

Belief through Thick and Thin

WESLEY BUCKWALTER
University of Waterloo

DAVID ROSE
Rutgers University, New Brunswick

JOHN TURRI
University of Waterloo

Abstract

We distinguish between two categories of belief—thin belief and thick belief—and provide evidence that they approximate genuinely distinct categories within folk psychology. We use the distinction to make informative predictions about how laypeople view the relationship between knowledge and belief. More specifically, we show that if the distinction is genuine, then we can make sense of otherwise extremely puzzling recent experimental findings on the entailment thesis (i.e. the widely held philosophical thesis that knowledge entails belief). We also suggest that the distinction can be applied to debates in the philosophy of mind and metaethics.

1. An Ambivalent Introduction

Imagine an elected official, Michael, who emphasizes the importance of giving to charity and helping the poor in all of his speeches. Michael commands all the relevant facts and figures about why it is good to help the less fortunate. He can explain exactly how giving will help members of his district, as well as quantify the net benefit of this for the overall well-being of society as a whole. But as the years passed, Michael made fewer and fewer speeches encouraging others to help the poor. In fact, on the day of a crucial vote that would pass new aid legislation, Michael does literally nothing to support it, even though he could easily offer some support at almost no cost to himself. In light of all that, does Michael *believe* that the aid legislation is important?

If you're like us, the question leaves you ambivalent. On the one hand, given his command of the facts about why it would be good to help the less fortunate, it seems that on at least some level Michael thinks that the aid legislation is important. That is, on some level he accepts that the proposition \langle the aid legislation is important \rangle is true. On the other hand, when given the opportunity, he doesn't support or vote the bill into law, which makes it seem that he doesn't believe that the legislation

is important. To put our ambivalence another way, it seems that on a purely intellectual level, Michael believes the proposition, whereas on a more practical level, he doesn't.¹

Or consider a physics student, Stewart, who pays close attention in class. He can explain exactly how the heliocentric theory of the solar system works, all the observational and mathematical evidence showing that planets orbit the sun and that the sun doesn't orbit the planets. On the final exam for the course, Stewart answers that 'the earth revolves around the sun' and gets a perfect score for his description of how the heliocentric theory works. Nevertheless, prior to starting at the university Stewart was home-schooled by his parents. His parents taught him that the earth is at the center of the universe. And Stewart has always accepted what his parents have taught him. In light of all that, does Stewart *believe* that the earth revolves around the sun?

Once again, the question leaves us ambivalent. On the one hand, given his command of the facts, it seems that on at least some level Stewart thinks that the earth revolves around the sun. That is, on some level he accepts that the proposition < the earth revolves around the sun > is true. After all, he has seen and understood the compelling evidence presented in his physics class regarding the movement of the earth around the sun. On the other hand, Stewart has always devoutly accepted what his parents have taught him. Holding on to these lessons from his childhood outside of physics class, it seems that he doesn't believe that the earth revolves around the sun. Once again it seems to us that on some level the agent believes the proposition, whereas on a more practical level, he doesn't.

In order to have convenient labels for these divergent senses of belief, let's call them *thin belief* and *thick belief*. Although we're not positioned to confidently defend a full-blown analysis of either category, we can helpfully characterize them in a way that maps onto the ambivalence we feel about the earlier examples.

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.²

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.³ Thick belief entails thin belief, but not vice versa.

We suspect that thin and thick beliefs, as characterized above, approximate genuine and genuinely distinct categories within folk psychology. That is, we suspect that the ambivalence we feel when considering the examples of Michael and Stewart reflects our implicit competence with two related but distinct folk psychological categories, each of which often finds expression with 'belief'. If we're right about

that, it would be an important fact about folk psychology with implications for multiple disciplines. For example, it would be relevant to research on belief and belief-ascription in cognitive psychology, to theory-of-mind research in developmental psychology, to philosophers of mind analyzing the concept of belief, and to epistemologists studying the norms of belief and its relationship to knowledge. We also think the distinction can contribute to substantive philosophical theorizing on belief (a topic we will return to in Section 6).⁴

Thin and thick belief both make appearances when philosophers and cognitive scientists characterize belief. Something approximating the category of thin belief that we have in mind (although not explicitly labeled as “belief”) dominates discussions in contemporary philosophy of mind and cognitive science on representationalist models of cognition. 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).⁵ (We discuss more examples of contemporary theorists employing the thin category of belief in Section 4 below.)

But the thick sense makes important appearances in the history of philosophy. 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,” fit to be a “governing principle of our actions.” Hume expresses himself concisely in the following passage:

The imagination has the command over all its ideas, and can join and mix and vary them, in all the ways possible. It may conceive fictitious objects with all the circumstances of place and time. It may set them, in a manner, before our eyes, in their true colors, just as they might have existed. But as it is impossible, that this faculty of imagination can ever, of itself, reach belief, it is evident, that belief consists not in the peculiar nature or order of ideas, but in the *manner* of their conception, and in their *feeling* to the mind. I confess, that it is impossible perfectly to explain this feeling or manner of conception. We may make use of words, which express something near it. But its true and proper name . . . is *belief*; which is a term, that every one sufficiently understands in common life.

On Hume’s view, belief *ordinarily understood* requires more than thin belief. On Hume’s view, belief *ordinarily understood* is thick belief.⁶

Ambivalence, suspicions, and implicit disagreement among luminous predecessors can work wonders in the context of discovery, but they’re inadequate in a context of justification. Having explained our motivation for positing a distinction between thin and thick belief, our goal in the remainder of this paper is to begin justifying it. Figuratively speaking, if hypotheses are cars, we’re making a

downpayment on ours. In order to accomplish this, we will *assume that the distinction between thin and thick belief is genuine* and show that it makes informative predictions and helps to explain otherwise puzzling observations. In short, we will *judge the hypothesis by its fruits*. We ask others to do the same.

Our point of departure in this paper is the long-held view in epistemology that knowledge entails belief, which we will call *the entailment thesis* (Lehrer 1968: 491). Among contemporary epistemologists, the entailment thesis is almost unanimously accepted as intuitively obvious and uncontroversial. In this respect, it is often viewed as roughly on par with *the factivity thesis*, which says that knowledge requires truth.⁷ Contemporary epistemologists fancy that this uncontroversial consensus spans virtually the entire history of Western philosophy, under the title, “the problem of the *Theaetetus*.” For example, here is how Roderick Chisholm, a titan of twentieth-century epistemology, matter-of-factly puts matters in the third edition of his influential textbook, *Theory of Knowledge*:

If you know that it is raining, then it *is* raining and you *believe* that it is raining. The point may be generalized by saying that, if you have knowledge, then you have true belief. But knowledge is more than mere true belief. For your belief that it is raining could be true even if you didn't *know* that it is raining. Perhaps you have simply made a lucky guess. What, then, must be *added* to true belief to get knowledge? Finding the answer to this question is sometimes called, “the problem of the *Theaetetus*,” since the question was first clearly formulated by Plato in his dialogue of that name. (Chisholm 1989: 90)

But recently several experimental philosophers have argued that the entailment thesis is neither obvious, uncontroversial, nor widely accepted (Myers-Schulz and Schwitzgebel 2013, Murray et al. forthcoming). Indeed, their results suggest that most people are perfectly willing to ascribe knowledge in the absence of belief. And so it may seem that a pall of doubt falls over the entailment thesis and, in the process, diminishes the probative value of “expert” intuition and consensus in philosophy.

Is this the beginning of an unhappy end for the entailment thesis? We think not, but it will take some hard work and ingenuity to effectively resist the experimental challenge. Equipped with our distinction between thin and thick belief, we will argue that the experimental results, when properly understood, don't cast doubt on the entailment thesis, when properly understood. The true lesson of the experimental results is that the entailment thesis should be understood in terms of *thin belief*, whereas the experimental results in question should be understood in terms of *thick belief*. It turns out that when people are primed to employ the category of thin belief, they respond exactly as we would expect if they accepted the entailment thesis (understood in our preferred way). Indeed, when understood properly, the entailment thesis is *as firmly rooted as the factivity thesis* in our ordinary ways of thinking and speaking about knowledge. Thus, putting the thick/thin distinction to work happily reconciles our folkways with the longstanding dominant view.⁸ This supports our hypothesis that *thin belief* and *thick belief* are genuine, distinct categories in folk psychology.

Here is the plan for the rest of the paper. Section 2 reviews the recent experimental challenge to the entailment thesis. Section 3 segues into our positive agenda. Section 4 motivates our preferred understanding of the entailment thesis, namely, the *thin entailment thesis*: knowledge entails thin belief. Section 5 reports five experiments that test whether the thin entailment thesis is widely accepted; the results strongly suggest that it is. Section 6 concludes the discussion by suggesting further, potentially important applications of the distinction between thick and thin belief.

2. The Experimental Challenge to the Entailment Thesis

Several recent experimental studies have challenged the entailment thesis.⁹ As already mentioned, the (simple, unadorned) entailment thesis says that knowledge entails belief. More specifically, the (simple, unadorned) entailment thesis says that *necessarily*, if you know that P, then you believe that P (see Armstrong 1969, p. 21). This section reviews the relevant experimental research.

2.1. *Quizzing Kate: Myers-Schulz and Schwitzgebel*

Why think that the entailment thesis is true? As Blake Myers-Schulz and Eric Schwitzgebel point out, a common strategy is to claim that the thesis is “obvious” and faces “no convincing counterexamples.” Some philosophers have claimed that there are convincing counterexamples (e.g. Radford 1966), but their claim has met stiff and widespread resistance (see e.g., Armstrong 1969, Cohen 1966, Jones 1971, Lehrer 1974). They’re accused of having peculiar, even perverse, intuitions about the relevant cases. Call these maligned intuitions *anti-entailment intuitions*.

Though the discussion over anti-entailment intuitions largely died out in the 1970’s, with the dominant response being Armstrong’s (1969) dismissal of anti-entailment intuitions, Radford claimed that the issue turned largely on what English speakers would or would not say about knowledge (Radford 1966, p. 5). The investigation of what English speakers would or would not say was never taken up until recently. Myers-Schulz and Schwitzgebel—suspecting that anti-entailment intuitions are neither perverse nor peculiar—ran experiments to see what laypeople thought about several cases, including a version of Colin Radford’s Unconfident Examinee case. According to the story, Kate studies very hard for an exam on English history, and she is doing very well on the exam “until she comes to the final question,” which asks, “What year did Queen Elizabeth die?” Kate “reviewed this date many times” in her studies and “had even recited” it “just a few hours earlier.” Moreover, when she reads the final question, she is initially “confident” that she will “recollect the answer.” But she is interrupted by a warning that time is almost up; she loses confidence and becomes nervous; she tries to recall “but nothing comes to her”; finally, “with a sigh of disappointment,” she says to herself, “I’ll just have to guess the answer.” So she writes down “1603,” which is in fact the correct answer.

In a between-subjects design, Myers-Schulz and Schwitzgebel asked half of their participants whether Kate “knows” that Queen Elizabeth died in 1603, and they

asked the other half whether Kate “believes” that Queen Elizabeth died in 1603. Participants could answer either ‘Yes’ or ‘No’. Myers-Schulz and Schwitzgebel observed that 87% answered that Kate knew, but only 37% answered that Kate believed. Myers-Schulz and Schwitzgebel ran further studies with several other cover stories, and twice they found similar large disparities between knowledge ascription and belief ascription. The crucial point is that in multiple cases rates of knowledge ascription far surpassed rates of belief ascription. Myers-Schulz and Schwitzgebel conclude that anti-entailment intuitions are neither peculiar, perverse, nor uncommon. Though perhaps not a majority, a substantial number of people seem to exhibit anti-entailment intuitions in certain sorts of cases. Advocates of the entailment thesis must do better than just sit back and claim that the entailment thesis is obvious.

Of course, Myers-Schulz and Schwitzgebel’s experimental challenge could have turned out otherwise. It could have turned out that rates of knowledge ascription and belief ascription matched. Had it turned out that way, it would have instead provided some evidence that philosophers were right to claim that the entailment thesis is obvious.

2.2. *Entailment, dispositional style: Rose and Schaffer*

David Rose and Jonathan Schaffer (forthcoming) argue that Myers-Schulz and Schwitzgebel’s studies don’t challenge the entailment thesis, for at least two reasons. First, Rose and Schaffer point out that a *within-subjects* design would have been more probative than a between-subjects design since what one really wants to know is whether the same person would attribute knowledge but deny belief. But when they re-ran the studies using a within-subjects design, the results were favorable to the anti-entailment camp: over 50% of participants ascribed knowledge but denied belief to Kate. Moreover, 64% of those who ascribed knowledge denied belief. Second, and more importantly, Rose and Schaffer argue that the entailment thesis should be understood in terms of *dispositional* belief: necessarily, if you know that P, then you dispositionally believe that P. However, Myers-Schulz and Schwitzgebel’s probes tended to elicit an occurrent reading of ‘believe’. So their results don’t pose a challenge to the entailment thesis, properly understood.

Occurrent belief is conscious endorsement. Dispositional belief is information stored in the mind available for endorsement. Dispositional beliefs are the kind of beliefs that sleeping people have. When Harry is fast asleep, he doesn’t consciously endorse that his cousin, Dudley, is despicable. Nevertheless, when Harry falls asleep, he doesn’t lose his belief that Dudley is despicable. Rather, he retains that belief in dispositional form, as information stored in mind and available to him for later endorsement.

Rose and Schaffer’s critique naturally suggests a more pointed empirical test for whether laypeople share the anti-entailment intuition: re-run Myers-Schulz and Schwitzgebel’s studies using probes specifically designed to elicit the dispositional reading of ‘believe’. If most participants still ascribe knowledge but deny dispositional belief to Kate, then people probably don’t find the entailment thesis intuitive, even when understood dispositionally.

In order to cue the dispositional sense of belief, Rose and Schaffer developed several novel techniques. For example, they directly asked participants whether Kate has the relevant disposition: *Is Kate disposed to believe that Queen Elizabeth died in 1603?*¹⁰ They also used a parenthetically qualified probe: *Did Kate still believe (in the sense that she still held the information in her mind, even if she could not access it) that Queen Elizabeth died in 1603?* When questioned these ways, roughly 60% of participants ascribed both knowledge and belief to Kate. Moreover, of those who ascribed knowledge, approximately 70% ascribed belief. This is exactly what we would expect if people generally accepted the entailment thesis.¹¹ Rose and Schaffer also applied these techniques to the other cases where Myers-Schulz and Schwitzgebel observed results consistent with a widely shared anti-entailment intuition. In each case, cueing the dispositional sense of belief reversed the earlier results.¹²

2.3. *Of dogs and gods: Murray, Sytsma and Livengood*

Dylan Murray, Justin Sytsma and Jonathan Livengood present additional experimental evidence that the anti-entailment intuition is widely shared. Moreover, responding to Rose and Schaffer, Murray et al. argue that their findings apply equally well to the dispositional entailment thesis. That is, Murray et al. argue that their findings avoid Rose and Schaffer's criticism of earlier findings on the anti-entailment intuition.

Murray et al. used stories featuring diverse agents, including a cash register, a border collie named 'Cassie', and God. They designed their stories and probes to avoid cueing a narrowly occurrent reading of 'believe'. For example, here is one their stories:

Researchers have found that some breeds of dog are surprisingly intelligent. Amongst the most intelligent are border collies. One border collie, Cassie, is even able to solve simple mathematical problems: If you ask her "what is $2 + 2$," for example, she will bark four times; similarly, if you ask her "what is $4 + 5$ " she'll bark nine times!

In a within-subjects design, participants were asked two questions (counterbalanced to avoid order effects): *'Does Cassie know that $2 + 2 = 4$?' and 'Does Cassie believe that $2 + 2 = 4$?' A non-trivial number of participants (28%) answered that Cassie knows but does not believe that $2 + 2 = 4$.¹³ Moreover, of those who ascribed knowledge, a majority (53%) did not ascribe belief, and of those who did not ascribe belief, nearly half (46%) ascribed knowledge. These are noteworthy results, suggesting that anti-entailment intuitions are widely shared. Importantly, Murray et al. claim, for two reasons, that it's unlikely that their participants were merely denying occurrent belief in this case. First, unlike in Myers-Schulz and Schwitzgebel's story about Kate, nothing "blocks" Cassie from "consciously accessing" the relevant information (as Rose and Schaffer suggest regarding the Unconfident Examinee case). Second, participants weren't asked about "any *particular occasions*" where Cassie might fail to consciously endorse the proposition.*

Murray et al. observed comparable results in three other experiments. One in particular provided the most striking evidence that the entailment intuition isn't

widely shared. This experiment featured their *Geocentrist case*, which inspired our earlier example of Stewart in Section 1.¹⁴ Here is the Geocentrist case:

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.

The results from Geocentrist are very surprising. Nearly half of the participants (46%) 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. Distinguishing between occurrent and dispositional belief seems unlikely to explain the results. Murray et al. conclude that philosophers should neither take the entailment thesis for granted any longer, nor assume that people generally find it intuitively compelling. Rather, the entailment thesis "stands in need of positive support." We agree that Murray et al.'s results tend to support these conclusions, and we applaud them for moving the discussion forward so productively. Complacent appeals to intuition are no longer respectable options for proponents of the entailment thesis.

After presenting their data and drawing their critical conclusions, Murray et al. respond to several objections and tentatively propose a hypothesis to explain their findings. They propose that belief, but not knowledge, requires "conviction" or "mental assent." People are at least implicitly sensitive to this fact, which explains why they view certain agents as having knowledge but not belief. Part of Murray et al.'s motivation for this proposal is that, to them, the availability of information to the mind is "intuitively" insufficient for belief. Murray et al. don't commit to a specific account of what assent requires, but they mention a few possibilities, including "a particular type of phenomenology—a subjective feeling of conviction."¹⁵

This strikes us as a potentially promising explanation of the extant findings used to question the entailment thesis. The attentive reader will have already noted how much the "conviction account" of belief resembles our earlier characterization of *thick belief*. We return to this in the next section.

3. Interlude

Each set of studies discussed thus far teaches a valuable lesson. First, Myers-Schulz and Schwitzgebel woke contemporary epistemology from its dogmatic slumber about the entailment thesis by offering provisional evidence that the anti-entailment intuition is neither peculiar nor perverse. Second, Rose and Schaffer emphasized

the importance of properly framing the entailment thesis and ensuring that experimental stimuli cue the appropriate sense of belief; otherwise, we're in danger of misinterpreting the results. Third, Murray, Sytsma and Livengood renewed the experimental challenge to the entailment thesis and proposed the conviction account of belief to explain why the entailment thesis is false.

We are no longer slumbering dogmatically; we are sensitive to the need to appropriately frame the entailment thesis; and we are alert to the importance of conviction to (the folk psychology of) belief. In short, we have taken these lessons to heart. In what follows, we seek to advance the discussion in light of them.

4. A Thinly Veiled Thesis

Recall that Rose and Schaffer distinguished dispositional from occurrent belief and argued that we should understand the entailment thesis dispositionally. This enabled them to reverse previous results which seemingly threatened the entailment thesis. In addition to distinguishing occurrent from dispositional belief, we will now apply a similar maneuver in response to the latest results from Murray et al. by further distinguishing between thick and thin belief.

Recall that 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. By contrast, thick belief requires more than a cognitive pro-attitude. It also involves *emotion* or *conation*. Thick and thin beliefs can be dispositional or occurrent in nature. Thick belief entails thin belief, but not vice versa.

We propose that the following thesis is true and reflects the way epistemologists have understood the relationship between knowledge and belief:

Thin entailment thesis: necessarily, if you know that P, then you thinly believe that P.

Following Rose and Schaffer, we also want to allow for a dispositional reading of thin belief. Thus, the thin entailment thesis should also be understood in terms of implying thin occurrent *or* thin dispositional belief.

We think that the thin entailment thesis is true and that it can be supported by two simple arguments. Here is the first argument. Knowledge itself is a cognitive pro-attitude. That is, to know that P is a way of regarding or taking P to be true. (Put another way: *knowing that P is true* is a way of *thinking that P is true*.) Therefore knowledge is a form of thin belief. Therefore the thin entailment thesis is true.

Here is the second argument. Suppose a speaker makes an assertion. As many have argued in the norms of assertion literature, if we respond by saying either "You don't know that," or "That isn't true," or "You don't believe that", then we are interpreted as rejecting the speaker's authority to make the initial assertion (Unger 1975, Bach and Harnish 1979, Williamson 2000, Weiner 2005). It's a very interesting fact that all three of these possible responses function this way. What explains this very interesting fact? The entailment thesis follows from the best explanation: (i) knowledge is the norm of assertion (Williamson 2000: ch. 11, Turri 2010, Turri 2011b, Turri 2013b), and (ii) knowledge entails true belief, so (iii) by challenging either knowledge, truth or belief, one thereby challenges authority to

assert (Turri 2010, Turri forthcoming). An unparsimonious alternative, which we find implausible, is that knowledge, truth and belief all function independently as norms of assertion.¹⁶

We think that the thin entailment thesis reflects the way epistemologists have understood the relationship between knowledge and belief because of passages like the following. Consider how Richard Feldman, a leading epistemologist, explains the entailment thesis in his recent textbook:

If you know something, then you must believe it or accept it. If you do not even think that something is true, then you do not know it. We are using ‘belief’ in a broad sense here: anytime you take something to be true, you believe it. (Feldman 2003: 13)

Feldman’s “broad sense” of belief is thin belief. Consider also how Linda Zagzebski, another leading epistemologist, explains some points of “philosophical consensus” about knowledge in her recent textbook:

Knowing is a relation between a conscious subject and an object, where the object is . . . some portion of reality. The relation is cognitive. That is to say, the subject thinks, not just senses or feels the object. More specifically, knowing includes believing.¹⁷ (2009: 3)

Zagzebski also adds, “there is a consensus that knowing is a form of believing.”

Notice that Feldman and Zagzebski both speak of *thinking that something is true* as a way of elucidating the relationship between knowledge and belief. It’s natural to construe the relevant sense of belief in terms of ‘think’.¹⁸ This will be important below.

5. Thicker Plot, Thinner Probe

The take-home lesson of this section is that the thin entailment thesis is widely accepted. But our goal is not to dispute any of the prior anti-entailment results reported in Sections 2.1. or 2.3. Setting aside the occurrent/dispositional distinction made by Rose & Schaffer, we think that both Myers-Schulz and Schwitzgebel, and Murray et al., have admirably shown that knowledge may not entail (thick) belief. We now wish to add to these prior discoveries the further claim that knowledge does entail (thin) belief—the sense of belief that some cognitive scientists and contemporary epistemologists have endorsed—together with evidence for these two distinct folk psychological categories of belief.

More specifically, we will argue for two main points. First, participants in Murray et al.’s studies tended to ascribe knowledge without belief mainly because they understood “belief” in the thick sense, while the storyline didn’t make clear that the protagonist had a thick belief. Accordingly, their results don’t suggest that the thin entailment thesis is widely rejected. Second, participants who understand “belief” in the thin sense will respond just as we would expect if the thin entailment thesis is widely accepted. If the thin entailment thesis is widely accepted and participants understand “belief” in the thin sense, then *more participants will ascribe belief than knowledge*,¹⁹ and *very few participants will both ascribe knowledge and deny belief*. (We will operationalize ‘very few’ in Section 5.1. below.)

To demonstrate our first point, we will keep Murray et al.'s original probes but adjust their vignettes to make it more plausible that the protagonist has a thick belief. Call this *the plot-thickening manipulation*. We predict that a significant increase in belief ascription will accompany the plot-thickening manipulation. To demonstrate our second point, we will keep Murray et al.'s original vignettes but adjust the probes so it's clear that we're asking about thin belief. Call this *the probe-thinning manipulation*. We predict that three things will accompany the probe-thinning manipulation: (1) belief ascription will rise significantly, (2) belief ascription will surpass knowledge ascription, and (3) the percentage of participants who both ascribe knowledge and deny belief will decrease significantly. Indeed, regarding (3) we expect this percentage to be at or near floor (again, we operationalize this notion in Section 5.1. below).

Of course by just thickening the plot on its own, one might naturally expect belief ascription to rise without the need to posit additional senses of belief.²⁰ But we present our plot-thickening and probe-thinning manipulations in tandem, as one pincer motion. That is, we will show that holding fixed the plot but thinning the probe results in greater belief ascription. And conversely, we will show that holding fixed the probe but thickening the plot also increases ascription. We present these two results together as evidence in support of both thick and thin senses of belief.

In making our points, we will focus on Murray et al.'s three most promising cases: Geocentrist, Dog and God.²¹

As indicated earlier, there could be many ways to "thicken" belief. Thus there could be many ways to augment the original stimuli so it's clearer to participants that the protagonist has a thick belief. For the sake of clarity and definiteness, choices must be made. In the experiments below, our strategy is to have the protagonists *endorse the relevant proposition* by explaining their choice or communicating the information to a social partner. Perhaps belief ascription in the presence of this kind of explicit endorsement is not a surprising result, or a result that would be unique to our view. Nonetheless it is a prediction of our account of thick and thin belief that this feature is more likely to cue thick belief ascriptions. Conversely, it would be problematic for our view if endorsement of this kind did not result in higher rates of thick belief ascription. Thus we include thick belief cases as a crucial test of the hypothesis. We designed a common plot-thickening manipulation for Geo and Dog in order to show that endorsing a proposition is a cue to thick-belief ascription across different cover stories.

There could be many ways to cue the thin sense of belief. Taking our cue from Feldman's and Zagzebski's natural use of 'think' in the passages quoted earlier, we asked participants whether the protagonist "thinks" that the relevant proposition is true. Moreover, bearing in mind Rose and Schaffer's findings, we wanted to allow for a dispositional reading of the question. To accomplish this, we included the adverbial phrase 'at least on some level'.²² Combining those two desiderata, we asked participants, "At least on some level, does S think that P?", appropriately substituting for 'S' and 'P'.

Having explained our strategic choices, we now turn to the experiments. Let's thicken the plot, thin the probe, and then see what happens.

5.1. Experiment 1: Geocentrist Case

5.1.1. *Method.* Participants ($N = 329$)²³ were randomly assigned to one of three conditions. Participants in the Original Geo condition received Murray et al.'s original Geocentrist case:

[*ORIGINAL GEO*]²⁴ 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 homeschooled 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.

Participants in Original Geo were then asked two questions (options in brackets):²⁵

Does Karen know that the earth revolves around the sun? [Yes/No]

Does Karen believe that the earth revolves around the sun? [Yes/No]

Participants in the Thick Geo condition received the exact same treatment as Original Geo, except for one important difference: their vignette included an additional sentence designed to provide participants with evidence that Karen has a thick belief (i.e. has "conviction" or "assents to" the relevant proposition). To accomplish this in the addition below, Karen explicitly endorses the proposition that the earth revolves around the sun:

[*THICKENED GEO*] Later, when Karen explains her answer to her friends, she tells them, "Of course, I want to respect my parents, but I cannot deny the evidence we learned about in class that the earth revolves around the sun."

Participants in the Thin Geo condition read the original vignette, ORIGINAL GEO, but they answered a different question designed to cue thin belief, while leaving open the possibility of dispositional belief. These were their questions:

Does Karen know that the earth revolves around the sun? [Yes/No]

At least on some level, does Karen think that the earth revolves around the sun?
[Yes/No]

5.1.2. *Results and discussion.* We made four main predictions. First, rates of belief ascription in Thick Geo will be significantly higher than in Original Geo. Second, rates of belief ascription will be significantly higher in Thin Geo than in Original Geo. Third, in Thin Geo knowledge ascription will not be higher than belief ascription. Fourth, the percentage of participants who both ascribe knowledge and deny belief will be significantly lower in Thin Geo than in Original Geo. For simplicity, let's call participants who both ascribe knowledge and deny belief *the anti-entailment crowd*. The smaller the anti-entailment crowd in Thin Geo, the better for our view.

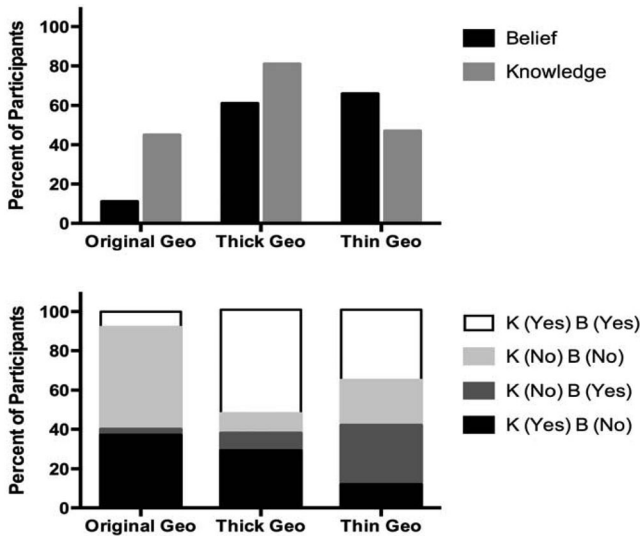


Figure 1. Results of Experiment 1. *Top panel:* percentage of participants in Original Geo, Thick Geo, and Thin Geo conditions attributing knowledge and belief. *Bottom panel:* overall response patterns; the K (Yes) B (No) pattern shows anti-entailment.

Previous studies have shown that approximately 10–20% of people are willing to ascribe knowledge of falsehoods (*the anti-factivity crowd*).²⁶ We will feel vindicated if the anti-entailment crowd isn't significantly larger than the anti-factivity crowd, which we stipulatively define at 15%.²⁷

All four predictions were true. Only 11% of participants in Original Geo ascribed belief, whereas 61% in Thick Geo and 66% in Thin Geo did.²⁸ In Thin Geo, belief ascription (66%) was significantly higher than knowledge ascription (47%).²⁹ The anti-entailment crowd was significantly smaller in Thin Geo (12%) than in Original Geo (37%).³⁰ Moreover, the anti-entailment crowd in Thin Geo (12%) was smaller than the 15% target. These results are represented in Figure 1.

To summarize, these results so far suggest that the protagonist in the original Geocentrist cases has a *thin* belief but not a *thick* belief about heliocentrism. They also suggest that when the original case is thickened by having the protagonist explicitly endorse the heliocentric proposition to peers, this cues ascriptions of thick belief regarding the heliocentric proposition.³¹ Lastly, since the entailment thesis is best understood in terms of thin belief, these results do not constitute a challenge to traditional entailment.

5.1.3. Objection and follow-up: matching knowledge probe. Before continuing, we would like to pause and address an initial concern with intuitions in Thin Geo. We take it that results for our probe-thinning manipulation begin to show that people are using a different sense of belief in Thin Geo than in Original or Thick Geo. However for the purposes of directly testing *thin belief entailment* specifically, a more challenging test might be to compare ascription rates for “at least on some

level thinks that P” to rates for “at least on some level knows that P”³² All the better for our defense of the thin entailment thesis if, even in this alternative test, people continue to ascribe knowledge at rates that do not surpass belief.

To find out, we ran a follow-up experiment. We re-administered the Thin Geo case verbatim, except this time we presented participants ($N = 64$) with the following two questions:³³

At least on some level, does Karen know that the earth revolves around the sun?
[Yes/No]

At least on some level, does Karen think that the earth revolves around the sun?
[Yes/No]

Here we found that response patterns mirrored those of our earlier study. The majority of participants ascribed both thin belief (81%) and knowledge (80%). And only 8% of participants ascribed knowledge without belief, which was marginally below the 15% anti-entailment threshold.³⁴ This result suggests that thin belief entailment continues to be supported even when knowledge probes feature matching “at least on some level” phrasing.

5.2 Experiment 2: Dog case

5.2.1. *Method.* Participants ($N = 322$)³⁵ were randomly assigned to one of three groups. Participants in the Original Dog condition received Murray et al.’s original dog case:³⁶

[ORIGINAL DOG] Researchers have found that some breeds of dog are surprisingly intelligent. Amongst the most intelligent are border collies. One border collie, Cassie, is even able to solve simple mathematical problems: If you ask her “what is $2 + 2$,” for example, she will bark four times; similarly, if you ask her “what is $4 + 5$ ” she’ll bark nine times!

Participants in Original Dog were then asked two questions:

Does Cassie know that $2 + 2 = 4$? [Yes/No]

Does Cassie believe that $2 + 2 = 4$? [Yes/No]

Participants in the Thick Dog condition received the exact same treatment as Original Dog, except for one important difference: their vignette included two additional sentences designed to provide participants with evidence that Cassie has a thick belief. To accomplish this in the addition below, Cassie explicitly endorses the proposition that $2 + 2 = 4$:

[THICKENED DOG] In addition to all this, when other dogs are asked the answers to simple math problems but they are unable to respond correctly, Cassie will bark the correct answers in order to help the other dogs solve the problems. It’s like she’s really trying to communicate the answer to them!

Participants in the Thin Dog condition read the original vignette, ORIGINAL DOG, but they answered a different question designed to cue the category of thin

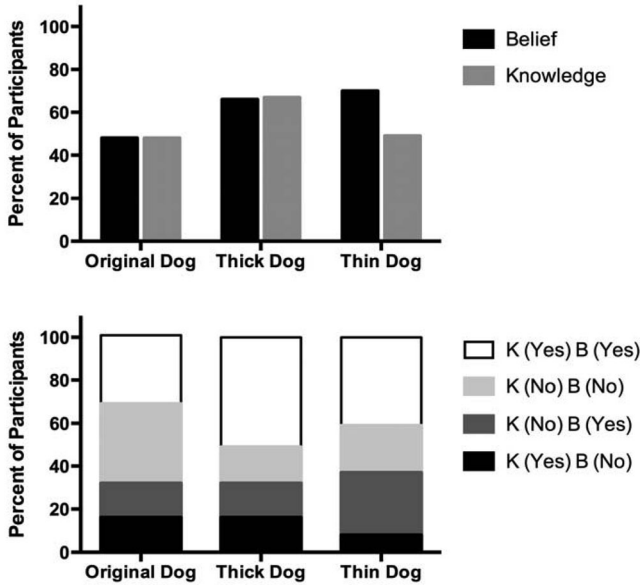


Figure 2. Results of Experiment 2. *Top panel:* Percentage of participants in Original Dog, Thick Dog, and Thin Dog conditions attributing knowledge and belief. *Bottom panel:* overall response patterns; the K (Yes) B (No) pattern shows anti-entailment.

belief, while leaving open the possibility of dispositional belief. These were their questions:

Does Cassie know that $2 + 2 = 4$? [Yes/No]

At least on some level, does Cassie think that $2 + 2 = 4$? [Yes/No]

5.2.2. Results and discussion. We made four main predictions. First, rates of belief ascription in Thick Dog will be significantly higher than in Original Dog. Second, rates of belief ascription will be significantly higher in Thin Dog than in Original Dog. Third, in Thin Dog knowledge ascription will not be higher than belief ascription. Fourth, the anti-entailment crowd will be significantly smaller in Thin Dog than in Original Dog and not significantly higher than 15%.

All four predictions were true. Whereas 48% of participants in Original Dog ascribed belief, 66% in Thick Dog and 70% in Thin Dog did.³⁷ In Thin Dog, belief ascription (70%) was significantly higher than knowledge ascription (49%).³⁸ The anti-entailment crowd was smaller in Thin Dog (8%) than in Original Dog (16%) by a marginally significant amount.³⁹ (Note that the anti-entailment crowd in Original Dog in our study was significantly smaller than in Murray et al.'s original study (28%).)⁴⁰ Moreover, the anti-entailment crowd in Thin Dog was significantly *less* than 15%.⁴¹ These results are displayed in Figure 2.

To summarize, these results suggest that Cassie, in the original Dog cases, has a *thin* belief but not a *thick* belief about the sums. They also suggest that when

the original case is thickened by having Cassie explicitly endorse the proposition to peers, this cues ascriptions of thick belief. And again, since the entailment thesis is best understood in terms of thin belief, these results do not constitute a challenge to entailment.

5.3 Experiment 3: God case

5.3.1. Method. For our final experiment we revisited Murray et al.'s God case, limiting our investigation to the response pattern associated with the thin belief probe.⁴² Participants ($N = 102$)⁴³ were asked two questions:

Does God know that $2 + 2 = 4$? [Yes/No]

At least on some level, does God think that $2 + 2 = 4$? [Yes/No]

Participants were then taken to another screen (they couldn't go back and change their answer), where they were asked:

Do you believe in God? [Yes/No]

5.3.2. Results. We made two predictions. First, knowledge ascriptions would not be significantly higher than belief ascriptions. Second, the anti-entailment crowd would not be significantly larger than 15%. Both predictions were true. Knowledge ascription was not significantly higher than belief ascription—indeed, they were exactly the same (61%).⁴⁴ And the anti-entailment crowd (7%) was significantly smaller than 15%.⁴⁵

We found that personal belief in God predicted significantly higher rates of both belief ascription (90/34%) and knowledge ascription (94/30%).⁴⁶ This is unsurprising because atheists don't think that God exists, and it's reasonable to suppose that only existing beings believe and know things. Thus atheists should be less likely to ascribe belief and knowledge. Nevertheless, even when we control for personal belief in God, we detect no threat to the thin entailment thesis. Neither theists nor atheists ascribed knowledge at rates significantly higher than belief.⁴⁷ Even theists don't come close to a significantly higher rate of knowledge ascription.

Murray et al.'s original study of the God case used a between-subjects design, whereas we used a within-subjects design. So a direct comparison across the two studies is fraught. Nevertheless, it's at least worth noting that theists in our study were significantly more likely than participants in Murray et al.'s study to ascribe belief to God (90/73%),⁴⁸ even though rates of knowledge ascription among theists in the two studies didn't differ significantly.⁴⁹ These results are shown in Figure 3.

5.4 General discussion of Experiments 1–3

A clear pattern emerges from the three experiments. First, it is probable that Murray et al.'s original results are partly due to the fact their probes cued the thick sense of belief, while their stories didn't make it clear that the protagonist had a thick belief. By thickening the plot, we were able to significantly increase rates of belief ascription. Second, whenever we cued the thin sense of belief, participants respond just as we would expect if the thin entailment thesis is widely accepted. That is,

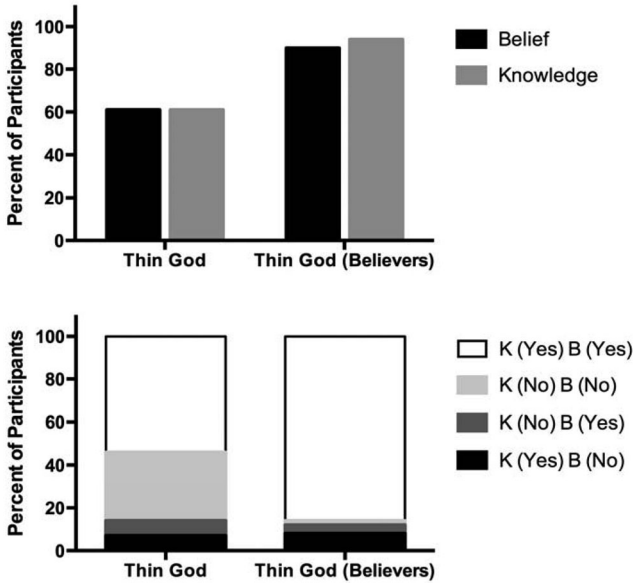


Figure 3. Results of Experiment 3. Percentage of participants in Thin God and Thin God (Believers, filtered for a personal belief in God) conditions attributing knowledge and belief top; overall response patterns below. The K (Yes) B (No) pattern shows anti-entailment.

when we thin the probe, participants consistently ascribe belief at higher rates than knowledge and very few join the anti-entailment crowd (i.e. ascribe knowledge but deny belief).

All of our results suggest that the thin entailment thesis is widely accepted. Indeed, it appears to be as firmly rooted as the factivity thesis in the ordinary view of knowledge. That is, *thin belief* is as good a candidate as *truth* for being a folk requirement on knowledge!⁵⁰

Recall that although our aim in the present section was to defend the (thin) entailment thesis, the larger goal of this paper has been to present evidence for the thin and thick categories of belief. The debate over ordinary intuitions and entailment has merely served as a convenient point of departure toward achieving this latter goal. So do the results taken as a whole support this conclusion? Though more empirical work needs to be completed, we think that the experiments above begin to demonstrate a real distinction that should be made in both cognitive science and epistemology between thick and thin categories of belief. We now turn to examine the evidence in light of this larger question, as well as consider one major objection to this claim.

For evidence in favor of the larger claim regarding thick and thin belief more generally, we conducted a post hoc test which collapses over Original and Thick conditions in the Geo and Dog experiments above and compares combined belief and knowledge scores with ascriptions across all of the Thin conditions in Geo,

Dog, and God. We choose to collapse responses in this way in order to compare the results of all belief ascription (using the unadorned word ‘belief’ in our probes) with thin belief ascription (or using the adorned ‘at least on some level thinks’ method of collecting *thin* belief ascription).

When looking at these results as a whole, we continued to see significant differences supporting the claim that these are genuine and genuinely distinct categories. Specifically, 46% of participants attributed belief across all of the Original and Thick conditions combined, whereas 65% attributed belief across all of the combined Thin conditions combined. And we saw the opposite pattern in knowledge ascription across all of the Original and Thick conditions (60%) when compared to the Thin conditions (52%).⁵¹

We also continued to find significant differences in how individual participants within each of these combined conditions ascribed knowledge and belief. There was significantly more knowledge ascription than belief ascription in Original and Thick conditions, and significantly more belief ascription than knowledge ascription in Thin conditions.⁵²

5.5 Experiment 4: *Within-subjects replications*

In our discussion of Myers-Schulz and Schwitzgebel’s original studies against the entailment thesis we mentioned one response by Rose and Schaffer (Section 2.2.) suggesting that a *within-subjects* design would have been more probative than a between-subjects design. This is why in our experiments on entailment (Sections 5.1.–5.3) we asked participants both questions about either thick or thin belief together with questions about knowledge, to see whether the same person would attribute knowledge but deny thick or thin belief. But what about thick and thin belief themselves? Will evidence for these distinct categories persist when the same person is asked about *both* categories of belief?

As an added dimension of support in favor of the distinction between thick and thin belief in folk psychology, we sought to replicate the effects for thick and thin belief using the Geocentrist and Dog cases from Experiments 1 and 2 *within subjects*, or among the very same groups of participants. In this follow-up experiment, participants ($N = 100$)⁵³ were presented with either the Original Geo or Original Dog vignette, followed by verbatim thick and thin belief probes used above (see appendix for materials).

In both cases, the result was a replication of the earlier findings showing significant differences in thin and thick belief attributions.⁵⁴ Thin belief ascription was at 52% in Original Geo compared to 20% thick belief ascription. Similarly, 68% ascribed thin belief in Original Dog as compared to 46% ascribing thick belief. These results are shown in Figure 4.

5.6 Experiment 5: *The limits of thin belief*

One potential worry one might have is that our strategy for cuing thin beliefs across Experiments 1–3 does not rise to the minimal level sufficient for belief. In fact, a version of this same basic worry was anticipated by Myers-Schulz and Schwitzgebel in their original paper. When speaking of some additional experiments they had




Figure 4. Results of within-subjects experiments. Percentage of participants in Original Geo and Original Dog attributing thin belief and thick belief.

conducted using ‘thinks’ as an alternative measure for belief ascription (see footnote 18) they write, “We hypothesize that some respondents may have interpreted ‘think’ in these scenarios as something like guess or suspect rather than believe.” It would be problematic for our account of thin belief if at least one version of Myers-Schulz and Schwitzgebel’s hypothesis turns out to be correct, namely that people who *randomly, equivocally, or blindly guess that P* are viewed as “at least on some level, thinking that P.” In other words, it would be problematic if people attributed thin belief using our metric in a maximally inclusive and indiscriminate way.

To begin addressing this worry about the minimum level for thin belief, we conducted a follow-up study. Participants ($N = 107$)⁵⁵ were presented with a new case very similar to the original Unconfident Examinee first administered by Myers-Schulz and Schwitzgebel. However in our case the protagonist, Sloan, was not a model student like Kate:

Sloan is a real slacker. He enrolled in Astronomy 101 this semester but he never attended a single class and he never did any of the assigned reading. Sloan is completely unfamiliar with the course material, but the final exam is multiple-choice and he’s hoping that he can get lucky and pass. Question 72 on the exam asks, “Which of the following best approximates the value of the gravitational constant?” The options are:

- A. $6.7 \times 10^{-21} \text{ m}^3/\text{kg s}^2$
- B. $6.7 \times 10^{-11} \text{ m}^3/\text{kg s}^2$
- C. $6.7 \times 10^{-12} \text{ m}^2/\text{kg s}^3$
- D. $6.7 \times 10^{-11} \text{ m}^2/\text{kg s}^3$
- E. $6.7 \times 10^{-10} \text{ m}^3/\text{kg s}^2$

Sloan randomly selects option B, which is the correct answer.

Participants then saw one of three questions to gauge thick belief, thin belief, and knowledge ascription:

Does Sloan believe that $6.7 \times 10^{-11} \text{ m}^3/\text{kg s}^2$ is the best approximation?

At least on some level, does Sloan think that $6.7 \times 10^{-11} \text{ m}^3/\text{kg s}^2$ is the best approximation?

At least on some level, does Sloan know that $6.7 \times 10^{-11} \text{ m}^3/\text{kg s}^2$ is the best approximation?

If our method for collecting thin belief does not rise to a minimal level necessary for belief—but rather resembles something akin to mere guessing—then we would expect high thin belief attribution to Sloan. If this is the case then we would also expect thin belief ascription at rates that significantly exceed ascriptions of thick belief and knowledge. In other words, if thin belief is ascribed more indiscriminately than thick belief in situations of mere guessing, we would expect higher thin belief ascriptions relative to rates of thick belief ascription.

Contrary to these predictions, however, we found that participants were generally unwilling to ascribe thin belief, thick belief and knowledge to Sloan. Responses to all three questions were significantly below chance.⁵⁶ Only 28% of participants in this case ascribed thin belief, or answered that “at least on some level” Sloan thinks that his answer is correct. And most crucially, thin belief was not ascribed at significantly higher rates than either thick belief (23%)⁵⁷ or knowledge (17%).⁵⁸

These results suggest that our thin probe does cue minimal features sufficiently rising to the level of belief. That is, these results provide evidence that the folk psychological category of thin belief requires more than guessing. Some may find this result surprising. After all, it might be thought, “doesn’t someone who randomly guesses that P represent P as true?” But the vast majority of participants did not attribute thin belief. In fact, participants ascribed thin belief at rates similar to thick belief and knowledge.⁵⁹

Despite these results, we also acknowledge that more empirical work than space currently allows is needed to investigate the lower bounds of the folk psychological category of thin belief. For instance, even though the vast majority of participants did not attribute either thick or thin belief to Sloan in the control above, it’s still puzzling that 23–28% of participants attributed belief. Further work could profitably investigate why this proportion of people indiscriminately ascribes belief, both thin and thick.

While this is an interesting question, it’s important to emphasize that these results do not detract from our argument for thick and thin categories of belief. Nor do these results suggest that a rump of indiscriminate thin-belief ascription can explain our earlier results. For instance, the Geocentrist case (Section 5.1.) displays a very large shift from the percent of participants who said that “Karen believes that the earth revolves around the sun” in the original condition (11%) to the percent who said that she “at least on some level thinks” in the Thin Geo condition (66%). It’s implausible that that difference could be explained in terms of participants allowing guessing to count as thin belief. For in order for such an explanation to work, participants would need to indiscriminately apply thin belief *but not thick belief*. However, our follow-up results suggest that thin and thick belief are indiscriminately ascribed at the level of guessing at similarly low rates.

We would also like to point out that independent of these discussions, the minimal pro-attitude necessary for belief is an interesting philosophical question in its

own right, itself open for debate. Related to this, note that there are more or less discriminating categories of guessing. For example, some of the more discriminating or less inclusive concepts of guess are ‘conjecture’, ‘extrapolation’, or ‘educated assumption’. To use our terminology, these attitudes might well suffice to represent information as true in a way we suspect many would find sufficient for thin belief. For more on this question, and for additional experiments on the lower bounds of thin belief, see additional work by Buckwalter and Turri (under review a) on the thick/thin distinction.

6. Conclusion

We interpret our results to reflect a genuine distinction between thin and thick belief in folk psychology. Once these two categories of belief are properly accounted and controlled for, the way is clear to a proper understanding and defense of the entailment thesis. The entailment thesis should be understood thinly, and when participants are effectively probed, they reveal a strong commitment to that thesis. Epistemologists can now enjoy an empirically informed justification for including a belief condition in their theories of knowledge. The thick/thin distinction enabled this impressive result. We humbly submit that this is an adequate downpayment on our hypothesis.

We will conclude by briefly suggesting some further, potentially important applications of the thick/thin distinction. Earlier we observed that discussions in contemporary philosophy of mind and cognitive science have typically focused on thin belief. More often than not in these discussions, thin beliefs are treated as information about the world that is internally represented and stored in the mind. But although it seems possible to have a partial belief or a full belief of many different strengths, mental representations appear to be discrete—you either have one in your mind or you don’t. So we take it that one historical challenge to traditional representational views of belief (as seen in theories ranging from Fodor 1990 to Searle 1983) is that they must somehow account for the continuous nature of belief.

Noting the apparent difficulties that representational theories have with partial or degrees of belief, one might reject the traditional representational view and develop an entirely new account of belief (for this see Schwitzgebel 2001 and 2002).⁶⁰ However, rejecting the traditional representation view of belief on the ground that it cannot accommodate the continuous nature of belief might be premature. For the traditional representational picture can be aided by distinguishing two categories of belief: thick and thin belief. Drawing on these two categories, one might make a plausible case that thin beliefs, whether occurrent or dispositional, are associated with discrete representations, while thick beliefs have additional properties that are continuous. And we suspect that the features that cue intuitions about partial belief and degrees of confidence along this continuum—things like the associated phenomenology of belief, or one’s willingness to assert or endorse a belief as the basis for action—will be features that typically also cue judgments about thick beliefs. If

correct, this distinction may potentially aid representationalists in accounting for intuitive features of belief like the possibility of partial or in-between belief.

Thinking thick and thin about the nature of belief may also contribute to ongoing debates regarding doxastic norms and the ethics of belief formation. For instance, one such debate involves the question of whether beliefs are formed voluntarily or involuntarily, and subsequently, how this question bears on whether agents can be held normatively responsible for forming them. A number of different positions have been defended in this literature (ranging from Pascal 1670 to Weatherston 2008) regarding what has come to be known as *the doxastic involuntarist challenge*. In short, the challenge is to explain how one can be held responsible for an action (forming a belief) that is beyond one's control.

Here too the thick/thin belief distinction may be of service. For it seems that thin beliefs are typically formed involuntarily, on the basis of perceiving and representing information about the world. Thick beliefs, on the other hand, are typically formed voluntarily, on the basis of more reflective cognitive pro-attitudes, conation, or endorsement. If so, then distinguishing between beliefs may allow philosophers and cognitive scientists to retain the basic idea that (thin) belief formation is predominantly involuntary, while also clearing the way for the normative evaluation of (thick) beliefs.⁶¹

Lastly, applying the thick/thin distinction to questions in the ethics of belief formation may lead to applications in moral philosophy more generally, where moral agents' beliefs often feature prominently. For example, consider the ongoing debate between motivation externalists and motivational judgment internalists in metaethics (e.g. Stocker 1979, Korsgaard 1986, Smith 1994). Internalists contend, roughly, that sincere moral judgment entails a motivation to act in accordance with that moral judgment, while externalists deny this. Although the connection between moral judgment and motivation has been spelled out a number of different ways among motivational externalists and internalists, it is frequently done so in terms of either *belief* that a certain moral proposition is true, or the motivation that such a *belief* purportedly entails.

Though we have now traveled some distance from entailment, we think that drawing on the thick/thin categories of belief may help resolve disagreements between motivational externalists and internalists, insofar as the connection between moral judgments and motivation hinges, in part, on substantive views of belief. For instance, if motivational judgment internalists cast their view in terms of thick beliefs, it seems very plausible to conclude that endorsement or conviction that P would entail motivation to act in accordance to P. Alternatively, if the notion of belief at play is thin belief, then motivational externalism seems quite plausible. Simply representing information as being the case, or believing in the thin sense, does not seem to require any further conviction or motivation to act in accordance with P (for more on this application of thick and thin belief see Buckwalter and Turri under review a). We present this and the other examples above not as an argument for one of these views, but simply as a further illustration of how clarifying and applying these features of (thick and thin) belief might prove crucial to ongoing theoretical arguments both within traditional epistemology, and well beyond.⁶²

Notes

¹ This case has also been instrumental in motivating intuitions in debates about moral motivation, see Stocker 1979, Buckwalter and Turri under review a.

² This is an initial approximation of the category of thin belief for purposes of introducing the discussion. In Section 5 we discuss some considerations which might lead to strengthening the minimal requirements of thin belief. For now, and the essential point we wish to uphold is that whatever the minimum requirements of thin belief beyond mental representation, they clearly demand less than the minimum requirements of thick belief. We thank an anonymous reviewer for encouraging us to clarify this point.

³ For more on the different ways that thin beliefs can be thickened, see Section 5.

⁴ Also see theoretical work on what the “fragmentation” of concepts might mean for a series of epistemic debates by Greco forthcoming.

⁵ Dretske 1988 also focuses on the explanatory role of belief in thought and behaviour. On standard representationalist models of cognition, it might seem unusual to say that very simple organisms or computers representing information that P “believe” P is true, or that the contents of the early visual system are “believed” to be true. On the other hand, we suggest that it is much less unusual to say that these entities “on some level think” that P. In fact, this comports well with the motivation for our choice of terminology to cue the category of thin belief below. We thank an anonymous reviewer for discussion on this point.

⁶ See Millgram 1995 and Sayre-McCord 2006 for more on Hume’s theory of belief and its potential mismatch to some of the things latter-day “Humeans” hold about belief. We thank an anonymous reviewer for encouraging us to clarify this point. For related views, see James 1889: 21 and Peirce 1958, 1976. James says, “In its inner nature belief, or the sense of reality, is a sort of feeling more allied to the emotions than to anything else,” and, “The emotion of belief” is “perfectly distinct, but perfectly indescribable in words.” Peirce says that belief is a kind of “controlled and contented habit of acting.”

⁷ For some recent controversy surrounding the factivity thesis, see Hazlett 2010 and 2012, Turri 2011a, Stokke 2013, Buckwalter and Turri under review a, and Buckwalter under review.

⁸ See Travis 2005 for a discussion of a particularly compelling rejection of the entailment thesis by Cook Wilson. We thank an anonymous reviewer for this reference. Also see Section 4 for some evidence suggesting that epistemologists have tended to elucidate the relationship between knowledge and belief in terms of a thin category of belief.

⁹ Or, at least, they challenge the view that the entailment thesis is either widely held, intuitively compelling to most people, or a principle of folk psychology. We acknowledge the difference between this challenge and challenging the entailment thesis *per se*, even though the two challenges aren’t unrelated. Our main argument doesn’t hinge on always explicitly distinguishing the two challenges, so for expository purposes we won’t always explicitly distinguish them in what follows.

¹⁰ In their experiments, Rose and Schaffer treat judgments about ‘what one is disposed’ to believe as one proxy for measuring what one dispositionally believes. Though for a possible distinction between these two things see Audi 1994.

¹¹ To be clear, Myers-Schulz & Schwitzgebel do not claim that a majority of people exhibit anti-entailment intuitions, only a substantial minority. We invite readers to decide for themselves if the 30% in Rose & Schaffer’s experiment on dispositional belief constitute such a minority.

¹² Myers-Schulz & Schwitzgebel also present a *Prejudiced Professor* case involving a well-intentioned liberal professor who harbours implicit biases towards certain students. They find that 63% of experimental participants agree that the professor knows that these students are just as capable as her other students, but only 23% agreed that she believes it. In response, Rose & Schaffer argue that this finding is an instance of a larger phenomenon in experimental philosophy known as the *epistemic side-effect effect* (see Beebe & Buckwalter 2010, Buckwalter 2013) whereby moral considerations can influence ascriptions of knowledge and belief (see Beebe forthcoming). However, the best explanation of this effect remains extremely controversial, since it is unclear whether this effect is due to a motivational bias (Alicke and Rose 2010) or whether it is a feature of our competence (Knobe 2010, Knobe, Prasada & Newman 2013, Szabó & Knobe 2013). Since the proper interpretation of these effects is still an open question

(see Knobe 2010 and Knobe et al. 2012 for further issues surrounding the proper interpretation of these effects), we omit this case from our discussion due to its entanglement with these larger issues.

¹³ *Non-triviality* was stipulated by Murray et al. to be average response patterns that were significantly greater than 20%, where responses that were not significantly different from 20% were interpreted to reflect mere “noise.”

¹⁴ Our focus is on the Geocentrist case in part because it has been offered by others as one of the strongest cases against entailment. Note that there are a number of worries about this case independent of our hypothesis. For one, there is a risk that the story primes a specifically religious sense of “belief” (though see below). For another, the vignette includes belief terms like ‘accepts’ and ‘holds’. Also, knowledge ascription could be artificially high due to the basic nature of the claim, or the assumption that “*everyone* knows the earth revolves around the sun”. We acknowledge these worries but set them aside in the ensuing discussion.

¹⁵ They also mention two other possible views about what assent requires. On one view, assent requires “believing a proposition on the basis of deliberation regarding its justification.” On the other view, assent requires “explicitly *reflective* deliberation.” Murray et al. don’t note it, but Augustine arguably held a related view of belief, what we might call *the deliberative view*. In *On the Predestination of the Saints* (bk. 1, ch. 5), Augustine wrote, “No one believes anything unless he has first thought that it is to be believed. For however suddenly, however rapidly, some thoughts fly before the will to believe . . . it is yet necessary that everything which is believed should be believed after thought has preceded.” We think it highly unlikely that any such deliberative view of belief is either correct or a genuine principle of folk psychology, though it’s arguably less implausible as a view about the requirements of *appropriate* belief. In any event, we set it aside.

¹⁶ Technically, there are two other alternatives that are parsimonious but far more implausible, namely: (a) belief is the norm of assertion and belief entails truth and knowledge; (b) truth is the norm of assertion and truth entails belief and knowledge. We assume that no one will take either of these alternatives seriously. For more on the connection between belief and assertion see Rose, Buckwalter and Turri under review.

¹⁷ We altered this quote by turning a numbered list into a simple paragraph.

¹⁸ Myers-Schulz and Schwitzgebel 2013 also seem to acknowledge this point. They speculate that “‘think’ is often used in ordinary English to ascribe what philosophers would call beliefs,” and proceed to use ‘think’ in additional experiments as an alternate way of measuring belief ascription, with mixed results.

¹⁹ Technically, this consequent clause should read, ‘then the percentage of participants who ascribe knowledge won’t be significant higher than the percentage of participants who ascribe belief.’ But since it’s harder to have knowledge than to have a belief, it’s natural to expect that belief ascription will outstrip knowledge ascription, which is how we frame matters in the main text.

²⁰ We thank Josh Knobe for discussion on this point.

²¹ We omit their cash register case because it raises deep questions about the nature of cognition and mindedness, which we can’t properly address in the present context. Relatedly, it also raises interpretative questions about how participants view mental-state ascriptions, and questions about such ascriptions, when it comes to machines and instruments.

²² This phrase also has one added bonus: it clears the way for ascribing *unconscious* belief too. (If Freud is correct, then unconscious belief can be occurrent and thick (indeed, *very* thick), and unconscious belief is not merely a species of thin belief.) The conscious/unconscious distinction hasn’t featured prominently in the recent controversy over the entailment thesis. But, for the record, we want to allow that thin, dispositional, unconscious belief suffices to satisfy the belief condition on knowledge.

²³ One-hundred three female, aged 18–65, $M = 27$ years, $SD = 8.25$ years. As with the experiments reported below, participants were recruited using Amazon Mechanical Turk, tested online using Qualtrics survey software, and compensated \$.25 for approximately 2–3 minutes 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 throughout the United States. Ninety-seven percent reported English as their native language. They filled out a brief demographic survey after testing.

²⁴ We use caps to name narrative elements, and we often name narrative elements after the experimental conditions they were used in. This eases exposition and helps readers keep track of which stories appeared in which conditions, while avoiding any undesirable confusion between the experimental conditions and the stories. (Participants never saw the labels.)

²⁵ In all experiments reported here, questions and options were rotated randomly.

²⁶ Starmans and Friedman 2012, Turri 2013a, Buckwalter under review.

²⁷ Recall that Murray et al. stipulated anti-entailment responses falling at or below 20% should be interpreted as mere “noise” (see footnote 6). We have lowered the bar thereby providing a more *liberal* standard for anti-entailment.

²⁸ Binomial, both $ps < 0.001$. For all reported inferential statistics: alpha level = 0.05, 95% confidence intervals, and all tests two-tailed, unless otherwise noted. As a control we also included the question “do you consider yourself to be a religious person?” We found no evidence of an association between answers to this question and participant ascriptions of belief (Fisher’s, $p = 0.51$) or knowledge (Fisher’s, $p = 0.67$) in Original Geo. Thus it seems unlikely that the reluctance to ascribe belief is the result of the application of at least any particular religious sense of belief.

²⁹ McNemar’s, $p = 0.018$, one-tailed.

³⁰ Fisher’s exact, $p < 0.001$, one-tailed. Overall response patterns for Original Geo, Thick Geo, and Thin Geo, respectively, were as follows: No knowledge no belief (52%, 10%, 23%), no knowledge belief (3%, 9%, 30%) knowledge no belief (37%, 29%, 12%), knowledge and belief (8%, 53%, 36%). Note roughly the same basic effect for anti-entailment (37%) as found by Murray et al. (46%).

³¹ Increases in knowledge ascription in Thick Geo may also suggest a relationship between endorsement and knowledge judgments to be pursued in future work.

³² We thank Josh Knobe for discussion on this point.

³³ Twenty-nine female, aged 18–62, $M = 29.9$ years, $SD = 9.13$ years. Questions were counterbalanced for order. There were no significant order effects on attributions of knowledge (Fisher’s exact test, $p = 0.54$) or thin belief (Fisher’s exact test, $p = 0.75$).

³⁴ Binomial, $p = 0.068$, one-tailed. Overall response patterns were as follows: No knowledge no belief (11%), no knowledge belief (9%) knowledge no belief (8%), knowledge and belief (72%).

³⁵ One hundred twenty-two female, aged 18–62, $M = 28$ years, $SD = 8.33$ years.

³⁶ Some philosophers may argue that dogs cannot have beliefs. But insofar as this case is presented as evidence for folk anti-entailment intuitions, we set this worry aside.

³⁷ Binomial, both $ps < 0.001$. As a control we also included the question “do you own a pet?” We found no evidence of an association between answers to this question and participant ascriptions of belief (Fisher’s, $p = 0.83$) or knowledge (Fisher’s, $p = 0.53$) in Original Dog.

³⁸ McNemar’s, $p < 0.001$, one-tailed.

³⁹ Fisher’s exact, $p = 0.059$, one-tailed.

⁴⁰ Binomial, $p = 0.002$, one-tailed.

⁴¹ Binomial, $p = 0.021$, one-tailed. Overall response patterns for Original Dog, Thick Dog, and Thin Dog, respectively, were as follows: No knowledge no belief (37%, 17%, 22%), no knowledge belief (16%, 16%, 29%) knowledge no belief (16%, 16%, 8%), knowledge and belief (32%, 51%, 41%).

⁴² There are four reasons for this. First, we were confident by this point that our plot-thickening manipulation was working effectively as predicted. Second, Murray et al.’s original God case was, unfortunately, a between-subjects design in which participants were simply probed with one question (either about knowledge or about belief). Consequently, there was no plot to thicken! Third, we agree with Rose and Schaffer that in testing entailment one ultimately wants to know whether the same person ascribes knowledge but not belief. A between-subjects design is thus inadequate for this. Lastly, we’re ultimately most interested in what will happen when we use the thin belief probe.

⁴³ Thirty-three female, aged 18–58, $M = 27$ years, $SD = 8$ years.

⁴⁴ McNemar’s, $p = 1.00$.

⁴⁵ Binomial, $p = 0.01$, one-tailed. Overall response patterns for Thin God, and Thin God (Believers), respectively, were as follows: No knowledge no belief (32%, 2%), no knowledge belief (7%, 4%), knowledge no belief (7%, 8%), knowledge and belief (54%, 86%).

⁴⁶ Fisher’s, both $ps < 0.001$, one-tailed.

⁴⁷ McNemar’s, both $ps \geq 0.687$.

⁴⁸ Binomial, $p = 0.008$.

⁴⁹ Binomial, $p = 0.15$.

⁵⁰ Our results provide intuitive support for the traditional thesis that knowledge entails (thin) belief. However, the thin entailment thesis is neutral with respect to a series of further psychological questions, including the psychological priority of knowledge and belief. For more on whether knowledge judgments are based on prior or accompanying judgments about belief, see Rose under review and Buckwalter & Turri under review b.

⁵¹ Knowledge, $X^2(1) = 4.82$ $p < 0.05$. Belief, $X^2(1) = 26.94$ $p < 0.001$.

⁵² McNemar's test, $ps < 0.001$ for all comparisons.

⁵³ Thirty-one female, aged 19–75, $M = 29$ years, $SD = 9.32$ years.

⁵⁴ McNemar's, $ps < 0.01$ for all comparisons.

⁵⁵ Forty-nine female, aged 18–71, $M = 32$ years, $SD = 11.11$ years.

⁵⁶ Binomial, all comparisons $ps < 0.05$.

⁵⁷ Fisher's exact, $p = 0.79$.

⁵⁸ Fisher's exact, $p = 0.40$.

⁵⁹ See Turri, Blouw and Buckwalter under review for more on the connection between knowledge and luck.

⁶⁰ Also see Schwitzgebel 2010, 546–548 for pragmatic arguments in favor of “thicker” propositional and reactive attitudes more generally.

⁶¹ This is reminiscent of responses by Audi 2008 and Cohen 1992 distinguishing between belief and acceptance for the purposes of normatively evaluating beliefs.

⁶² For helpful feedback and discussion, we thank Josh Alexander, Peter Blouw, Carolyn Buckwalter, Ori Friedman, Kareem Khailfa, Joshua Knobe, Jonathan Livengood, Charles Millar, Dylan Murray, Blake Myers-Schulz, Jonathan Schaffer, Eric Schwitzgebel, Angelo Turri, and two anonymous reviewers. This research was supported by the Social Sciences and Humanities Research Council of Canada and an Ontario Early Researcher Award.

References

- Alicke, M., and Rose, D. (2010). Culpable Control or Moral Concepts? *Behavioral and Brain Sciences*, 33, 330–31.
- Armstrong, D. (1969). Does Knowledge Entail Belief? *Proceedings of the Aristotelian Society*, 70, 21–36.
- Armstrong, D. (1973). *Belief, Truth, and Knowledge*. Cambridge: Cambridge Press.
- Audi, R. (1994). Dispositional Beliefs and Dispositions to Believe. *Noûs*, 28, 419–434.
- Audi, R. (2008). Belief, Faith, and Acceptance. *International Journal for Philosophy of Religion*, 63, 87–102.
- Augustine. [428]. *On the Predestination of the Saints*. Trans. P. Holmes, R.E. Wallis; rev. B.B. Warfield; rev. K. Knight. Retrieved from <http://www.newadvent.org/fathers/15121.htm>, 31 December 2012.
- Bach, K. (2008). Applying Pragmatics to Epistemology. *Philosophical Issues*, 18, 68–88.
- Bach, K., and Harnish, R.M. (1979). *Linguistic Communication and Speech Acts*. Cambridge: MIT Press.
- Beebe, J. (Forthcoming). A Knobe Effect for Belief Ascriptions. *Review of Philosophical Psychology*.
- Beebe, J., and Buckwalter, W. (2010). The Epistemic Side-Effect Effect. *Mind & Language*, 25, 474–98.
- Buckwalter, W. (Under review). Factive verbs and protagonist projection.
- Buckwalter, W. (2013). Gettier Made ESEE. *Philosophical Psychology*, DOI:10.1080/09515089.2012.730965.
- Buckwalter, W., and Turri, J. (Under review a). In the Thick of Moral Motivation.
- Buckwalter, W., and Turri, J. (Under review b). Action, Truth and Knowledge.
- Campbell, C. A. (1967). Towards a Definition of Belief. *The Philosophical Quarterly*, 17, 204–220.
- Chisholm, R. (1989). *Theory of Knowledge*, 3rd ed. Prentice-Hall.
- Cohen, J. (1966). More about Knowing and Feeling Sure. *Analysis*, 27, 11–16.
- Cohen, L. J. (1992). *An Essay on Belief and Acceptance*. Oxford: Clarendon Press.

- Dretske, F. (1983). The Epistemology of Belief. *Synthese*, 55 (1): 3–19.
- Dretske, F. (1988). *Explaining Behavior*. Cambridge: MIT Press.
- Feldman, R. (2003). *Epistemology*. Prentice Hall PTR. pp. 12–13
- Fodor, J. A. (1990). *A Theory of Content*. Cambridge: MIT Press.
- Gorman, M. (1993). Hume's Theory of Belief. *Hume Studies*, 19, 89–102.
- Greco, D. (Forthcoming). Iteration and Fragmentation. *Philosophy and Phenomenological Research*.
- Hansen, S. J. (1988). Hume's Impressions of Belief. *Hume Studies*, 14, 277–304.
- Hazlett, A. (2010). The Myth of Factive Verbs, 80 (3): 497–522.
- Hazlett, A. (2012). Factive Presupposition and the Truth Condition on Knowledge. *Acta Analytica*, 27 (4): 461–478. doi:10.1007/s12136-012-0163-3.
- Horgan, T., and Tienson, J. (2002). The Intentionality of Phenomenology and the Phenomenology of Intentionality. *Philosophy of Mind: Classical and Contemporary Readings*, ed. Chalmers: 520–33. Oxford University Press.
- Hume, D. 1748 [1977]. *An Enquiry Concerning Human Understanding*. (E. Steinberg, Ed.) (2nd ed.). Indianapolis: Hackett.
- 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.
- Knobe, J. (2010). Person as Scientist, Person as Moralizer. *Behavioral and Brain Sciences*, 33, 315–29.
- Knobe, J., Buckwalter, W., Nichols, S., Robbins, P., Sarkissian, H., and Sommers, T. (2012). Experimental Philosophy. *Annual Review of Psychology*, 63, 81–99.
- Knobe, J., Prasada, S., and Newman, G. (2013). Dual Character Concepts and the Normative Dimension of Conceptual Representation. *Cognition*, 127, 242–257.
- Korsgaard, C. (1986). Skepticism about Practical Reason, *Journal of Philosophy*, 83, 5–25.
- Lehrer, K. (1968). Belief and Knowledge. *The Philosophical Review*, 77, 491–499.
- Lehrer, K. (1974). *Knowledge*. Clarendon Press.
- Lycan, W. (1988). Judgment and justification. Oxford: Oxford University Press.
- Millgram, E. (1995). Was Hume a Humean? *Hume Studies*, 21 (1): 75–94.
- Murray, D., Sytsma, J., and Livengood, J. (Forthcoming). God Knows (But does God Believe?). *Philosophical Studies*.
- Myers-Schulz, B. and Schwitzgebel, E. (2013). Knowing that P without Believing that P. *Noûs*, 47 (2): 371–48.
- Pascal, B. (1670). *Pensées*, A. Kreilsheimer, trans. and ed., New York: Penguin.
- Peirce, C. S. (1958). *Collected Papers of Charles Sanders Peirce*. Ed. Charles Hartshorne, Paul Weiss, and Arthur W. Burks, Harvard University Press, Cambridge, Massachusetts.
- Peirce, C. S. (1976). *The New Elements of Mathematics*, Ed. Carolyn Eisele, Mouton Publishers, The Hague, 2, 312.
- Price, H. H. (1969). *Belief*. Allen & Unwin.
- Radford, C. (1966). Knowledge—By Examples. *Analysis*, 27, 1–11.
- Rose, D. (Under review). The Psychological Case for the Primacy of Knowledge.
- Rose, D., Buckwalter, W., and Turri, J. (Under review). When Words Speak Louder Than Actions: Delusion, Belief and the Power of Assertion.
- Rose, D., and Schaffer, J. (Forthcoming). Knowledge Entails Dispositional Belief. *Philosophical Studies*.
- Ryle, G. (1949). *The Concept of Mind*. London: Hutchinson.
- Sayre-McCord, G. (2006). Hume on Practical Morality and Inert Reason. In *Oxford Studies in Metaethics*, Vol. 3, ed. R. Shafer-Landau, 299–320. Oxford: Oxford University Press.
- Schwitzgebel, E. (2001). In-Between Believing. *Philosophical Quarterly*, 51, 76–82.
- Schwitzgebel, E. (2002). A Phenomenal, Dispositional Account of Belief. *Noûs* 36, 249–75.
- Schwitzgebel, E. (2010). Acting Contrary to Our Professed Beliefs, or the Gulf Between Occurrent Judgment and Dispositional Belief. *Pacific Philosophical Quarterly*, 91, 531–553.
- Searle, J. R. (1983). *Intentionality*. Cambridge: Cambridge UP.
- Smith, M. (1994). *The Moral Problem*. Oxford: Basil Blackwell.
- Starmans, C., and Friedman, O. (2012). The Folk Conception of Knowledge. *Cognition*, 124 (3): 272–83.
- Steup, M. (2006). The Analysis of Knowledge. *Stanford Encyclopedia of Philosophy*.

- Stocker, M. (1979). Desiring the Bad: an Essay in Moral Psychology. *The Journal of Philosophy*, 76, 738–53.
- Stokke, A. (2013). Protagonist Projection. *Mind and Language*, 28 (2): 204–232.
- Szabó, Z. G., and Knobe, J. (2013). Modals with a Taste of the Deontic. *Semantics and Pragmatics*, 6, 1–42.
- Travis, C. (2005). A Sense of Occasion. *Philosophical Quarterly*, 55, 286–314.
- Turri, J. (2010). Prompting Challenges. *Analysis*, 70 (3): 456–462. doi:10.1093/analysis/anq027.
- Turri, J. (2011a). Mythology of the Factive. *Logos & Episteme*, 2 (1): 143–152.
- Turri, J. (2011b). The Express Knowledge Account of Assertion. *Australasian Journal of Philosophy*, 89 (1): 37–45.
- Turri, J. (2013a). A Conspicuous Art: Putting Gettier to the Test. *Philosophers' Imprint*.
- Turri, J. (2013b). The Test of Truth: An Experimental Investigation of the Norm of Assertion. *Cognition*, 129 (2): 279–91.
- Turri, J. (Forthcoming). You Gotta Believe. *Epistemic Norms: New Essays on Action, Belief and Assertion*. Ed. C. Littlejohn and J. Turri. Oxford UP.
- Turri, J., Blouw, P., and Buckwalter, W. (Under review). Knowledge and Luck.
- Unger, P. (1975). *Ignorance: a Case for Skepticism*. Oxford. Oxford UP.
- Weatherston, B. (2008). Deontology and Descartes' Demon. *Journal of Philosophy*, 105, 540–569.
- Weiner, M. (2005). Must we know what we say? *Philosophical Review*, 114 (2): 227–251.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press.
- Zagzebski, L. (2009). *On epistemology*. Wadsworth.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's web site:

Original Geocentrist (5.1.1-2)

Thick Geocentrist (5.1.1-2)

Thin Geocentrist (5.1.1-2)

Geocentrist Follow-up: Matching Knowledge Probe (5.1.3)

Original Dog (5.2.1-2)

Thick Dog (5.2.1-2)

Thin Dog (5.2.1-2)

God Thin (5.3.1-2)

Original Geo Within Subjects (5.5)

Original Dog Within Subjects (5.5)

Guess Control – Thin Sloan (5.6)

Guess Control – Matching Knowledge Sloan (5.6)

Guess Control – Thick Sloan (5.6)