotomy born of ambiguity. Second, we provide further evidence for a crucial distinction between two different categories of belief in folk psychology: thick belief and thin belief. Third, we demonstrate how careful attention to deep features of folk psychology can help diagnose and defuse seemingly intractable philosophical disagreement in metaethics.

1 The Concept of Moral Judgment

Michael works for minimum wage at a local sandwich shop. The job doesn't pay well — it's minimum wage. But Michael takes pride in a job well done. He conscientiously completes every task assigned and cheerfully serves every customer. Most importantly, Michael believes that he ought to be honest to his employer. In particular, he believes that he ought to tell the truth about how many hours he works.

It's completely natural to expect Michael's moral judgment to be accompanied by motivation to tell the truth. Of course, we don't necessarily expect the motivation to be infeasible; other factors might undermine or override it. But because this expectation is so elementary and utterly natural, it's easy to overlook its centrality in commonsense psychology.

To see this, just imagine that an observer learns that Michael believes he ought to tell the truth and comes to the opposite conclusion. The observer expects that Michael will be motivated to lie and that, other things being equal, he will lie. Needless to say, such

✉ Wesley Buckwalter
wesleybuckwalter@gmail.com

John Turri
john.turri@gmail.com

¹ Department of Philosophy, University of Waterloo, 200 University Avenue West, Waterloo, ON N2L 3G1, Canada

² Cognitive Science Program, University of Waterloo, Waterloo, ON, Canada

an expectation is exceedingly puzzling. We'd double-check to make sure the story was heard correctly. After convincing ourselves that it hadn't simply been misheard, we'd conclude that this expectation is simply confused.

And what about observers who, given Michael's moral judgment, form no expectations about Michael's motivation? It's not that they expect him to act contrary to his moral judgment. Rather, it just doesn't lead them to expect anything. Such observers aren't seriously confused. But there seems to be something important missing from their repertoire of mindreading skills. They aren't fully competent in interpreting other agents or applying the concept of moral judgment.

This raises an interesting question. Why might full competence in this respect require us to expect motivation to accompany moral judgment? A venerable theory in moral psychology, known as motivational judgment internalism ('internalism' for short), offers one answer (Williams 1937; Falk 1947). Full competence requires this because it is a conceptual necessity that motivation accompanies moral judgment. On this view, it's just part of our concept of moral judgment that motivation accompanies it.

Fast forward a few years. Michael is completely frustrated with his job at the sandwich shop. Petulant customers disrespect him; the boss takes him for granted; and he is unfairly passed over for promotion several times. Michael still remembers his past and all the reasons why he should tell the truth about how many hours he works. He still thinks that he ought to tell the truth, but he no longer has any motivation to do so. He's not the least bit dissuaded from lying about his hours.

If motivational internalism is correct, then the previous paragraph describes an impossible situation: if Michael has no motivation to tell the truth, then he doesn't sincerely believe that he ought to tell the truth. But is it really impossible? Many find it difficult to accept that it's impossible. Instead, they have the intuition that the situation is possible.

One response would be to weaken the internalist thesis while retaining some conceptual connection between moral judgment and motivation. Perhaps the motivation entailed by sincere moral judgment only holds for some moral judgments. For example, perhaps it's a conceptual necessity that moral judgment motivates in psychologically normal people (Dreier 1990; Blackburn 1998), or in rational people, or in morally virtuous people (Smith 1994). Or perhaps it's a conceptual necessity that a linguistic community couldn't have the concept of moral judgment if they didn't also tend to ascribe motivation in tandem with the ascription of moral belief (compare Dreier 2006). Opponents argue that all of these weaker forms of internalism face problems (Copp 1997; Miller 2003; Roskies 2003; Sayre-McCord 1997).¹

But for present purposes the important point is that, one way or another, to a greater or lesser extent, weakened versions of internalism try to preserve a conceptual connection between moral judgment and motivation. Philosophers who advocate these views are trying to identify what they take to be the kernel of conceptual truth in internalism, while honoring the intuition that we must make room for cases like Michael's (see

¹ For example, it's not clear that we made Michael abnormal or irrational in the second half of the story. Rather, as Stocker put the point, such a "course of life" seems "typical" (1979, p. 741).

Björklund et al. 2012). And make no mistake about it: lots of very smart people are convinced that some such kernel exists.

A completely different response is to break with internalism entirely and accept the opposite view, motivational judgment externalism ('externalism' for short). Externalists propose that any connection between moral judgment and motivation is contingent (Stocker 1979; Mele 1996; Brink 1997; Svavarsdottir 1999). Nothing in the concept of moral judgment entails corresponding motivation. It needn't betray conceptual incompetence or confusion to attribute a moral belief while explicitly abstaining from attributing corresponding motivation (Zangwill 2008). Indeed, just the opposite in some cases. One reason why the crimes of some psychopaths, such as Ted Bundy, strike as so monstrous is that they know that what they're doing is wrong, but this moral knowledge doesn't deter them in the least (Roskies 2003). Their moral knowledge provides them with no motivation — and that's scary.² But since knowledge entails belief, if they know that they ought not to do such things, then they must also believe that they ought not to do such things. In drawing this conclusion, it seems doubtful that we betray any conceptual confusion. And, once again, lots of very smart people are convinced that they're making no such mistake about what the concept of moral judgment requires.

So which is it? Is there a conceptual connection between moral judgment and motivation, or isn't there?

2 An Impasse?

As several philosophers have observed, the debate between internalists and externalists continues with little end in sight. For example, "At this point in the dialectic," writes Connie Rosati, "internalists and externalists tend to produce additional cases and probe our intuitions further in an effort to overcome what seems an impasse" (2006). Given that we seem to have reached an impasse, how best to proceed?

David Hume once argued that ambiguity best explains persistent disagreement among parties to a longstanding debate.

From this circumstance alone, that a controversy has been long kept on foot, and remains still undecided, we may presume that there is some ambiguity in the expression, and that the disputants affix different ideas to the terms employed in the controversy. (Hume 1748: section 8.1)

The debate between internalists and externalists has been long kept on foot and remains still undecided. We think it wise to follow Hume's counsel, "It may, therefore, be worthwhile to consider, what is real, and what is only verbal, in this controversy" (Hume 1772, p. 72).

But uncovering a verbal dispute is only incidental to our primary goal. Our primary goal is to illuminate a deep and important truth underlying the verbal dispute. Philosophers were wrong to assume that there was any special connection between moral judgment and motivation. Motivational internalism about moral judgment

² For empirical work on moral judgments made by psychopaths (or people with so-called "acquired sociopathy") see Roskies (2006, 2008). For empirical work on ordinary judgments about moral motivation about psychopaths see Nichols (2002, 2004).

merely reflects a more general phenomenon with deep origins in the folk psychology of belief. Specifically, our primary goal is to test whether key intuitions involved in this debate could be due to two separate conceptions of belief. Just as many have suspected, when understood according to one concept, motivational internalism expresses a conceptual truth. It was a mistake to weaken internalism (much). Nevertheless, when understood according to the other concept of belief, motivational externalism is also true. The upshot is that the professional philosophical debate between internalists and externalists could be, just as Hume would have predicted, largely verbal.

The remainder of this paper makes good on those bold claims. Section 3 reviews relevant recent work from the epistemology of belief. This work locates within folk psychology a distinction between thick belief and thin belief. Section 4 shows that these two categories can be applied to explain crucial intuitions in the debate between motivational internalists and externalists. Section 5 charts some of the implications of this finding for several foundational issues in metaethics.

3 The Folk Psychology of Belief Meets Moral Motivation

There is considerable disagreement in contemporary epistemology over the nature of knowledge, perhaps more so than among moral psychologists about the nature of moral judgment. But there is virtual consensus among epistemologists on this point: knowledge requires true belief. That is, necessarily, if you know that P, then you believe that P (see Armstrong 1969, p. 21). This is known as the entailment thesis. If asked why one might accept the entailment thesis, the answer is that it's obvious and there are no counterexamples.

And then something interesting happened. A pair of experimental philosophers remembered that there had been controversy over the entailment thesis. Back in the 1960s and 70s, some philosophers presented what they took to be counterexamples to the entailment thesis (for debate see Radford 1966, p. 5; and Armstrong 1969; Cohen 1966; Jones 1971; Lehrer 1974). No one suspected these contrarians of conceptual incompetence. But their contrarian concerns were brushed off as confusion of some sort. Professional epistemologists returned to the business of discovering what needed to be added to true belief to get knowledge. And that was the end of that, until the pair of experimental philosophers — Blake Myers-Schulz and Eric Schwitzgebel — did what any politician in a democracy would do: they took their case to the people (2013).

Their results were shocking. Many people read the case and answered that the protagonist in the story knew the proposition in question (call it 'Q'). But many more people read the case and answered that the protagonist did not believe Q. The end result was that people were far more likely to classify it as a case of knowledge than as a case of belief. The result was replicated using a range of different stories. It began to look like the entailment thesis was very far from being an obvious conceptual truth about knowledge. Instead, it began to look like, on the ordinary concept of knowledge, belief was not required.

This unsurprisingly led to further investigation. Upon further investigation, another pair of experimental philosophers — David Rose and Jonathan Schaffer — argued that the initial studies failed to account for a distinction between occurrent and dispositional belief. This second round of studies did account for the occurrent/dispositional

distinction. The studies yielded results that were fully in line with the orthodox position on the entailment thesis: more people ascribed belief than ascribed knowledge (Rose and Schaffer 2013).

But then a third round of studies was conducted by yet another group of researchers (Murray et al. 2013). This third round of studies also accounted for the occurrent/dispositional distinction but found results generally consistent with those of Myers-Schulz and Schwitzgebel. They once again found that, across a wide range of cases, people ascribed knowledge at significantly higher rates than they ascribed belief.

Here is a representative case from the third round of studies:

[GEOCENTRIST] Karen is a first-year student at a prestigious university. She is a good student and has been doing very well in her classes. One of the classes she is taking is introduction to physics. One of the topics covered in this class is the place of the earth in the solar system. For example, one of the things that Karen has been taught is that the earth revolves around the sun. Prior to starting at the university, however, Karen was home-schooled by her parents. Karen's parents taught her that the earth is at the center of the universe. Karen accepts what her parents have taught her. In particular, despite what she has been taught in her physics class, she holds that the earth does not revolve around the sun. One of the questions on the final exam in Karen's physics class is the following: "True or false: The earth revolves around the sun." Karen answers "true" on this question. She gets the question correct and ends up scoring 100 % on the exam. (Murray et al. 2013, p. 94)

Nearly 50 % answered that Karen knows but does not believe that the earth revolves around the sun. Of those who ascribed knowledge, the vast majority (85 %) did not ascribe belief. This is exactly the opposite of what we would expect if participants accepted the entailment thesis.

What's going on here? Unless the majority of English-speaking adults are guilty of gross conceptual incompetence, it appears that it's not a conceptual truth that knowledge requires belief. But then how could the vast majority of epistemologists and cognitive scientists be so wrong about the concept of knowledge vis-a-vis the entailment thesis?

Further research revealed that the disagreement was merely apparent and had its roots in a deep feature of the folk psychology of belief (Buckwalter et al. 2015). There are at least two categories within folk psychology that are expressed by the term 'belief'. One of these categories — thin belief — is implicitly assumed in the contemporary professional literature on knowledge and belief. The other category — thick belief — enjoys wide, but not exclusive, currency in ordinary life.

A thin belief is a bare cognitive pro-attitude. To have a thin belief that P, it suffices that you represent that P is true, regard it as true, or take it to be true. Put another way, thinly believing P involves representing and storing P as information. It requires nothing more. In particular, it doesn't require you to like it that P is true, to emotionally endorse the truth of P, to explicitly avow or assent to the truth of P, or to actively promote an agenda that makes sense given P.

A thick belief requires more than a bare cognitive pro-attitude. As a rough first approximation, thick belief also involves emotion or conation. There might be many ways to thicken belief beyond a bare cognitive pro-attitude. For example, in addition to

representing and storing P as information, you might also like it that P is true, emotionally endorse the truth of P, explicitly avow or assent to the truth of P, or actively promote an agenda that makes sense given P. Thick belief entails thin belief, but not vice versa.

The thin category of belief dominates contemporary philosophy of mind and cognitive science. For example, Fred Dretske characterizes belief merely as an “inner representation of the world,” something “represented . . . in the internal language of thought” (Dretske 1983: 14). Moreover, this view of belief as merely “saying something in the internal language of thought” now “dominates philosophical thinking” (Dretske 1983: 3; see also Fodor 1990).

But the thick sense makes important appearances in the history of philosophy (Hume 1748.; James 1889; Peirce 1958). For example, in the *Enquiry Concerning Human Understanding*, Hume asks, what distinguishes mere imagination from belief? (All Hume quotes are from Hume 1748: §5.2.) The distinction can’t be in terms of representational content, Hume argues, because whatever we believe we can equally well imagine. Neither is the distinguishing feature one of storing, or being disposed to store, the content in memory. Rather, on Hume’s view, a distinctive “feeling” distinguishes belief from imagination. Belief is a “sentiment” with an “influence on the passions,” a “governing principle of our actions.” On Hume’s view, belief ordinarily understood requires more than thin belief. On Hume’s view, belief ordinarily understood is thick belief — this is what we mean by belief “in common life.”³

With this distinction between thin and thick belief in mind, return now to the Geocentrist example above. Researchers found that when they primed people to answer in terms of thick belief, they were able to replicate the earlier results that threatened the entailment thesis: many more people ascribed knowledge than belief. However, researchers also found that when they primed people to answer in terms of thin belief, they got dramatically different results: many more people ascribed belief than knowledge.

In order to prime people to answer in terms of thick belief, researchers asked whether Karen “believes” the proposition in question. In order to prime them to answer in terms of thin belief, researchers asked whether Karen “thinks” that the proposition is true. In order to allow for a reading that includes dispositional thin belief, researchers also asked, “At least on some level, does Karen think” that the proposition is true?⁴

So, in the end, it turns out that the professionals were right when they proposed the entailment thesis as a conceptual truth. But unless we’re willing to conclude that most people lack basic competence in applying the concept of knowledge, the entailment thesis must be understood in terms of thin belief.

We submit that, in light of these results, any program of philosophical or psychological research on “belief” should take into account the difference between thick belief and thin belief. Of course, the results don’t show that thick and thin belief share no

³ Gendler (2008a, b) draws an orthogonal distinction between ‘belief’ and ‘alief’. Further work is required to determine whether this interacts with our classification of thick and thin belief. For how the alief/belief distinction might bear on moral motivation and moral phenomenology, see Kriegel (2012). For yet another theoretical account of the different psychological states that might be implicated in the debate between internalists and externalists see Kauppinen (2015).

⁴ Buckwalter et al. (2015) replicate this effect for thick and thin belief across several experiments using a number of different cover-stories.

interesting features in common. All beliefs, thick and thin, share important core properties. But it's beyond doubt that they also differ in important ways too. When intuitions are split on whether "belief" must be a certain way, we should immediately ask whether one set of intuitions is keyed to thick belief, while the other set is keyed to thin belief.⁵

The careful reader will already have anticipated what this might mean for the debate between motivational internalists and externalists. Prior to adding one more iteration to that debate, we should investigate whether internalist intuitions are correctly keyed to thick belief, while externalists intuitions are correctly keyed to thin belief. If internalists and externalists are appealing to two distinct kinds of psychological states, then their disagreement might be largely verbal.⁶ We investigate this possibility in the next section.

4 Experiments in Moral Motivation

This section reports four experimental studies we have conducted examining thin and thick belief ascription in the absence of moral motivation. Experiment 1 begins by adapting what is perhaps the most influential thought experiment to date in the literature on motivational internalism and externalism. We find that people view the case differently depending on how they are questioned. In particular, they view the case differently when they are questioned in terms of thick belief rather than thin belief. Experiment 2 replicates those results and shows that the effect persists even after taking into account two potential objections. Experiment 3 once again replicates the basic finding and extends it to cases involving non-moral evaluative judgments. Experiment 4 rules out a possible objection to the stimuli used in prior experiments. Cumulatively, these results strongly suggest that (i) thick belief and thin belief have very different motivational profiles, and (ii) the motivational internalism/externalism debate simply evaporates in light of the distinction. We conclude Section 4 by reviewing a compatible finding from others in the recent experimental literature.

4.1 Stocker's Case: Experiment 1

To ensure that our discussion connects directly to the issues at the heart of the internalism/externalism debate, our first experiment uses vignettes inspired by Michael Stocker's renowned case of the jaded politician. We begin by adapting Stocker's politician thought experiment as a test case partly because of its influence on the moral motivation debate, and partly because we suspect it will prove among the most challenging cases to detect differences between thin and thick belief ascription (for reasons that will become clear in Section 4.3).

⁵ For a similar and insightful conceptual move in the debate about free will see May (2014).

⁶ Theorizing about the nature of verbal disputes must be saved for another occasion. See Chalmers (2011, p. 522) for one general working definition that "a dispute over S is (broadly) verbal when for some expression T in S, the parties disagree about the meaning of T, and the dispute over S arises wholly in virtue of this disagreement regarding T."

The Jaded Politician case convinced many people to either abandon or weaken internalism.⁷ Here is the case in Stocker's own words:

Recently, I read a story of what might be taken as typical of one course of life. It was said of this political figure that, in his youth, he cared a lot about the suffering of people in all parts of the world and devoted himself to making their lives better. But now he concerns himself only with the lives and fortunes of his close family and friends. He remembers his past, and he knows that there is still a lot he could do to help others. But he no longer has any desire so to do. (Stocker 1979, p. 741)

We presented participants ($N=164$, aged 18–66 years, mean age=30.49 years, 62 female)⁸ with the following vignette, very similar in structure to Stocker's original politician case above:

[JADED POLITICIAN] Michael has been a politician for decades. For years he cared deeply about poor people and worked tirelessly on anti-poverty legislation because it was the right thing to do. Then Michael slowly became completely focused on the welfare of his close friends and family. He remembers his past and all the reasons why he should help poor people. But he no longer has any motivation to help them. When new anti-poverty legislation is proposed, Michael does not support it.

All participants were then asked the following two comprehension and control questions:

Michael is a: [politician/aid worker/insurance salesman]

Michael _____ motivation to help poor people. [still has some/no longer has any]

Then participants were assigned to one of two conditions. Participants in the Thick Belief condition were asked:

Does Michael believe that he ought to help poor people? [Yes/No]

Alternatively, participants in the Thin Belief condition were asked the following question:

At least on some level, does Michael think that he ought to help poor people?
[Yes/No]

⁷ If we weren't joining an ongoing controversy in medias res, we would use different, simpler cases. Part of the reason, as we discuss below, is that certain details of Stocker's original politician case suggest that it is not a fair test of internalism.

⁸ As with the experiments reported below, participants were recruited using Amazon Mechanical Turk, tested online using Qualtrics survey software, and compensated \$0.15 for approximately 2–3 min of their time. Participants were not allowed to re-take any survey reported here, and participants who had taken previous similar surveys were excluded by their AMT Worker ID. Participants were located in the United States. Preliminary analysis revealed no effects of participant age or gender on the dependent variables, so the analyses that follow collapse across those demographic variables. The same is true in all the other experiments reported in this paper.

Upon answering one of the dichotomous belief questions, participants in both thick and thin conditions were asked to rate how confident they were in their answer. These responses were collected on a ten item scale, where “1” was anchored with “not at all confident” and “10” was anchored with “completely confident”.

We predicted that rates of thin-belief ascription will be significantly higher than rates of thick-belief ascription. Our prediction was true. The rate of belief ascription was higher in the Thin Belief condition (78 %) than in the Thick Belief condition (60 %), $\chi^2(1, n = 157) = 5.76, p < .05$, Cramer's $V = .192$ (all tests two-tailed, unless otherwise noted).⁹ Participants ascribed thin belief at rates that exceeded what could be expected by chance, binomial test, $p < .01$, test proportion = .5. By contrast, participants in the Thick Belief condition ascribed belief at rates that did not differ significantly from chance, binomial test, $p = .11$, test proportion = .5. These results are shown in Fig. 1.¹⁰

4.2 The Very Jaded and the Elite: Experiment 2

Experiment 1 provides some evidence for our thesis. However two objections immediately surface. The first worry is that although the vast majority of participants in Experiment 1 answered that “Michael no longer has any motivation to help poor people,” it might be objected that the explicit text of the story doesn't obviously entail that Michael lacks the relevant motivation. Perhaps some participants implicitly assumed that Michael would still help if it was easy and convenient enough for him.¹¹ In order to address this worry and inhibit such a reading of the case, we tested the following story, which makes Michael's utter lack of motivation even clearer:

[VERY JADED POLITICIAN] Michael has been a politician for decades. For years he cared deeply about poor people and worked tirelessly on anti-poverty legislation because it was the right thing to do. Then Michael slowly lost all interest in that set of issues and became completely focused on making money and accumulating a vast personal fortune at the cost of others. He remembers his past and all the reasons why he should help poor people. But he no longer has any motivation to help them. Now whenever new anti-poverty legislation is proposed, Michael does literally nothing to support it, even though he could easily offer some support at almost no cost to himself.

This addresses the first worry by making it explicitly clear that Michael does literally nothing to help poor people, even when it is easy and doing so would cost him next to nothing. As the saying goes, Michael would not lift a finger to help.

⁹ Seven participants were removed from the analysis for failing a comprehension question.

¹⁰ In order to employ more powerful statistical tests to detect patterns in belief ascription, we also created a belief-confidence measure. This measure took answers given to the dichotomous belief questions and combined them with scalar confidence judgments (e.g. “How confident are you in the answer you just gave to the previous question?”). Using this more sensitive measure, we detected the same pattern of results reported on dichotomous scales throughout the paper and, thus, omit these analyses for simplicity.

¹¹ Dreier (1990, p. 12) notes a similar point when he writes of Stocker's original politician case that being asked to imagine a person who “would not lift a finger, as it were, to achieve what he believes to be good” would “strain our imaginative power.”




Fig. 1 Experiment 1. Percent of participants ascribing thick belief and thin belief in the Jaded Politician case

The second worry regards the thin-belief probe. Perhaps people are interpreting “on some level” in a very broad and indiscriminate way, so as to include basically everyone regardless of their beliefs. “After all,” it might be thought, “doesn’t everyone think we ought to help the poor, on some level?” To mitigate this worry, we developed a storyline about an elitist politician who has always opposed anti-poverty legislation:

[ELITIST POLITICIAN] Michael has been a politician for decades. For years he has cared deeply about wealthy people and worked tirelessly on big-business legislation because it was the best way to serve their interests. Michael has never supported any kind of anti-poverty legislation. He remembers his past and all the reasons why serving wealthy people will make him money. So he does not have any motivation to help poor people. When new anti-poverty legislation is proposed, Michael does not support it.

This addresses the second worry by offering participants a chance to reveal that they attribute the relevant thin belief indiscriminately to anybody.

To test these two worries, we conducted another experiment looking at thin and thick belief ascription in these two cases above. Participants ($N=169$, aged 18–63 years, mean age=28.59, 54 female) in Experiment 2 were divided up into four groups. Participants were presented with either the Very Jaded Politician or the Elitist Politician cover story. All participants answered comprehension and motivational control questions similar to those from Experiment 1. As in Experiment 1, participants were assigned to either the Thick Belief condition or the Thin Belief condition and asked one of these two questions:

Does Michael believe that he ought to help poor people? [Yes/No]

At least on some level, does Michael think that he ought to help poor people?
[Yes/No]

These questions were then followed by the same confidence measure used in Experiment 1.

We made four predictions, two for the Very Jaded Politician and two for the Elitist Politician story. First, for the Very Jaded Politician story, rates of thin-belief ascription will be significantly higher than rates of thick-belief ascription.

Second, rates of thin-belief ascription in Very Jaded Politician will continue to exceed chance. Third, for the Elitist Politician story, rates of thin-belief ascription will not differ significantly from rates of thick-belief ascription. Fourth, for the Elitist Politician story, rates of thin-belief and thick-belief ascription will both be significantly below chance.

Again, all of our predictions were true, addressing both objections listed above. Beginning with the Very Jaded Politician stories, first, belief ascription was significantly higher in the Thin Belief condition (69 %) than in the Thick Belief condition (46 %), $X^2(1, n=101)=5.29, p<.05$, Cramer's $V=.229$.¹² Second, thin-belief ascription was significantly greater than chance, binomial test, $p<.05$, test proportion = .5, whereas thick-belief ascription didn't differ from chance, binomial test, $p=.672$, test proportion = .5. Moving on to the results for the Elitist Politician stories, third, rates of thin-belief ascription (6 %) did not differ significantly from rates of thick-belief ascription (0 %), $X^2(1, n=67)=1.89, p=.17$. Fourth, rates of thin-belief and thick-belief ascription were both significantly below chance, binomial test, $ps<.01$, test proportion = .5. These results are shown in Fig. 2.

4.3 Lying in a Comma: Experiment 3

Experiment 2 replicates the findings of Experiment 1 regarding the role of thin and thick categories of belief. Results also seem to mitigate the worries presented in Section 4.2. The Very Jaded cases show that thin belief continues to be ascribed to agents without the corresponding motivation to act. The Elitist cases show that thin belief is not ascribed indiscriminately, but rises to a minimal level sufficient for belief. However Experiments 1–2 still used the same basic cover-stories, closely resembling Stocker's original thought experiment. In Experiment 3, we sought to expand our results to cases not only beyond politics and charitable giving, but also to cases involving non-moral evaluative judgments.

Participants ($N=123$, aged 18–67, mean age = 30.43 years, 39 female) in Experiment 3 were divided up into four groups. Members of the first two groups saw the following case involving a traditional moral judgment that should be very familiar from the opening example of this paper:

[LIAR] Michael works for minimum wage at a local sandwich shop. Even though the job did not pay well, Michael had always been an extremely conscientious and honest worker. He took pride in every sandwich he made, doing his best to serve every customer. But over the years, Michael slowly became frustrated at the lack of respect and appreciation he received from his boss. After being passed up for promotion several times, Michael decides to begin lying about how much overtime he works. He remembers his past and all the reasons why he should be honest. But he no longer has any motivation to be honest. Now any time Michael

¹² One participant was removed from the analysis for answering that the protagonist “still has some” motivation for helping the poor.




Fig. 2 Experiment 2. Percent of participants ascribing thick belief and thin belief in the Very Jaded and Elitist Politician cases

is confident he won't be caught, he lies to his employer about how much overtime he works.

Each participant presented with the Liar case was then asked comprehension check and motivational control questions:

Michael is a: [sandwich shop employee/aid worker/insurance salesman]

Michael _____ motivation to tell his employer the truth about how much overtime he works. [still has some/no longer has any]

As in Experiment 2, participants were then assigned to either the Thick Belief condition or the Thin Belief condition and asked one of these two questions:

Does Michael believe that he ought to tell his employer the truth about how much overtime he works? [Yes/No]

At least on some level, does Michael think that he ought to tell his employer the truth about how much overtime he works?[Yes/No]

A third and fourth group of participants saw the following case involving a non-moral judgment involving a conventional norm:

[COMMA] Connie works as a copy editor for a university press. Even though it's a difficult job, Connie has always been extremely attentive and conscientious, especially when it comes to identifying comma splices. But over the years, Connie slowly became frustrated at the lack of respect and appreciation she received from manuscript authors. After having them reverse her helpful suggestions many times, Connie decides to stop correcting comma splices. She remembers her training and the reasons why comma splices are ungrammatical. But she no longer has any motivation to correct them. Now whenever a manuscript contains a comma splice, Connie doesn't change it.

Each participant presented with the Comma case was then asked comprehension check and motivational control questions:

Connie is a: [copy editor/aid worker/insurance salesperson]

Connie _____ motivation to correct the comma splices. [still has some/no longer has any]

Participants were again assigned to either the Thick Belief condition or the Thin Belief condition and asked one of the these two questions:

Does Connie believe that she ought to correct the comma splices? [Yes/No]

At least on some level, does Connie think that she ought to correct the comma splices? [Yes/No]

Thick and thin belief questions in all conditions were then followed by the same confidence measure used in Experiments 1–2.

This time we made two predictions. First, we predicted that we would replicate the earlier findings showing significantly more ascription for thin belief than thick belief, even in cases that had nothing to do with politics and helping the poor. Second, we predicted that we would continue to see differences between thick and thin belief ascriptions in the non-moral case.

Both of these predictions were borne out.¹³ Among the participants passing comprehension and motivation checks, 87 % ascribed thin belief in the Comma case, compared to only 65 % ascribing thick belief, $X^2(1, n=62)=4.31, p<.05$, Cramer's $V=.264$. And in Liar, 76 % of participants ascribed thin belief. But only 17 % ascribed thick belief, $X^2(1, n=58)=20.03, p<.01$, Cramer's $V=.588$. Liar thin-belief was significantly greater than chance, binomial test, $p<.01$, test proportion=.5. Liar thick-belief was significantly below chance, binomial test, $p<.01$, test proportion=.5. Comma thin-belief was significantly greater than chance, binomial test, $p<.01$, test proportion=.5. Comma thick-belief was not significantly greater than chance, binomial test, $p=.150$, test proportion=.5. These results can be seen in Fig. 3.

4.4 An Internalist Objection: Experiment 4

Experiment 3 replicates the findings of Experiment 2 and extends them to new cover stories featuring an additional moral duty involving honesty and a non-moral duty involving grammar. However, one objection is that these stimuli do not rule out additional moral beliefs about acting otherwise.¹⁴ In the Liar case, for instance, although the story states that Michael “remembers his past and all the reasons why he should be honest,” it does not explicitly state that Michael is unaware of new moral reasons against being honest. It is unclear how this additional belief would explain

¹³ Three participants were removed from the analysis for failing comprehension or motivation checks in Experiment 3.

¹⁴ We credit an anonymous referee for raising this objection.

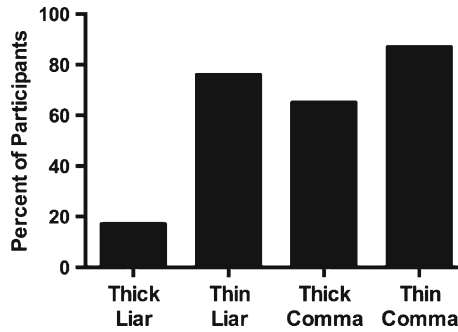


Fig. 3 Experiment 3. Percent of participants ascribing thick belief and thin belief in the Liar and Comma cases

differences in ascriptions observed in prior studies or favor internalist or externalist positions. Nonetheless, we conducted a follow-up experiment to test whether differences in thick and thin ascription can be explained by additional moral beliefs to act otherwise.

Participants ($N=80$, aged 19–67, mean age = 35.5 years, 34 female) in Experiment 4 were given adapted versions of the Liar case. The only change to the case was the antepenultimate sentence explicitly stating that Michael has not adopted any new moral reasons against being honest:

[UNMOTIVATED LIAR] He remembers his past and all the reasons why he should be honest. And he is aware of no morally valid reason against being honest. But he no longer has any motivation to be honest. Now any time Michael is confident he won't be caught, he lies to his employer about how much overtime he works.

After seeing this case, participants were assigned to the thick belief condition or thin belief condition and answered questions according to the same procedure used in Experiment 3.

Our prediction was that we would replicate the earlier findings showing significantly more ascription for thin belief than thick belief. This prediction was true. Among the participants passing comprehension and motivation checks, 55 % ascribed thin belief, compared to only 16 % ascribing thick belief, $X^2(1, n=77)=12.49, p<.01$, Cramer's $V=.403$.¹⁵ Thick belief ascription was significantly below chance, binomial test, $p<.01$, test proportion = .5. However thin belief ascription did not significantly rise above chance. We conclude that differences in thick and thin belief ascription do not depend additional moral beliefs to act otherwise.

4.5 Excursus: Perfect and Imperfect Duties

The results of Experiments 3–4 further support the hypothesis that the thick/thin belief distinction underlies the internalism/externalism debate about moral motivation.

¹⁵ Three participants were removed from the analysis for failing comprehension or motivation checks in Experiment 4.

However one thing happened in Experiment 3 that we did not predict. Although we saw the same basic pattern of belief ascription, the size of the difference between thin belief and thick belief ascription differed. More specifically, although we showed that thick and thin belief ascription differs in cases involving moral as well as non-moral judgments, responses in the original Stocker case about charitable giving seemed to most closely resemble responses in the Comma (or non-moral) case than they did in the Liar (or straightforwardly moral) case. We think this is an interesting question to pursue in future work. And although it is largely orthogonal to our present goal — our task was to demonstrate the role of thick and thin belief — we offer a tentative hypothesis about why this might be so.

Let us begin by returning to Stocker's original article and the politician thought experiment. When considering possible explanations for the jaded politician's behavior towards the poor, Stocker writes, "Perhaps the politician is not attracted to helping those people now because he believes he has already done enough for them" (p. 742). We think that this interpretation of the case is very natural, and we suspect that many participants in the experiments might also judge that the politician has already done enough to help the poor.

Why is Stocker's proposed reading of the case natural? Because aiding the less fortunate is an imperfect duty. The text of the story makes it seem likely that the politician has done enough such service for a lifetime. If that's the case, then he has already displayed enough motivation to satisfy any reasonable internalist requirement on moral judgment. By contrast, the duty to not lie is perfect. And whereas fulfillment of an imperfect duty is sensitive to past performance, fulfillment of a perfect duty is not. Donating to charity 10 weeks in a row plausibly excuses one from donating on the eleventh week. But even if a worker has told the truth 10 weeks in a row about how many hours they worked, that won't excuse lying on the eleventh week. In order for Stocker's thought experiment to count as a legitimate counterexample to internalism, the agent must make a sincere moral judgment without the corresponding motivation to act. But one plausible hypothesis is that intuitions about whether or not this is the case will depend, in part, on whether judgments involve perfect or imperfect duties.

If the intuitions reported also track this feature of the cases, then we would expect higher rates of belief ascription among the cases involving perfect duty than in the imperfect duty cases. We would also predict a greater disparity between rates of thick- and thin-belief ascription in cases featuring perfect duties than in cases featuring imperfect duties. The reason is that, all things considered, factoring in past performance increases the likelihood that the politician's behavioral profile exhibits sufficient motivation. And conversely, since the perfect duties are insensitive to past performance, participants will be less likely to judge that the liar is sufficiently motivated to be honest. And that is exactly what happens: we see larger effects in the liar case than in the politician or comma cases. As Stocker suspected, perhaps the politician has done his fair share.

One final note about the politician case. We have attempted to show that cuing thick or thin categories of belief plays an important role in generating externalist and internalist intuitions. But our additional hypothesis about the variation between cover stories might also shed new light on Stocker's basic politician case independently of this larger goal. If charitable giving is an imperfect duty that the politician has plausibly

fulfilled through his prior activities or legislative legacy, then it is unclear how the intuition that he “believes he ought to help the poor” without “motivation to do so at the very end of the story” poses any problem for the internalist in the first place.

4.6 Prior Experimental Findings

We would like to close this experimental section by commenting on how our theory relates to some other experimental results on moral motivation. Specifically, we note that our findings are very similar in spirit to the findings of Strandberg and Björklund (2013) on internalism and ordinary intuitions. In their study, Strandberg and Björklund present participants with a basic story in which various agents named “Anna” have no motivation to give to charity. They then ask participants the following question:

Could it be the case that Anna thinks she is morally required to give some of her money to the starving even if she is not motivated at all to do so?

They report finding “surprisingly little support to the most celebrated versions of internalism” (p. 319). The majority of participants across most of their scenarios gave the externalist-friendly answer. They answer “yes,” it is possible for Anna to think that she is morally required to donate despite lacking any corresponding motivation.

These results fit well with our hypothesis and results, for two main reasons. First, notice that the question above asks participants about what “Anna thinks” rather than what “Anna believes”. Given that “thinks” is more likely to cue thin-belief ascription, our view predicts agreement that, all else equal, moral judgment does not require motivation. And that’s exactly what happens. Just like our politician and sandwich shop employee cases, thin belief ascription in Strandberg and Björklund cases aligns well with motivational externalism. We would caution, however, that this doesn’t tell the full story.

Second, Strandberg and Björklund’s cases all involve charitable donation, which is an imperfect duty. If our additional hypothesis about charitable giving (Section 4.3) is on the right track, we would also predict a generally tendency towards “yes” answers in these cases. The reason, to reiterate, is that if charitable giving is an imperfect duty that can be fulfilled by past performance, then agents might have already exhibited enough motivation to satisfy the internalist requirement. Thus we would urge caution in interpreting results about charitable giving as evidence against internalism. An appropriate precaution would be to use a wider range of cover stories involving different types of duty, as we did above.

In sum, we find much to applaud in Strandberg and Björklund’s enlightening work. We take their studies to independently confirm our results on thin-belief ascription.¹⁶ We also take their studies to support our hypothesis that moral beliefs about perfect duties have a more demanding motivational profile than moral beliefs about imperfect duties. However, for these same reasons, we doubt that their results should inspire “skepticism” about internalism. Instead Strandberg and Björklund have identified but one ordinary way of thinking that is friendly to externalism. As our results show, this is

¹⁶ Insofar as participants’ judgments about the possibility of moral requirements in Strandberg and Björklund’s experiments serve as proxy for sincere moral judgments in our own.

only half the story. One can also obtain the opposite result by cuing another ordinary way of thinking that is friendly to internalism through thick-belief ascription, especially in cases of perfect duties. Thus by calling attention to this distinction we hope not only to contribute to the ongoing internalist/externalist debate but also to future experimental work in philosophy on the connection between belief and moral judgment.

5 Moral Motivation: a Question about Belief

Hume tells us that “All . . . of sound reason are disgusted with verbal disputes, which abound so much in philosophical and theological enquiries.” However Hume also proposes a remedy for such situations. For this we must turn to “clear definitions, from the precision of those ideas which enter into any argument, and from the strict and uniform use of those terms which are employed” (Hume 1779: section 12.7).

The Humean cure works wonders in the present case. By clearly distinguishing the categories of thick belief and thin belief, we were able to identify an underlying source of persistent disagreement between motivational internalists and externalists. When people are probed in terms of thick belief, they are significantly less likely to ascribe moral belief to an unmotivated agent. Conversely, when people are probed in terms of thin belief, they are much more likely to ascribe moral belief to an unmotivated agent. We observed this general pattern in cases that very closely resemble Stocker’s original jaded politician thought experiment. We replicated this effect using completely different cover stories. In fact, in the experiment involving lying, we were able to completely reverse people’s judgments by switching probes. When people are encouraged to respond in terms of thick belief, belief-ascription was almost as low as it could get, which fits perfectly with the internalist view. By contrast, when people are encouraged to respond in terms of thin belief, belief-ascription was almost as high as it could get, which fits perfectly with the externalist view. It’s hard to imagine better evidence than this for our diagnosis of the internalism/externalism debate. Lastly, we’ve also shown that this effect persists in cases involving non-moral evaluative judgments. We found similar differences in our adaptations of Stocker’s original thought experiment as we did in cases involving the correction of grammatical errors.

Turning back to Stocker’s original thought experiment for a moment, Dreier (1990, p. 10) writes:

Of course, there are different ways of characterizing the politician’s situation. Maybe he has ceased to *believe* that it is a good thing to help others. That would be a way of describing the politician’s state that does not jeopardize our simple principle. But Stocker’s insight is that it makes perfect sense to describe the uncaring man another way. Couldn’t he say, “I have not changed my moral *beliefs*; I still think the morally good thing to do is to help the unfortunate. But I am no longer interested in doing what is morally good.” Stocker wants us to notice that it is a perfectly intelligible, all too familiar thing to say. “Lack of [the desire to do what one *believes* to be good] is commonplace (emphasis added).

Later in the same passage, Dreier writes that cases like the jaded politician are “highly abnormal” and that imagining “a man who would not lift a finger, as it were, to achieve what he believes to be good” would “strain our imaginative power” (p. 12).

Dreier’s important observation is right in line with our suggestion that key intuitions in thought experiments — like the jaded politician — that have perpetuated the debate between internalist and externalists persist due to a general issue about the nature and categories of belief. It strains the imagination to suppose that radically disaffected people such as the politician or the sandwich shop employee hold thick moral beliefs. By contrast, it doesn’t strain the imagination to suppose that such people hold thin moral beliefs.

Of course, there certainly still are deep, illuminating, and important research questions that remain to be investigated at the intersection of philosophy and psychology when it comes to the nature of moral judgments and their bearing on other mental states and behaviors. But we humbly submit that, in the words of Hume, “clear definitions” and “the strict and uniform use” of the categories of beliefs allows us to bypass the perennial metaethical debate on moral motivation.

To clarify, we are not saying that philosophers could not construct further thought experiments on moral motivation that serve as counterexamples to internalism or externalism. We are not suggesting that appeals to commonsense or intuition exhaust the evidence in favor of internalism or externalism. Nor do we deny that metaethicists have offered theoretical arguments for internalism or externalism that are relatively independent of intuitions about particular cases. Rather, we are simply pointing out that some basic features of folk psychology — our ordinary theory of mind which informs mental-state attribution to other people — can straightforwardly explain why there is such fundamental intuitive disagreement between motivational internalists and externalists in the first place. We are suggesting that the division among philosophers derives from a fault line running deep in the folk psychology of belief familiar to us from everyday life.¹⁷

To that end, we submit that we have also made significant progress in further elucidating two familiar categories from folk psychology: thick and thin belief. Building on past findings for the role of thick and thin belief in the debate in epistemology on belief entailment, our results indicate that thick belief ascriptions are more sensitive to an agent’s behavioral profile. People tended to ascribe thick belief that P when agents actively promoted an agenda that makes sense given P. By contrast thin belief ascriptions were largely insensitive to an agent’s behavioral profile. People ascribed thin belief that P in situations in which agents actively failed to promote an agenda that makes sense given P. These findings suggest that their use in predicting and evaluating action and assertion marks an important difference between thick and thin belief.¹⁸

¹⁷ To that extent our discussion assumes that philosophers are people too and, thus, aren’t immune to the influence of folk psychology. For similar results to this effect see Vaesen et al. (2013) suggesting that epistemic intuitions of professional epistemologists are sensitive to native language, and results by Schulz et al. (2011) suggesting that certain personality differences among experts predict intuitions about free will and moral responsibility. For studies involving heuristics and biases among professionals see Machery (2012); Schwitzgebel and Cushman (2012, 2015); Knobe and Samuels (2013); Tobia et al. (2013).

¹⁸ For more on the role of assertion in demarcating categories of thick and thin belief see Rose et al. (2014).

These findings underscore the usefulness of experimental cognitive science for adjudicating perennial debates in philosophy (for an overview of other examples, see Buckwalter and Turri [forthcoming](#)). If we are correct, then crucial intuitions perpetuating the internalist/external track general features of the folk-psychological categories of thick and thin belief. It is doubtful that the source of intuitions perpetuating this debate would have been properly diagnosed through introspection or uncontrolled social observations alone.

Lastly, further work is required to answer many important psychological and philosophical questions related to our discussion. On the psychological side, it remains to be seen whether the categories of thin and thick belief are fundamental in folk psychology, or whether there are even further, more specific categories. Although we doubt this is the case for thin belief, we do not doubt that it is true for thick belief. For example, it's possible that folk psychology recognizes an important distinction between those propositions whose truth one approves of or emotionally endorses, on the one hand, and those propositions that we are willing to explicitly assert, on the other. The experiments reported here weren't designed to investigate these possibilities. Further work is needed on this front.

On the philosophical side, the category of thin belief in ordinary moral cognition might be taken to support motivational externalism. Internalists argue that moral judgment entails corresponding motivation, while externalists deny that moral judgment requires this. Rejecting entailment requires only one case in which moral belief is present but motivation is absent. Thus the fact that thin moral beliefs do not require motivation may be used to construct counterexamples to motivational internalism by demonstrating at least one type of moral belief that lacks motivation.

Our findings also have significant implications for several fundamental debates in contemporary metaethics. For example, debates over the correct semantic theory of moral discourse turn partly on how the motivational internalism/externalism debate is resolved (Dreier 1990, 2006). Non-cognitivism and relativist cognitivism have been motivated by appeal to some version of motivational internalism (for variations on this theme see discussion in Blackburn 1984; Railton 1986; Brink 1989; Shafer-Landau 2003). And debates of the correct moral semantics in turn have consequences for the correct ontological theory of moral attributes (Mackie 1977; Korsgaard 1996). Future work is needed to trace the ramifications of our findings for these other core metaethical questions.

References

- Armstrong, D. 1969. Does knowledge entail belief? *Proceedings of the Aristotelian Society* 70: 21–36.
- Björklund, F., G. Björnsson, J. Eriksson, R. Francén Olinder, and Caj Strandberg. 2012. Recent work on motivational internalism. *Analysis* 72(1): 124–137.
- Blackburn, S. 1984. *Spreading the word*. Oxford: Oxford University Press.
- Blackburn, S. 1998. *Ruling passions*, 59–68. Oxford: Clarendon.
- Brink, D. 1989. *Moral realism and the foundations of ethics*. Cambridge: Cambridge University Press.
- Brink, D. 1997. Moral motivation. *Ethics* 108: 4–32.
- Buckwalter, W., D. Rose, and J. Turri. 2015. Belief through thick and thin. *Noûs* 49(4): 748–775.
- Buckwalter, W., and J. Turri. *forthcoming*. Moderate scientism in philosophy. *Forthcoming in* *Scientism: prospects and problems*. In eds. de Ridder, J., Peels, R., and von Woudenberg, R. Oxford University Press.

- Chalmers, D.J. 2011. Verbal disputes. *Philosophical Review* 120(4): 515–566.
- Cohen, J. 1966. More about knowing and feeling sure. *Analysis* 27: 11–16.
- Copp, D. 1997. Belief, reason, and motivation: Michael Smith's the moral problem. *Ethics* 108: 33–54.
- Dreier, J. 1990. Internalism and speaker relativism. *Ethics* 101: 6–26.
- Dreier, J. 2006. Moral relativism and moral nihilism. In *Oxford handbook of ethical theory*, ed. D. Copp, 240–264. New York: Oxford University Press.
- Dretske, F. 1983. The epistemology of belief. *Synthese* 55(1): 3–19.
- Falk, W.D. 1947. "Ought" and motivation. *Proceedings of the Aristotelian Society* 48: 111–138.
- Fodor, J.A. 1990. *A theory of content*. Cambridge: MA. MIT Press.
- Gendler, T.S. 2008a. Alief and belief. *Journal of Philosophy* 105(10): 634–663.
- Gendler, T.S. 2008b. Alief in action (and reaction). *Mind & Language* 23(5): 552–585.
- Hume, D. 1748 [1977]. In *An enquiry concerning human understanding*, (2nd ed.), ed. E. Steinberg. Indianapolis: Hackett.
- Hume, D. 1772 [1992]. Of the dignity or meanness of human nature. In *Essays and Treatises on Several Subjects*. London: Basil.
- Hume, D. 1779 [2007]. *Dialogues concerning natural religion*. Ann Arbor: University of Michigan Library. URL <<http://name.umd.umich.edu/004895521.0001.000>>.
- James, W. 1889. The psychology of belief. *Mind* 14(55): 321–352.
- Jones, O.R. 1971. Knowing and guessing: By examples. *Analysis* 32: 19–23.
- Kauppinen, A. 2015. Intuition and belief in moral motivation. In *Moral internalism*, ed. G. Björnsson, C. Strandberg, R.F. Olinder, J. Eriksson, and F. Björklund. Oxford: Oxford University Press.
- Knobe, J., and R. Samuels. 2013. Thinking like a scientist: Innateness as a case study. *Cognition* 126: 72–86.
- Korsgaard, C.M. 1996. *The sources of normativity*. Cambridge: Cambridge UP.
- Kriegel, U. 2012. Moral motivation, moral phenomenology, and the alief/ belief distinction. *Australasian Journal of Philosophy* 90(3): 469–486.
- Leher, K. 1974. *Knowledge*. Oxford: Clarendon.
- Machery, E. 2012. Expertise and intuitions about reference. *Theoria* 73: 37–54.
- Mackie, J. 1977. *Ethics: Inventing right and wrong*. New York: Penguin.
- May, J. 2014. On the very concept of free will. *Synthese* 191(12): 2849–2866.
- Mele, A. 1996. Internalist moral cognitivism and listlessness. *Ethics* 106: 727–753.
- Miller, A. 2003. *An introduction to contemporary metaethics*. Cambridge: Polity Press.
- Murray, D., J. Sytsma, and J. Livengood. 2013. God knows (but does God believe?). *Philosophical Studies* 166(1): 83–107.
- Myers-Schulz, B., and E. Schwitzgebel. 2013. Knowing that P without believing that P. *Noûs* 47(2): 371–384.
- Nichols, S. 2002. On the genealogy of norms: A case for the role of emotion in cultural evolution. *Philosophy of Science* 69: 234–255.
- Nichols, S. 2004. *Sentimental rules: On the natural foundations of moral judgment*. Oxford: Oxford University Press.
- Peirce, C.S. 1958. In *Collected papers of Charles Sanders Peirce*, ed. Charles Hartshorne, Paul Weiss, and Arthur W. Burks. Cambridge: Harvard University Press.
- Radford, C. 1966. Knowledge — by examples. *Analysis* 27: 1–11.
- Railton, P. 1986. Moral realism. *Philosophical Review* 95: 163–207.
- Rosati, Connie S. 2006. Moral motivation. In *The stanford encyclopedia of philosophy*, (Fall 2008 Edition), ed. Edward N. Zalta. URL = <<http://plato.stanford.edu/archives/fall2008/entries/moral-motivation/>>.
- Rose, D., W. Buckwalter, and J. Turri. 2014. When words speak louder than actions: Delusion, belief and the power of assertion. *Australasian Journal of Philosophy* 92(4): 683–700.
- Roskies, A. 2003. Are ethical judgments intrinsically motivational? Lessons from 'acquired sociopathy'. *Philosophical Psychology* 16: 51–66.
- Roskies, A. 2006. Patients with ventromedial frontal damage have moral beliefs. *Philosophical Psychology* 19: 617–627.
- Roskies, A. 2008. Internalism and the evidence from pathology. In *Moral psychology*, vol. 3, ed. W. Sinnott-Armstrong, 191–206. Cambridge: The MIT Press.
- Rose, D., and J. Schaffer. 2013. Knowledge entails dispositional belief. *Philosophical Studies* 166(1): 19–50.
- Sayre-McCord, G. 1997. The metaethical problem. *Ethics* 108: 55–83.
- Schulz, E., E.T. Cokely, and A. Feltz. 2011. Persistent bias in expert judgments about free will and moral responsibility: A test of the expertise defense. *Consciousness and Cognition* 20: 1722–1731.
- Schwitzgebel, E., and F. Cushman. 2012. Expertise in moral reasoning? Order effects on moral judgment in professional philosophers and non-philosophers. *Mind & Language* 27: 135–153.

- Schwitzgebel, E., and F. Cushman. 2015. Philosophers' biased judgments persist despite training, expertise and reflection. *Cognition* 141: 127–137.
- Shafer-Landau, R. 2003. *Moral realism. A defence*. Oxford: Clarendon.
- Smith, M. 1994. *The moral problem*. Oxford: Basil Blackwell.
- Stocker, M. 1979. Desiring the bad: an essay in moral psychology. *The Journal of Philosophy* 76: 738–753.
- Strandberg, C., and F. Björklund. 2013. Is moral internalism supported by folk intuitions? *Philosophical Psychology* 26(3): 319–335.
- Svavarsdottir, S. 1999. Moral cognitivism and motivation. *The Philosophical Review* 108: 161–219.
- Tobia, K., W. Buckwalter, and S. Stich. 2013. Moral intuitions: Are philosophers experts? *Philosophical Psychology* 26(5): 629–638.
- Vaesen, K., M. Peterson, and B.V. Bezooijen. 2013. The reliability of armchair intuitions. *Metaphilosophy* 44(5): 559–578.
- Williams, D.C. 1937. The meaning of 'good'. *Philosophical Review* 46(4): 416–423.
- Zangwill, N. 2008. The indifference argument. *Philosophical Studies* 138: 91–124.