

**Evidence for multiple kinds of belief in theory of mind**

**Alejandro Vesga**

Assistant Professor  
Department of Philosophy  
University of Houston  
ajvesgaplazas@uh.edu

**Neil Van Leeuwen**

Professor  
Department of Philosophy  
Florida State University  
nvanleeuwen@fsu.edu

**Tania Lombrozo**

Arthur W. Marks '19 Professor  
Department of Psychology  
Princeton University  
lombrozo@princeton.edu

Datasets, code scripts, and supplemental materials are available in OSF with DOI identifier [10.17605/OSF.IO/38YGN](https://doi.org/10.17605/OSF.IO/38YGN). (Please refer to the published version in JEP: General)

### Abstract

People routinely appeal to ‘beliefs’ in explaining behavior; psychologists do so as well (for instance, in explaining belief polarization and learning). Across three studies (N = 1,843, U.S.-based adults), we challenge the assumption that ‘belief’ picks out a single construct in people’s theory of mind. Instead, laypeople attribute different kinds of beliefs depending on whether the beliefs play predominantly epistemic roles (such as truth-tracking) or non-epistemic roles (such as social signaling). We demonstrate that epistemic and non-epistemic beliefs are attributed under different circumstances (Study 1) and support different predictions about the believer’s values (Study 2) and behavior (Study 3). This differentiation emerges reliably across three distinct signatures of attributed belief, and even when the believed content and attributed level of certainty about that content are held constant across cases. Our findings call for a more fine-grained characterization of theory of mind and provide indirect support for the hypothesis that human cognition itself features multiple varieties of belief.

*Keywords:* belief; theory of mind; epistemic reasoning; cognitive attitudes; mental state attribution.

*Availability Statement:* Datasets, code scripts, and other materials are available in OSF with DOI identifier [10.17605/OSF.IO/38YGN](https://doi.org/10.17605/OSF.IO/38YGN).

### **Statement of Public Relevance**

People routinely appeal to ‘beliefs’ in explaining and predicting behavior. Psychologists do so as well – for instance, they describe *belief* polarization and explain placebo effects by appeal to beliefs. But are beliefs really a unitary phenomenon? Recent work suggests that beliefs come in different flavors: some beliefs are mainly about representing the world; others play important roles in social signaling, emotion regulation, and more. We find that this distinction has an analog in people’s “theory of mind.” When attributing beliefs to others and predicting behavior on the basis of those beliefs, participants systematically differentiated kinds of believing: beliefs that aim at truth and accuracy and beliefs that have other functions, such as social signaling or emotional regulation. These findings support a new and nuanced picture of how people represent and reason about the minds of others.

### **Authors’ Notes**

Some or all of the ideas and data reported in this manuscript were presented in the 2024 Conference of the Society of Philosophy and Psychology, held in Purdue University, and the 2024 Conference of the Cognitive Science Society, held in The Netherlands. The latter only included some of the data presented here, and the corresponding proceedings can be accessed in [reference omitted for anonymity.]

## Introduction

Notions of ‘belief’ play central roles in both scientific and intuitive theories of mind. Psychologists posit beliefs to explain various phenomena (Porot & Mandelbaum, 2021), from placebo effects, which are a consequence of individuals’ *beliefs* (Ossipov et al., 2010; Watson et al., 2012), to learning, which generates changes in *belief* (Hofer & Pintrich, 1997; Rivers, 2021; Yan et al., 2014). Some phenomena even have belief in their name, as in *belief* bias and *belief* polarization (Campbell & Kay, 2014; Cohen, 2003; Pronin, 2007). And, unlike many other constructs in psychological science, such as semantic memory or reward prediction error, ‘belief’ is also common in everyday folk explanations. People posit beliefs to explain others’ behaviors (“she took an umbrella because she believes it will rain”), inferences (“she accused the butler because she believes he was the one at the crime scene”), and social allegiances (“they get along because they both believe the 49ers are the best”) (Bendaña & Mandelbaum, 2021; Quilty-Dunn & Mandelbaum, 2018; Westra, 2023). Belief is so central to intuitive psychology that it claims half the billing in the paradigmatic framework for theory of mind: “belief-desire psychology” (Goldman, 2006; Wellman & Woolley, 1990).

Yet, despite the ubiquity of belief—or perhaps because of it—the task of characterizing this mental state remains elusive (Jong, 2018; Van Leeuwen & Lombrozo, 2023). Some research suggests that beliefs are responsive to evidence; other work suggests the opposite (Bergamaschi Ganapini, 2020; Bortolotti, 2010). Some findings suggest that counterevidence leads to belief polarization; other findings suggest otherwise (Kahan, 2015; Ranney & Clark, 2016). Some research suggests that beliefs are not under voluntary control; other research suggests that they are (Cusimano & Goodwin, 2019, 2020; Gilovich & Regan, 1986; Malle & Knobe, 1997; Turri et al., 2018). This heterogeneity is a challenge for efforts to define belief and renders it puzzling

how appeals to belief can and (often) do succeed in offering psychological explanations or predictions.

One approach to this heterogeneity has been gaining recent traction: positing different *kinds* of beliefs or belief-like “cognitive attitudes” – different flavors of believing (Apperly & Butterfill, 2009; Chinn et al., 2014; Davoodi & Lombrozo, 2022; Metz et al., 2023; Rutjens & Preston, 2020; Van Leeuwen, 2014, 2023; Westra, 2023). One version of this hypothesis, which we call the *Varieties of Belief* hypothesis, posits the existence of at least two kinds of cognitive attitude, where one is more closely aligned with truth-tracking (epistemic) functions, such as representing the world to support accurate predictions, and the other with social and motivational (non-epistemic) functions, such as group cohesion and emotion regulation. On this account, different kinds of believing are viewed as algorithmic-level processes shaped by different computational-level problems (Marr, 1982). That is, the different processes involved in generating and utilizing distinct kinds of beliefs can be understood in light of the epistemic or non-epistemic aims that those beliefs serve. Beliefs that serve epistemic aims (such as accurately representing the world) may fall short in achieving non-epistemic aims (such as maintaining emotional and social well-being) and *vice versa*. To best handle these conflicting epistemic and non-epistemic aims, cognition operates with more than one variety of believing.

In the current research, we explore an idea that parallels the Varieties of Belief hypothesis and arises in the context of “theory of mind,” or people’s mental toolkit for understanding minds. We call it the *Belief Pluralism* hypothesis. One might ask: If it is indeed likely that there are different kinds of belief serving different cognitive functions, do laypeople (not just researchers) track those differences in their own representations of people’s beliefs? In other words, do

everyday attributions of belief show differentiations that align with epistemic and non-epistemic functions? The Belief Pluralism hypothesis says that the answers to these questions is “yes.”

Of course, the Varieties of Belief hypothesis and its correlate in theory of mind, Belief Pluralism, are in principle logically independent claims. Nevertheless, support for the former motivates the latter, and support for the latter indirectly supports for the former. Theory of mind is typically understood as an aspect of human social cognition that emerges – through maturation or learning – for the purpose of predicting and explaining behavior. If the Varieties of Belief hypothesis is correct, such that human cognition in fact involves distinct varieties of believing, then we should expect that at least in some cases, behavior will be more successfully predicted and explained by attributing different varieties of belief. Hence, support for the Varieties of Belief hypothesis at least motivates the question of whether there is a correlate in people’s intuitive theory of mind; or, in our terms, whether the Belief Pluralism hypothesis is true. In turn, the finding that people attribute different kinds of beliefs along epistemic and non-epistemic lines would lend indirect support to the Variety of Belief hypothesis. If people systematically distinguish between epistemic and non-epistemic beliefs when reasoning about people’s minds, it is plausible that people are indeed tracking a real cognitive distinction.

Whether or not the Varieties of Belief hypothesis is ultimately true, finding that people’s theory of mind systematically differentiates beliefs with epistemic and non-epistemic functions would have important implications (see also Westra, 2023).<sup>1</sup> First, it would call for revision in a core construct in standard views of theory of mind: the notion of belief itself. Paradigmatic accounts of theory of mind posit a simple distinction between desire and belief. This distinction, as well as the unitary conception of belief that it seems to presuppose, has guided research in

---

<sup>1</sup> See Westra (2023) for an important precedent congenial with our questions. Westra distinguishes between *epistemic* and *symbolic* concepts of beliefs that are deployed differentially in intuitive theory of mind.

developmental psychology, comparative psychology, clinical psychology and beyond, where motivating questions include: Do children and non-human animals represent others' false beliefs? Do two populations (that differ in age, ASD diagnosis, or otherwise) reason about other minds in the same way? If the Belief Pluralism hypothesis is true, these questions must be refined to target specific varieties of belief. More generally, attributions of belief play an important role across the social sciences (for instance, in explaining attributions and misattributions of belief across political divides and across cultures), such that pluralism about belief would have widespread implications.

### **The Belief Pluralism Hypothesis**

The Belief Pluralism Hypothesis posits that people's theory of mind distinguishes between two kinds of beliefs: epistemic and non-epistemic. This is intended as a distinction in cognitive attitudes, or ways of "believing," not in belief contents. To illustrate, suppose Jane believes that a Republican will win the next election, and this belief comes from her foray into sociological research about American politics and informs her predictions about how economic markets will behave. That is, Jane's belief serves largely epistemic functions. Jack, on the other hand, may also believe that a Republican will win the next election (the same content), but does so because of the emotional weight of his political identity and the role this belief plays in signaling his commitments to his social network. In other words, Jack's belief serves non-epistemic functions. Observers might attribute "belief" that a Republican will win to both Jane and to Jack, but these attributions could involve different kinds of believing that license different expectations – for example, that in light of new evidence, Jane will more readily change her belief than Jack will change his.

While some prior work supports a distinction between epistemic and non-epistemic aspects of belief, this work has not controlled for differences in belief content (Shtulman 2013, Metz et al, 2023, Metz et al. 2020; Davoodi & Lombrozo 2022). As one example, Heiphetz et al. (2013) investigated differences between factual beliefs, ideological beliefs, and preferences. They found that both adults and 5-10 year-old children judged that two people who disagree about a factual belief (for instance, whether germs are very big or very small) cannot both be right, whereas two people who disagree about a preference (for instance, whether pink or green is the prettiest color) can be. Participants' judgments about ideological beliefs (for instance, whether there is one God or many) fell in between. These findings reveal systematic differentiation across kinds of beliefs in lay theory of mind, but the study was not designed to isolate differences in cognitive attitude, as posited by the Belief Pluralism hypothesis, so it remains in principle possible that the differences that surfaced reflect content alone. In our own studies, we aim to investigate differences in cognitive attitude while holding content fixed, as in our example of Jane and Jack.

### **Signatures of Distinct Belief Attributions**

To test the Belief Pluralism hypothesis, we identify diagnostic “signatures” that could plausibly differentiate attributions of epistemic versus non-epistemic beliefs. To motivate this approach, consider a parallel with accounts of human cognition that posit two systems, such as System 1 and System 2 (Evans, 2008; Sloman, 1996; Stanovich & West, 2000; Tversky & Kahneman, 1981). On these views, System 1 tends to be fast, automatic, and emotional, whereas System 2 is slower, controlled, and more logical. Measured processing speed, automaticity, and emotional involvement are therefore “signatures” that can be used to diagnose which system is in



operation. Of course, these signatures will not be perfect. On some occasions a slow response could still reflect the operation of System 1, and a fast response could reflect the (unusually speedy) operation of System 2. Responses could also reflect a combination of both systems. But in general, there will be some association between measured response time and system type, such that System 1 responses are more likely to be fast, and System 2 responses are more likely to be slow.

Similarly, we propose three signatures that help differentiate people's attributions of epistemic versus non-epistemic beliefs. Our signatures, unlike those for System 1 and System 2, reflect the socio-cognitive roles for different concepts of belief. As we explain further below, our signatures are binary vs. probabilistic construal, perceived directional versus non-directional control, and the use of "believe" versus "think" in natural language ascriptions. The signatures are not expected to perfectly trace the hypothesized distinction between cognitive attitudes, but we expect them to express patterns that are somewhat diagnostic of which cognitive attitude people are attributing, just as response time is somewhat diagnostic of cognitive system. If these signatures reliably differentiate attributions of epistemic versus non-epistemic beliefs, it should be possible to elicit judgments concerning these signatures as a way of diagnosing whether a belief attribution posits epistemic or non-epistemic belief, and to use these signatures to induce epistemic versus non-epistemic attributions.

### ***Binary vs. probabilistic construal***

Our first signature reflects the extent to which a belief is construed as a categorical or more probabilistic commitment. In particular, a belief can be construed as binary (one believes that  $p$  or not) or as a subjective probability (one believes with  $x\%$  probability that  $p$  is true).

Intuitively, if a belief's function is to accurately represent the world, it should be proportional to evidence and updated incrementally as new evidence is obtained (for instance, by following Bayes' rule). This suggests that an epistemic cognitive attitude should support or at least allow for a probabilistic (or more generally, graded) construal. However, a probabilistic construal seems less appropriate if a belief's function is to signal religious or moral convictions, or to signal group allegiance (compare: "I believe that a Republican will win the next election" with "I believe there is a 94% chance that a Republican will win the next election"). In such cases, a binary belief may function more effectively by obscuring uncertainty or hesitation.

Note that a belief's being preferentially construed in binary terms does not rule out the possibility that the person who has the belief *also* has or can assign some subjective probability or degree of certainty about the truth of the belief's contents, which they may or may not wish to reveal. Religious believers, for example, may generally report their religious beliefs in binary terms ("I believe!"), though they may still privately harbor more or less doubt about the truth of the professed belief (Davoodi et al., 2019; Luhrmann, 2012; Van Leeuwen, 2022). So—importantly for our experimental design—people attributing beliefs may tend to construe them in binary terms when they serve non-epistemic aims, while still being able to reach an independent judgment about the degree of certainty of the person holding the attributed (non-epistemic) belief (Nelson & Lombrozo, in prep).

Based on this signature, our first hypothesis is that attributions of epistemic cognitive attitudes are more likely than attributions of non-epistemic cognitive attitudes to involve probabilistic (vs. binary) construals. Returning to our example, while Jane might be naturally described as believing with some probability that a Republican will win the next election, Jack will more naturally be described as believing this categorically.

***Perceived Directional vs. non-directional control***

Our second signature concerns voluntary control. Voluntary control over belief is often taken to threaten belief's ability to reliably track the truth: If we can decide what to believe at will, then our beliefs will be determined by our preferences, not reality (Arpaly & Brinkerhoff, 2018; Kelly, 2003). For example, a president who lost an election could make themselves feel better by simply choosing to believe they actually won. This motivates the prediction that people will associate epistemic cognitive attitudes with limited voluntary control. In contrast, non-epistemic cognitive attitudes that support social signaling or identity preservation may well be under the influence of directional motives, since their need for accuracy is lower. If, for example, the president's "belief" about the election does not serve the psychological role of accurately representing the world, and instead serves the role of signaling values or mobilizing followers, it may be more voluntary and perceived as such.

For those steeped in the philosophical literature on doxastic voluntarism, control over belief might seem like an implausible attribution – the orthodox view is that direct control over belief is absent or limited (Alston, 1988; Audi, 2001; Hieronymi, 2006). By contrast, empirical work has found that laypeople are generally quite willing to endorse various forms of control over belief. For instance, recent work by Cusimano and colleagues has found that people judge others to have considerable control over their beliefs, and indeed sometimes endorse directional influences on belief, such as believing in line with one's moral obligations (Cusimano & Goodwin 2020; Cusimano & Lombrozo, 2020, 2023; Cusimano et al., 2024; see also Turri et al., 2018). Combining these trains of thought, we developed the expectation that participants would find it natural to classify beliefs as resulting from *different* forms control: what we call *non-*

*directional control* (one decides whether to come to a belief about something, but without presupposing a conclusion) versus *directional control* (one decides *to* believe something).

Our second hypothesis is that in people's belief attributions, epistemic cognitive attitudes are more likely than non-epistemic cognitive attitudes to be associated with non-directional control ('*S* decided to believe whether *p*'), as opposed to directional control ('*S* decided to believe *p*'). To return to our running example, it should be natural to judge that Jane can exert *non-directional* control in deciding to form a belief about whether or not a Republican will win the next election—for instance, she can decide to give it more thought or to gather more evidence. What she cannot (or at least should not) do is exert *directional* control by simply deciding to believe that a Republican will win the next election. Such directional control would threaten the link between her belief and reality, and thus threaten the belief's epistemic status. By contrast, it might be natural to judge that after joining his young conservatives club, Jack exerted *directional* control in deciding to believe *that* a Republican will win the next election. If Jack's belief is non-epistemic, then directional influences on belief will not necessarily undermine the belief's (non-epistemic) functions.

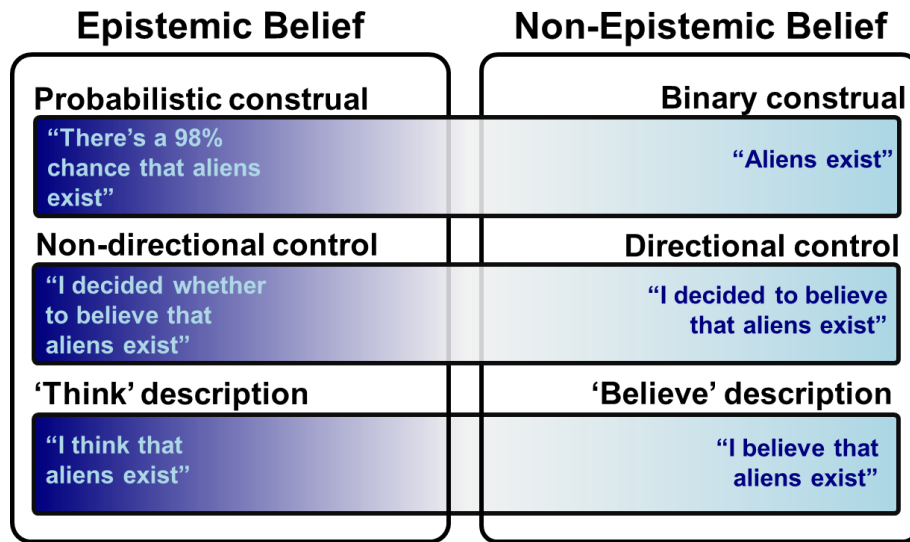
### ***'Believe' vs 'think' descriptions***

Our third and final signature concerns the use of "believes" versus "thinks." Recent research has found that people tend to use different words for religious and non-religious beliefs in a number of different languages (Heiphetz et al., 2021; Van Leeuwen et al., 2021). In English, for instance, beliefs with religious contents or held for religious reasons are more likely to be reported using "believe" than "think" (e.g., "she believes that Jesus turned water into wine" but "she thinks that the wine is a merlot"). This pattern is also found in other languages, as with

‘dwen’ and ‘gyi dze’ in Fante (an Akan dialect common in Ghana) and 认为 (rènwéi) and 相信 (xiāngxìn) in Chinese. A plausible explanation of this effect is that these linguistic differentiations tend to track whether a belief is perceived to be held with an epistemic or a non-epistemic cognitive attitude. If so, the differential pattern of use of “thinks” vs. “believes” should also surface in ascriptions of beliefs with epistemic versus non-epistemic roles more generally (whether or not those non-epistemic roles are religious).

Hence, our third hypothesis is that ascriptions of epistemic cognitive attitudes are more likely than ascriptions of non-epistemic cognitive attitudes to involve “think” (e.g., ‘S thinks that *p*’) versus “believe” (e.g., ‘S believes that *p*’). So while Jane will be more likely to be judged as *thinking* that a Republican will win the next election, Jack will be more likely to be judged as *believing* that a Republican will win the next election.

Importantly, this signature of cognitive attitude – like those above – is not deterministic. There are likely to be some uses of “thinks” that correspond to a non-epistemic attitude, and some uses of “believes” that correspond to an epistemic attitude (just as response time will be a noisy indicator of whether some judgment is the output of System 1 or System 2). ‘Thinks’ and ‘believes’ overlap extensively in their usage, and they play roles beyond those we focus on here. For instance, both can be used as a hedge to indicate uncertainty (“I *think/believe* the party is on Friday”). Nonetheless, we predict that some signal can be extracted from uses of “thinks” versus “believes,” such that we will see an association between word choice and epistemic versus non-epistemic cognitive attitude (especially when controlling for belief certainty).



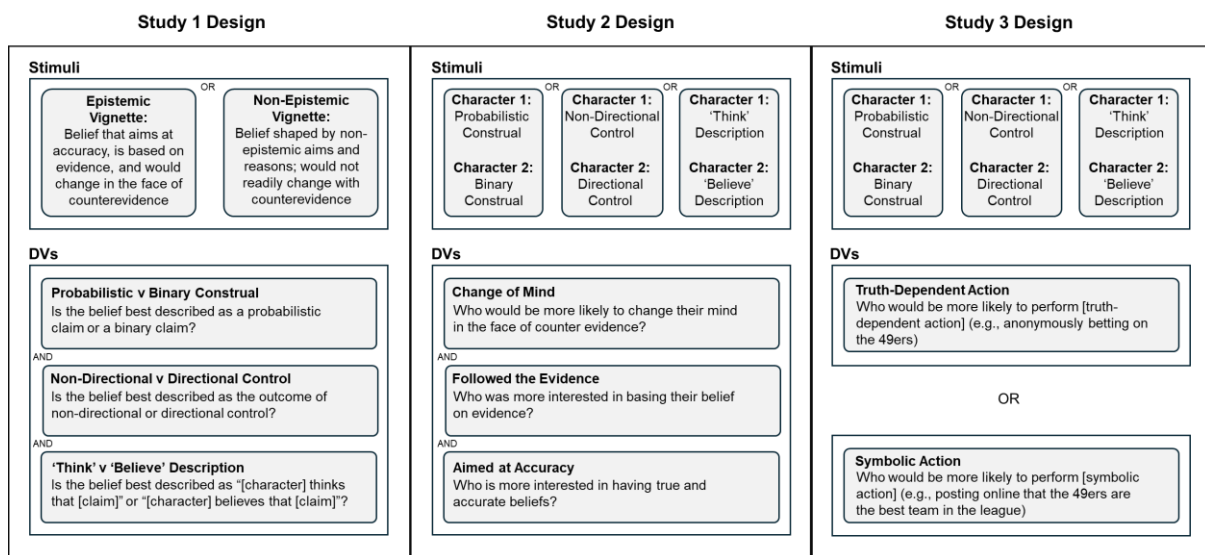
**Figure 1.** Three signatures to distinguish ascriptions of epistemic beliefs from ascriptions of non-epistemic beliefs: Probabilistic v Binary construals, Perceived Non-Directional v Directional control, and ‘Think’ v ‘Believe’ descriptions. The leftmost element of each signature is its epistemic pole; the rightmost element of each signature is its non-epistemic pole.

### Present research

Across three studies ( $N = 1,843$  U.S.-based adults), we use these hypothesized signatures of belief attributions to investigate whether laypeople differentiate epistemic and non-epistemic cognitive attitudes. In Study 1, we test whether beliefs that play epistemic roles are more likely than those that play non-epistemic roles to support a probabilistic (vs. binary) construal, to be perceived as involving non-directional (vs. directional) control, and to be described in terms of “think” (vs. “believe”). In Study 2, we test whether participants make the reverse set of inferences: that beliefs that are reported as probabilistic (vs. binary), achieved through non-directional control (vs. directional control), and expressed using “think” (vs. “believe) are more likely to be epistemic. Finally, in Study 3, we test an important consequence for theory of mind:

whether participants predict different kinds of behavior (truth-dependent vs. symbolic) from agents who express their beliefs in keeping with the signatures associated with epistemic vs. non-epistemic attitudes. See Figure 2 for the design scheme.

All studies were approved through the Institutional Review Boards of the authors' universities and preregistered using *aspredicted.com*. Preregistrations, data, materials, and supplements are available on [OSF](#).



**Figure 2.** Schematic representation of the designs of Studies 1-3.

### *Transparency and Openness.*

All studies adhere to the Transparency and Openness Promotion Guidelines. All experimental materials (including surveys and stimuli), deidentified data, preregistrations, and scripts for data processing and analysis necessary to reproduce the results are openly available at OSF with DOI identifier [10.17605/OSF.IO/38YGN](https://doi.org/10.17605/OSF.IO/38YGN). The analyses were conducted using R (v 4.4.2). The scripts include calls for required packages and were tested in a separate system to ensure reproducibility.

## Studies

### Study 1

In Study 1, we tested whether beliefs that play epistemic roles are more likely than those that play non-epistemic roles to receive a probabilistic (vs. binary) construal, to be perceived as involving non-directional (vs. directional) control, and to be described in terms of “think” (vs. “believe”).

### *Participants*

Study 1 had a total sample of 400 U.S.-based adults recruited through Prolific and compensated \$0.55 for a 3-minute study. We excluded 17 participants for failing one or more basic attention checks, for a total of 383 participants in our analyzed sample (age:  $m = 38.82$ ,  $\sigma = 13.39$ ; Male = 192, Female = 183, Other = 8). Sample sizes were determined by power analyses based on effect sizes obtained in pilot studies.

Although we did not collect additional demographic information from our samples, a study of online data quality recruiting a Prolific sample (Douglas et al., 2023) found that the majority of participants identified their ethnicity as white (72.38%), followed by Asian or Asian America (12.90%), Black or African American (9.88%) and Latino, Hispanic, Chicano or Puerto Rican (8.67%). Family income was distributed from less than \$10,000 per year (6.25%) to over \$150,000 per year (7.86%), with 47.57% in the \$10,000 to <\$60,000 range, and the remaining 37.90% in the \$60,000 to <\$150,000 range (.40% of the sample did not respond). With regard to highest education, the sample ranged from less than a high school education (.40%) to a graduate



degree (15.12%), with the modal response corresponding to a 4-year degree (36.29%), followed by having completed some college (24.40%).

### *Procedure*

Study 1 employed a fully between-subjects design. Participants were randomly assigned to one of eight vignettes describing a character with a given belief. Half of these vignettes corresponded to the epistemic condition. The other half to the non-epistemic condition. Below are two sample vignettes illustrating how the same belief content was preserved across the epistemic and non-epistemic conditions:

Adam and John were best friends in their childhood and teenage years. However, after graduating from high school John left to serve in Water Without Borders, a humanitarian program based in the Global South. Ten years later, after traveling the world with the organization, John came back to his town and rebuilt his friendship with Adam, who never left. However, one day, Adam hears a rumor that John came back because he was stealing money from the organization he was working for.

*Epistemic condition:* Adam has trouble accepting this, because it doesn't fit with what he knows about John. If Adam were presented with strong evidence, he would fully accept John's guilt. Ultimately, what's important to Adam is to accept whatever is true. But, based on the evidence he has, Adam currently \_\_\_ that John did not steal money from the organization he was working for.

*Non-Epistemic condition:* Adam has trouble accepting this, because he wants to be a loyal friend to John. Even if Adam were presented with strong evidence, he would have trouble accepting John's guilt. It would feel like wronging John. What's important to Adam is nurturing his friendship with John, and he wouldn't want to betray that. Given these feelings, Adam currently \_\_\_ that John did not steal money from the organization he was working for.

After reading their assigned vignette, participants were asked to make several judgments pertaining to the character's belief, expressed in the last sentence of the vignette (e.g., that "John did not steal money from the organization he was working for"). Our three primary dependent variables captured the signatures of cognitive attitude described in the introduction.

**Believe/Think.** Participants were asked: "What is a more natural way of completing the blank in the sentence?" They selected between "thinks" and "believes." This question was adapted from Van Leeuwen et al. (2021). After participants made a choice, the blank in the vignette was replaced with the word they selected.

**Binary/Probabilistic.** Participants were asked: "Do you think it is natural to describe [Character's] belief as corresponding to some probability (= "[Character] believes that there is an x% chance that [claim]")?" They responded either "No: [Character] simply believes that [claim] (as opposed to not believing this)" or "Yes: It is natural to describe [Character's] belief as corresponding to some probability."

**Directional/Non-Directional.** Participants were asked: "Which of the following is a better description of what happened?" They selected between "[Character] decided to believe that [claim]" and "[Character] decided whether to believe that [claim]."

Across our four pairs of vignettes, we varied the non-epistemic considerations relevant to the non-epistemic condition. Besides the concerns with loyalty in the John/Adam vignette illustrated above, we included a vignette concerning religious faith, one concerning a teacher's moral commitment to see potential in her students (motivated by Cusimano & Lombrozo, 2021), and one about an individual's identity and desire to be true to herself.

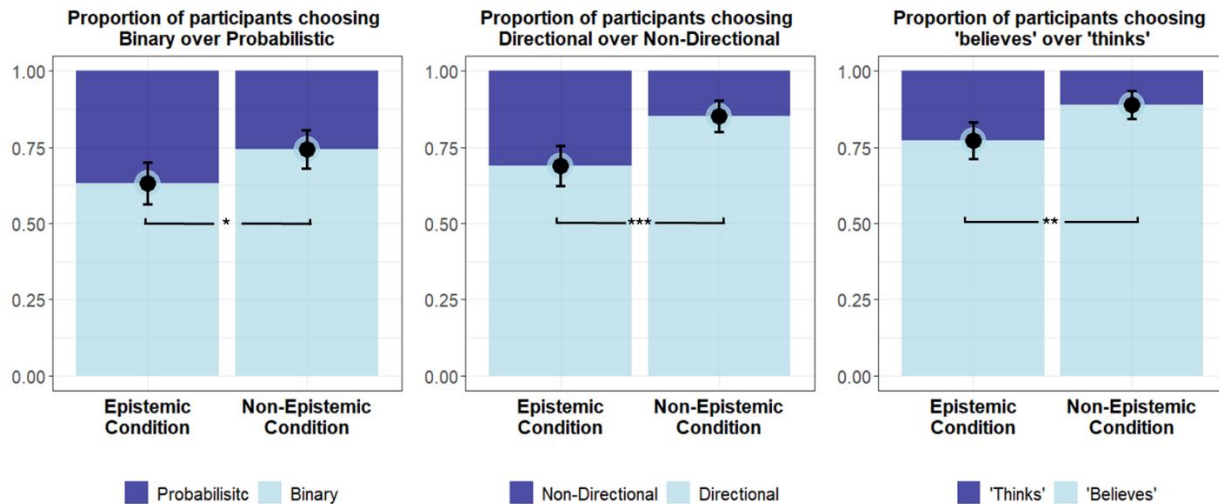
To ensure that differences across epistemic and non-epistemic conditions did not just reflect belief strength or confidence, participants were asked, "How certain do you think [Character] is that [claim]?", and they responded on a 7-point scale from -3 (Not Certain at all) to 3 (Completely Certain). Note that this measure is not redundant with our Binary/Probabilistic measure: someone could believe that it is more natural to describe a belief in binary terms even if the belief is not held with maximal certainty (see Nelson & Lombrozo, in prep).

For exploratory purposes, participants were also asked to report whether the character had good reasons for their belief and whether the belief was important to their identity; these questions and corresponding analyses are reported in [OSF](#).

### ***Results & Discussion***

To analyze our primary measures (Binary/Probabilistic, Directional/Non-Directional, and Believes/Thinks), we fit mixed-effects logistic models with condition (epistemic, non-epistemic) as a predictor and random intercepts for vignette to control for possible variability across vignettes. In all cases we coded the feature that we associated with the epistemic condition as 0 (Probabilistic, Non-Directional, Thinks) and the feature that we associated with the non-epistemic condition as 1 (Binary, Directional, Believes). All three measures conformed to our predictions (see Figure 1). Comparing the non-epistemic context to the epistemic context,

participants were significantly more likely to select a binary construal over a probabilistic construal ( $\beta = .527$ ,  $SE = .224$ ,  $p = .018$ ), to select a directional construal over a non-directional construal ( $\beta = .981$ ,  $SE = .259$ ,  $p < .001$ ), and to select ‘believes’ over ‘thinks’ ( $\beta = .8881$ ,  $SE = .289$ ,  $p = .002$ ).



**Figure 3.** Proportions of participant responses for our three signatures across the epistemic and non-epistemic conditions in Study 1. Dark blue bars represent participants choosing a probabilistic construal (first panel), non-directional control (second panel), and “thinks” (third panel). Light blue bars represent participants choosing a binary construal (first panel), directional control (second panel), and ‘believes’ (third panel). Dots represent mean proportions choosing binary, directional, and ‘believes’, with error bars representing 95% confidence intervals. \* < .05, \*\* < .01, \*\*\* < .001.

Given that the vignettes involved beliefs that were matched in content across the epistemic and non-epistemic conditions, we have some evidence that the signatures differentiate how a belief is held (the cognitive attitude) rather than belief content. However, an important concern is that the

characters across the paired vignettes might have been judged by participants as merely holding their beliefs with different levels of *certainty*. If our signatures track differences in certainty (for instance, with a binary belief simply judged to be more certain than a probabilistic belief), then our differences across conditions could reflect differences in the certainty with which a cognitive attitude is held, rather than a difference in the cognitive attitude itself. To address this concern, we analyzed our measure of certainty. A linear regression predicting certainty from condition (epistemic, non-epistemic) and vignette found that there was no significant difference between judgments about the character's certainty across the epistemic and non-epistemic conditions ( $\beta = 0.181$ ,  $SE = 0.121$ ,  $p = .138$ ). Most importantly, however, the differences that we observed across all three signatures remained significant when our certainty measure was added as a predictor to the original mixed-effects model reported above (Binary/Probabilistic:  $\beta = .488$ ,  $SE = .228$ ,  $p = .032$ ; Directional/Non-Directional:  $\beta = .942$ ,  $SE = .262$ ,  $p < .001$ ; Believes/Think:  $\beta = .881$ ,  $SE = .289$ ,  $p = .002$ ), indicating that the predicted effects were not an artifact of differences across conditions in inferred certainty.

## Study 2

Study 1 found that participants differentiated epistemic and non-epistemic beliefs in line with all three signatures. Study 2 tests whether participants make the reverse inference: When a belief is expressed as binary (vs. probabilistic), directional (vs. non-directional), or using “believes” (vs. “thinks”), are participants more likely to infer that the belief plays a non-epistemic role? To assess this we asked participants to draw inferences regarding the three aspects of functional role manipulated across the epistemic vs. non-epistemic conditions in Study

1: whether the belief aims at truth (vs. a non-epistemic value), is based on evidence, and would change in light of new evidence.

### *Participants*

Study 2 had a total sample of 732 U.S.-based adults recruited through Prolific and compensated at a rate of \$0.68 for a 4-minute study. We excluded 9 participants for failing one or more basic attention checks, for a total of 721 participants in our analyzed sample.<sup>2</sup> Sample sizes were determined by power analyses based on effect sizes obtained in pilot studies.

### *Procedure*

In a fully between-subjects design, participants read one of 15 possible vignettes, the result of crossing our three signatures (Binary/Probabilistic, Directional/Non-Directional, and Believe/Think) with five claims (“A Republican will win the next presidential election”, “Jesus was born in a Manger near Bethlehem”, “There is alien life on Earth”, “Chickens feel fear and anxiety before being sacrificed for food”, and “Artificial intelligence will never achieve human-like consciousness”). These claims were selected because they could plausibly support an epistemic or a non-epistemic cognitive attitude.

Each vignette involved two characters who agree that the given proposition is true, but express their belief differently. Below are sample vignettes illustrating the difference across contrasts:

---

<sup>2</sup> Study 2 is a composite of two studies. The first study (N=482, 8 excluded after attention checks) measured the Binary/Probabilistic and the Think/Belief conditions. The second study (N=250, 1 excluded after attention checks) measured the Directionality condition. Due to experimenter error, we only collected age and gender demographics for the latter study (age:  $m = 37.92$ ,  $\sigma = 13.00$ ; Male = 105, Female = 133, Other = 9).

Alex and Blaine were both called by the same survey company. The survey company asked them each several questions. One question concerned their take on whether [there is alien life on earth]. Here's what they both said:

Binary/Probabilistic:

Alex: “[there is alien life on earth].”

Blaine: “There’s a 98% chance that [there is alien life on earth].”

Believe/Think:

Alex: “I **believe** that [there is alien life on earth].”

Blaine: “I **think** that [there is alien life on earth].”

The vignette was slightly different for the Directional/Non-Directional contrast:

Alex and Blaine were both called by the same survey company. The survey company asked them each several questions. One question concerned their take on whether [there is alien life on earth]. Both Alex and Blaine agreed that [there is alien life on earth]. They were then asked whether any decision was involved in coming to have this belief.

Directional/Non-Directional:

Alex said, “I decided to believe that [there is alien life on earth]”

Blaine said, “I decided *whether* to believe that [there is alien life on earth]”

Participants were then asked to provide ratings concerning the epistemic and non-epistemic features of the characters' beliefs. Each participant evaluated the same three dimensions used to manipulate epistemic and non-epistemic context in Study 1: whether a character would change their mind in light of new evidence, based their belief on evidence, and aimed at truth. The lower bound of the 7-point scale was associated with the non-epistemic profile and the upper bound was associated with the epistemic profile.

*[Change of Mind]* We know that one of them (either Alex or Blaine) would be likely to change their view if presented with evidence indicating that it is not the case that [there is alien life on earth]. We also know that the other one (either Alex or Blaine) would resist changing their mind even in the face of such evidence.

Who do you think is the person who is more likely to change their mind if they were presented with evidence indicating that it is **not** the case that [there is alien life on earth]?  
(-3 definitely Alex to 3 definitely Blaine)

*[Followed the Evidence]* We know that, in arriving at their view, one of them (either Alex or Blaine) was most interested in following the evidence, no matter which view it led them to adopt. We also know that the other one (either Alex or Blaine) was most interested in adopting the particular view most consistent with their values – things like their morals, loyalties, and religious faith.



Who do you think is the person who was most interested in basing their view only on the evidence? (-3 definitely Alex to 3 definitely Blaine)

*[Aimed at Accuracy]* We know that one of them (either Alex or Blaine) is most interested in having true and accurate beliefs about the world. We also know that the other one (either Alex or Blaine) is more interested in holding the beliefs that best reflect other values, such as morality, loyalty, or faith.

Who do you think is the person who is more interested in having true and accurate beliefs about the world? (-3 definitely Alex to 3 definitely Blaine).

Because these were precisely the three aspects of functional role manipulated across the epistemic versus non-epistemic conditions in Study 1, this study was the mirror image of Study 1, looking not at inferences to each signature from epistemic versus non-epistemic role, but at inferences to epistemic versus non-epistemic role from each signature.

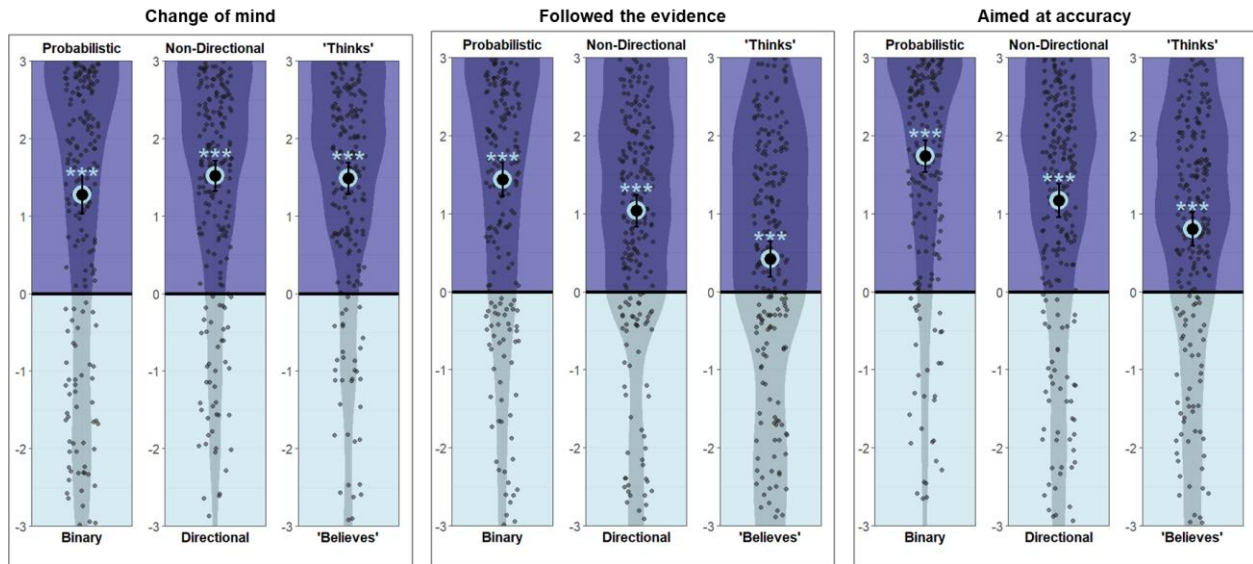
As in Study 1, participants also answered questions about the characters' certainty: "Who do you think is more certain that [there is alien life on earth]? (-3 definitely Alex to 3 definitely Blaine). We also included questions about which character had better reasons for belief and about centrality to identity. See [OSF](#) for all items and questions and corresponding analyses.

### ***Results & Discussion***

For all three measures, the character who expressed their belief using the epistemic pole of the corresponding signature (Probabilistic, Non-Directional, and 'Think') was seen as more

likely to exhibit the epistemic feature, as verified statistically with *t*-tests comparing responses to the scale midpoint. In other words, and as reflected in the positive means in Figure 4, participants judged it more likely that a character would change their mind in light of evidence, based their belief on evidence, and aimed at accuracy when they expressed their belief using the epistemic pole of our signatures (Probabilistic, Non-Directional, and ‘Thinks’) compared to the non-epistemic pole of our signatures (Binary, Directional, and ‘Believes’).

As in Study 1, we also analyzed certainty. T-tests revealed that the means for certainty were significantly below the midpoint for all conditions (Binary/Probabilistic:  $m = -1.183$ ,  $t = -28.214$ , 95% CI =  $-1.183 - -1.028$ ; Directional/Non-Directional:  $m = -1.004$ ,  $t = -112.53$ , 95% CI =  $-1.021 - -0.986$ ; Believe/Think:  $m = -1.535$ ,  $t = -258.32$ , 95% CI =  $-1.547 - -1.524$ ), indicating that participants judged the character at the non-epistemic pole (Binary, Directional, and ‘Believes’) to be more certain. Most importantly, however, the significant differences for each signature were maintained after regressing a generalized linear model with item and certainty scores as factors (to control for item variation and certainty scores) and running *t*-tests over the model’s projected data (see Table 1). As in Study 1, this suggests that the inferences drawn from each signature reflected the perceived epistemic versus non-epistemic roles of the attributed beliefs, and not their contents or levels of certainty. (See [OSF](#) for further statistics and supplementary measures.)



**Figure 4.** Mean ratings given by participants in Study 2, in the Binary/Probabilistic, Directional/Non-Directional, and Think/Believe conditions, organized by measure. Large points represent mean ratings; small points represent individual participant responses. Error bars represent 95% confidence intervals. Each row of three asterisks represents  $p < .001$  in a one-sample t-test (\*  $< .05$ , \*\*  $< .01$ , \*\*\*  $< .001$ ).

	t	n	m	CI	p
<b>Change of mind</b>					
Binary/Probabilistic	20.039	235	1,279	1.153 – 1.405	***
Directional/Non-Directional	33.232	246	1,522	1.432 – 1.612	***
Believe/Think	32.416	236	1,489	1.398 – 1.579	***
<b>Followed the evidence</b>					
Binary/Probabilistic	83.662	235	1,741	1.700 – 1.782	***
Directional/Non-Directional	66.657	246	1,174	1.139 – 1.208	***
Believe/Think	75.026	236	0,805	0.784 – 0.827	***
<b>Aimed at accuracy</b>					

Binary/Probabilistic	58.010	235	1,440	1.391 – 1.489	***
Directional/Non-Directional	46.448	246	1,040	0.996 – 1.084	***
Believe/Think	18.134	236	0,421	0.376 – 0.467	***

Table 1. Results of a one-sample t-test over projected data, modelled with a generalized linear model for each measure with item and certainty scores as factors:  $t$  = t-statistic,  $n$  = number of participants in the condition,  $m$  = mean, CI = 95% Confidence Interval,  $p$  =  $p$ -value significance (\* < .05, \*\* < .01, \*\*\* < .001).

### Study 3

Study 3 tests an important consequence of the hypothesis that theory of mind posits more than one kind of belief: epistemic and non-epistemic cognitive attitudes should be associated with different predictions about behavior. We test this by presenting participants with the signatures of cognitive attitude manipulated in Study 2 and eliciting judgments about behavior. In particular, we predicted that participants would be more likely to judge that characters who express their beliefs employing the epistemic poles of our signatures (Probabilistic, Non-Directional, and ‘Think’) would perform truth-dependent actions, as opposed to symbolic actions, than characters who express their beliefs employing the non-epistemic poles of our signatures (Binary, Directional, and ‘Believe’).

By truth-dependent actions, we mean actions whose success depends on the truth of the belief(s) on which they are based. For example, the success of anonymously betting \$1000 that the San Francisco 49ers will win the Super Bowl depends on whether it is ultimately true or false that the 49ers win the Super Bowl. The success of a symbolic action, on the other hand, depends on the signaling consequences of said action, independently of whether the belief on which it is based is true. For example, posting on Facebook that the San Francisco 49ers will win the Super

Bowl can be successful in signaling one's allegiances whether or not the 49ers win the Super Bowl. The idea of "symbolic actions" comes from research on cultural psychology, which suggests that certain culturally significant actions can gain symbolic functions that aim to negotiate, signal, or prompt culturally significant responses from relevant community members (Boesch, 1991; Sperber, 1995; Straub, 2021; see also Westra, 2023). Here we use the distinction between symbolic actions and truth-dependent actions to investigate whether and how attributing epistemic versus non-epistemic beliefs affects predictions about behavior, a core function of theory of mind.

### ***Participants***

Study 3 had a total sample of 751 U.S.-based adults recruited through Prolific and compensated \$0.33 for a 2-minute study. We excluded 11 participants for failing one basic attention check, for a total of 740 participants in our analyzed sample (age:  $m = 39.28$ ,  $\sigma = 13.13$ ; Male = 355, Female = 372, Other = 13). Sample sizes were determined by power analyses based on effect sizes obtained in pilot studies.

### ***Procedure***

We modified the paradigm from Study 2 to test whether our three signatures elicit different inferences about people's *truth-dependent* and *symbolic* actions, respectively.

Participants were presented with one of 30 vignettes, the result of crossing three contrasts (binary/probabilistic, directional/non-directional, believe/think), the same five propositions from Study 2, and two action types: truth-dependent or symbolic. As in Study 2, participants learned about a pair of characters who expressed their belief regarding some proposition following one

of our three signatures. Unique to Study 3, participants were then told that one of those characters performed an action (truth-dependent or symbolic), and they were asked to judge which of the two characters was more likely to have performed that action. Below is our behavior prediction measure illustrated in the Binary/Probabilistic condition for a particular proposition:

Alex and Blaine were both called by the same survey company. The survey company asked them each several questions. One question concerned their take on whether [there is alien life on earth]. Here's what they both said:

Binary/Probabilistic:

Alex: “[there is alien life on earth].”

Blaine: “There's a 98% chance that [there is alien life on earth].”

[Action prediction] We know that, sometime later, one of the two (Alex or Blaine) did the following:

*[Truth-Dependent]* Anonymously donated \$1000 to a scientific organization investigating how to combat potential threats to humans from alien life on earth.

*[Symbolic]* Became an avid spokesperson for the organization ‘We are not alone’ that supports alien enthusiasts all around the country.

Who do you think did this? (-3: Definitely Alex to 3: Definitely Blaine)

As a control, participants also answered the same certainty question as in Study 2. See [OSF](#) for all items and questions and corresponding analyses.

### ***Results & Discussion***

Our primary prediction was that participants would be more likely to judge that truth-dependent actions, as compared to symbolic actions, were performed by characters who expressed their beliefs using the epistemic signatures (Probabilistic, Non-Directional, ‘Think’). We therefore expected that, for each of our signature contrasts, the ratings for our action prediction measure would be significantly higher for truth-dependent actions, compared to symbolic actions.

To test the effect of action type in each of our signatures, we separated our sample into three subsamples that corresponded to the action prediction ratings given for each of our signatures (Binary/Probabilistic, Directional/Non-Directional, and ‘Think’/‘Believe’). Then, for each signature, we conducted separate mixed-effects models with the action prediction rating (-3: Alex – 3: Blaine) as a dependent variable, action type (truth-dependent coded as 0/symbolic coded as 1) as the fixed effect, and vignette as a random intercept to control for variability across vignettes. Results showed that for each signature, participants gave significantly higher ratings for truth-dependent actions than for symbolic actions, meaning that truth-dependent actions were judged closer to the epistemic pole of each signature than were symbolic actions (the coefficients for action type were as follows across conditions: Binary/Probabilistic:  $\beta = .762$ ,  $SE = .239$ ,  $p = .0016$ ; Directional/Non-Directional:  $\beta = .911$ ,  $SE = .220$ ,  $p < .001$ ; Think/Believe:  $\beta = .567$ ,  $SE =$

.219,  $p = .010$ ). Figure 5 illustrates this pattern, with mean ratings for the truth-dependent action higher (closer to the epistemic pole) than mean ratings for symbolic actions given each of our signatures.<sup>3</sup>

Notably, all mean ratings in Study 3 fell numerically below the scale midpoint, indicating that despite the predicted (and observed) variation across symbolic and truth-dependent actions, participants generally thought that the character who expressed the non-epistemic pole of our signatures (Binary, Directional, and ‘Believe’) was more likely to engage in both kinds of action. This is likely a consequence of attributing greater *certainty* to this character. T-tests comparing certainty ratings to the scale midpoint found that in all cases, participants judged the character who expressed the non-epistemic pole of our signatures (Binary, Directional, and ‘Believe’) to be more certain, (Binary/Probabilistic:  $M = -1.364$ ,  $SE = 0.128$ ,  $p < .001$ ; Directional/Non-Directional:  $M = -1.482$ ,  $SE = 0.110$ ,  $p < .001$ ; ‘Believes’/‘Thinks’:  $M = -1.743$ ,  $SE = 0.103$ ,  $p < .001$ ). Compensating for this approximately 1-point offset in certainty favoring the non-epistemic character, the mean judgments in Figure 5 would straddle the scale mid-point, with symbolic actions falling on the side of the non-epistemic character, but truth-dependent actions falling on the side of the epistemic character.

Importantly, we also verified that differences in certainty could not account for the difference across symbolic and truth-dependent actions. While we did find that participants gave significantly lower certainty ratings for truth-dependent actions than for symbolic actions in the Binary/Probabilistic condition ( $\beta = -0.524$ ,  $SE = .249$ ,  $p = .036$ ), the effects of condition


---

<sup>3</sup> In an exploratory analysis, we also investigated whether the effect of action type (Truth-Dependent vs. Symbolic) was moderated by signature (Binary/Probabilistic, Directional/Non-Directional, Think/Believe). To do so we fit a mixed-effects model with action prediction rating as the dependent variable, action type and signature as fixed effects, and vignette as a random intercept. This analysis included an interaction term between action type and signature. A Type III ANOVA on the interaction term revealed that it was not significant ( $F(2,730.63) = 0.727$ ,  $p = 0.484$ ), indicating that the effect of action type on ratings was comparable across the three signatures.



reported above remained significant when adding certainty to the model as a control

(Binary/Probabilistic:  $\beta = 1.038$ ,  $SE = .196$ ,  $p < .001$ ; Directional/Non-Directional:  $\beta = 1.091$ ,  $SE = .197$ ,  $p < .001$ ; Think/Believe:  $\beta = .427$ ,  $SE = .195$ ,  $p = .030$ ).



**Figure 5.** Mean ratings in Study 3 for Symbolic and Truth-Dependent actions in the Binary/Probabilistic, Directional/Non-Directional, and Believe/Think conditions. Error bars represent 95% confidence intervals. \*  $< .05$ , \*\*  $< .01$ , \*\*\*  $< .001$ .

### General Discussion

Is there more than one category of “belief” represented in people’s theory of mind? Study 1 suggests that people systematically differentiate beliefs that are held with epistemic versus non-epistemic cognitive attitudes: the former are more likely than the latter to be construed as probabilistic (vs. binary), judged to be the outcome of non-directional (vs. directional) control, and expressed using the verb “think” (vs. “believe”). Study 2 documents novel patterns in people’s inferences from these signatures to judgments about a belief’s epistemic versus non-epistemic role, and Study 3 extends these results to people’s predictions concerning truth-dependent versus symbolic action.

Importantly, we observed these effects despite controlling for both the *content* and *certainty* of beliefs, suggesting that our participants made distinctions between cognitive *attitudes* – between kinds of believing. That said, content and certainty may well be cues to cognitive attitude in many real-world cases. For instance, we expect that someone who asserts “the president is alive” will be judged to hold an epistemic belief more readily than someone who asserts “Jesus is alive,” and will be so judged using content as a cue. Our decision to hold belief content fixed – and to control for effects of certainty – was methodologically important in testing our hypothesis that theory of mind differentiates epistemic and non-epistemic cognitive *attitudes*, as posited by Belief Pluralism. This is a key aspect of our proposal, and one way in which it differs from prior theoretical and empirical work introducing distinctions between kinds of belief (Abelson, 1986; Apperly & Butterfield, 2009; Buckwalter et al., 2015; Skitka et al., 2021; Sperber, 1997). That said, our methodological choice to hold content fixed likely explains why our effects are relatively small. It is reasonable to expect that the differences across our signatures would be considerably larger if content and certainty were allowed to vary as well.

Considerations about content could also explain a notable pattern in our data. While our results indicate that participants differentiate epistemic and non-epistemic features of beliefs, in Study 1 a majority of participants described characters’ beliefs using the non-epistemic poles of our signatures – that is, a majority of beliefs were described as binary, under directional control, and using the word ‘believes’. This could be an artifact of our belief contents, which were selected for their plausibility in both epistemic and non-epistemic contexts (e.g., believing that a friend is innocent). This may have provided content-based cues that beliefs were non-epistemic, with our experimental manipulations resulting in systematic variation around judgments that were skewed non-epistemic. Had we instead used belief contents that do not readily support a

non-epistemic construal (“The year is 2025,” “The atomic number of Gold is 79,” “Water is wet,” etc.), we would expect judgments skewed in the other direction. And with the liberty to vary content alongside epistemic versus non-epistemic context, we would expect much larger effects, with epistemic beliefs generating a majority of responses at the epistemic poles (Probabilistic, Non-directional, and ‘Think’), and non-epistemic beliefs at the non-epistemic poles (Binary, Directional, and ‘Believe’). Exploring such interactions between content and other cues to cognitive attitude would be a valuable direction for future research.

### **Relationship to prior empirical work**

The present findings contribute to recent research that offers a refined picture of the psychological constructs involved in theory of mind (Ho et al., 2022; Navarro, 2022; Schaafsma et al., 2015; Spaulding, 2018; Wellman & Miller, 2008; Westra, 2018, 2023). In particular, our findings support pluralism about “belief” in intuitive psychology and offer one way to characterize the two broad belief kinds involved: in terms of their epistemic or non-epistemic roles.

The hypothesis that theory of mind involves more than one kind of belief can explain several findings from prior research. For example, many studies have found differences between religious and scientific beliefs, such that the former are more likely to be based on affiliative, moral, and intuitive considerations (all plausibly non-epistemic), while the latter are more often grounded in evidence (a paradigmatically epistemic consideration) (Metz et al., 2023; Shtulman, 2013). In both children and adults, religious beliefs are more likely than scientific or factual beliefs to be considered subjective, in the sense that two people who disagree can both be right (Heiphetz et al., 2013; Liquin et al., 2020; Metz et al., 2023; Wainryb et al., 2004). On the

assumption that religious beliefs, compared to scientific beliefs, typically play more non-epistemic roles, these findings suggest that beliefs are systematically differentiated along the lines that our distinction between epistemic and non-epistemic belief would predict. However, this prior work has contrasted beliefs that differ in content (i.e., religious versus scientific content), so while the evidence is consistent with our posited distinction in *cognitive attitudes* (or kinds of *believing*), the findings reported here are the first to find the posited differentiation while controlling for content, and thus offer the strongest support for a distinction in attitude.

Only two prior results, to our knowledge, have similarly isolated judgments of belief attitude from effects of belief content. In Study 4 of Heiphetz, Landers, and Van Leeuwen (2021), participants were presented with vignettes in which individuals endorse a given content, but where the context is matter-of-fact versus religious. For example, in one matter-of-fact vignette context Kerry believes that aspirin is not a cure for headaches because she had tried it many times before with no success; while in the related religious vignette context, Terry believes the same thing because the church she is a member of teaches that prayer and not medicine cures medical ills. Participants were asked to choose between “thinks” and “believes” to fill in the blank in the following sentence: “*Kerry/Terry always refused the aspirin her friends offered, because she \_\_\_\_\_ that aspirin is not a cure.*” Participants were much more likely to use “thinks” for matter-of-fact vignette contexts (62% of the time), but “believes” for religious vignette contexts (74% of the time), even though the attributed contents were matched (in this case, *that aspirin is not a cure*). Study 3 of Van Leeuwen et al. (2021) replicated this result across five languages and cultural contexts.

Our present work builds on these findings in three important ways: (i) by suggesting that the relevant (and more general) distinction is between a belief’s epistemic versus non-epistemic

roles, rather than being just between matter-of-fact versus religious contexts (which is arguably a special case of our more general distinction), (ii) by identifying two additional signatures of attitude type beyond uses of “thinks” vs. “believes,” and (iii) by investigating inferences from signatures to the epistemic and non-epistemic features and consequences of belief (as in Studies 2 and 3), not only the conditions under which different kinds of belief attitudes are attributed (as in Study 1).

### **Alternative hypotheses**

It is worth considering how the Belief Pluralism hypothesis differs from three alternatives. First, consider what we call the “masked” belief hypothesis. On this view, there is only one kind of belief attitude, but typical properties of this single belief type (such as responsiveness to evidence) can be “masked” by social or motivational pressures (Flores, n.d.; Helton, 2018). To translate this hypothesis to theory of mind, it could be that laypeople represent others as having a single type of belief, but recognize that said type of belief can be distorted or masked by non-epistemic considerations, such as emotion or conflicts of interest (cf. Pronin & Hazel, 2023). On this hypothesis, what we call attributions of non-epistemic beliefs are just attributions of distorted (masked) epistemic beliefs.

A second alternative can be called the “enhanced” belief hypothesis. On this view, people’s theory of mind does trace separable kinds of belief attitudes. However, beliefs fulfilling epistemic functions are more basic and can be “enhanced” into more complex beliefs that are sensitive to non-epistemic considerations. Thus, enhanced beliefs (which we are hypothesizing to be non-epistemic beliefs) entail the properties of basic epistemic beliefs but not the other way

around.<sup>4</sup> To translate into our terms, this hypothesis holds there are no distinctive non-epistemic beliefs represented in lay theory of mind: instead there are only epistemic beliefs with and without non-epistemic enhancement.

Finally, consider the “masquerading” belief hypothesis. On this view, people’s intuitive theory of mind does recognize a distinctive cognitive attitude responsive to non-epistemic factors. However, this attitude only passes as a belief without being belief in its own right.<sup>5</sup> Only purely epistemic doxastic attitudes deserve the label “belief.” On this view, only what we call an epistemic belief can be properly called a belief.

In different ways, these views all attempt to do without the existence of what we call non-epistemic beliefs. Applied to our study, the masked belief hypothesis would hold that our signatures do not track ascriptions of distinct varieties of belief, but rather people’s sensitivity to how much beliefs (conceived of as one sort of thing) may have been masked or corrupted by non-epistemic considerations. Similarly, the enhanced belief hypothesis would hold that our non-epistemic signatures are not tracking ascriptions of a differentiable non-epistemic belief attitude, but of an enhanced belief that builds upon, and thus entails, a more basic epistemic belief. Finally, the masquerading belief hypothesis would hold that our signatures distinguish genuinely different cognitive attitudes, only one of which can be called belief. Although our studies were not designed to differentiate these hypotheses, we think a few considerations speak in favor of Belief Pluralism, or the need to posit non-epistemic beliefs. (That said, we will not be arguing that there are no such things as masked, enhanced, or masquerading beliefs, or their theory of

---

<sup>4</sup> Arguably Buckwalter et al. (2015) defend this view. They posit notions of a “thick” kind of belief that adds non-epistemic features to a “thin” notion of belief that is characterized by epistemic features; thus the properties of thin belief are entailed by thick belief but not *vice versa*. A wrinkle to their version, however, is that they hold that thin belief is likely conceived as involuntary, but thick belief is (or can be) voluntary, which is puzzling since thick beliefs are supposed to also have the properties of the thin beliefs they entail (which presumably includes involuntariness). See below for more discussion of the enhanced belief view.

<sup>5</sup> We thank an anonymous reviewer for this suggestion.

mind representations; those may *also* exist. We are only arguing here that there are considerable data that the hypotheses under discussion struggle to explain, but that Belief Pluralism explains well.)

One reason to doubt the sufficiency of the masked and masquerading accounts comes from recent work demonstrating that (many) people classify non-epistemic reasons as legitimate reasons for at least some of their own or others' beliefs, not as illegitimate distortions that compromise a belief's importance or justification (Cusimano & Lombrozo, 2021b, 2021a, 2023; Metz et al., 2023). For instance, Cusimano and Lombrozo (2021a) found that many participants thought characters should adopt the belief that was more loyal, rather than the belief that was more strongly supported by evidence, and moreover judged a character's belief to be more *justified* when it aligned with loyalty rather than evidence. Metz et al. (2023) found that many participants judged the moral value of a belief to be a "good reason" for belief. These findings suggest that non-epistemic considerations are routinely recognized and even condoned in evaluating beliefs, challenging the masked belief hypothesis (on which non-epistemic influences *distort* belief) and the masquerading belief hypothesis (according to which masquerading beliefs would likely be judged by the same epistemic standards as the things they are masquerading as, i.e., epistemic beliefs). More generally, the masked and masquerading hypotheses on their own have difficulty making sense of the interesting variation in normative stances people take toward (what we call) epistemic versus non-epistemic belief.

One reason to doubt the sufficiency of the enhanced belief account comes from its contention that non-epistemic beliefs have all the properties of epistemic beliefs (plus a few more). Westra (2023), who advances a theoretical position similar to our own, notes that this claim is in tension with the results of Heiphetz et al. (2013) mentioned above. Recall that

Heiphetz et al. (2013) contrasted responses to disagreements involving factual beliefs, which are presumably epistemic, with those involving ideological beliefs, which are presumably non-epistemic. Participants often judged that two people who endorse conflicting ideological beliefs “can both be right.” This suggests that such beliefs are not seen as aiming at truth; thus they have features that depart from, rather than build on, those of epistemic beliefs.

In some ways the masquerading belief hypothesis is the closest to our own, insofar as it recognizes two functionally distinct attitudes. Indeed, if it admits that what we call non-epistemic belief is a functionally distinguishable attitude that can be held independently from epistemic beliefs, we fear the contention is merely terminological. By using the label ‘belief’ to denote these attitudes we aim to mark that people’s theories of mind recognize distinct doxastic states that fulfill some of the core functions of belief (e.g., they represent the world as actual, take in reasons, and serve in our inferences about people’s behavior) but are still different from the classic epistemic beliefs posited by philosophers and psychologists. Once this is accepted, not much more hangs on the decision to label it a kind of ‘belief’ or not.

### **Open questions**

Do people’s representation of epistemic and non-epistemic beliefs differ in kind, or only in degree? While we have framed our distinction between epistemic and non-epistemic beliefs as a difference in kind, we are sympathetic to the possibility that people’s representations of others’ epistemic and non-epistemic beliefs could instead describe two ends of a continuum (or two clusters in a multi-dimensional space), with many beliefs falling in between. In fact, some of our additional analyses, reported in Supplementary Materials, suggest that it may be useful to think of beliefs as falling within a multi-dimensional space, as opposed to lying along a single



continuum. For example, while we find that responses along our three signatures tend to be correlated, the associations are modest. In preliminary work, we also explored an additional dimension – whether a belief’s content is regarded as an objective matter of fact – and found that it did not pattern with choices between think vs. believe and binary vs. probabilistic construal (Vesga et al., 2024). These findings suggest that representations of belief are not unidimensional.

On the other hand, our exploratory measures revealed reliable relationships between our manipulation of belief type and additional epistemic versus non-epistemic considerations: we found that participants were more likely to judge the characters in the epistemic (vs. non-epistemic) condition to hold their belief for good reasons, and the characters in the non-epistemic (vs. epistemic) condition to regard the belief as important (see Supplementary Materials). These measures were also reliably associated with some of our signatures. This suggest that additional factors – such as justification and importance to identity – might reliably relate to the dimension or dimensions along which attributions of belief can vary. We suspect that in real-world contexts, where belief content, belief confidence, and other factors are not controlled, beliefs will tend to cluster along correlated dimensions (for instance, the beliefs that have our non-epistemic signatures will also tend to be high-stakes and important to identity), but that in more controlled scenarios, it will be possible to tease some of these dimensions apart.

A related question for future research concerns the relationships between epistemic and non-epistemic beliefs. Many or even most people hold beliefs to norms of consistency (one should not believe both  $P$  and  $not-P$ ). Yet it appears to be the case that inconsistency is psychologically common, both among beliefs in an individual’s mind (Porot & Mandelbaum, 2021; Sommer et al., 2022) and among people’s folk theories more generally (e.g., Davoodi & Lombrozo, 2023; Shtulman & Lombrozo, 2016). Does that mean that people are routinely in

violation of such norms? Importantly, if beliefs come in multiple varieties, it may be that the inconsistencies among “beliefs” are less evidence of irrationality than they are of attitude compartmentalization: while there may be inconsistencies across an agent’s non-epistemic and epistemic beliefs, her epistemic beliefs considered in themselves may be highly consistent. For example, Metz et al. (2020) found that some participants simultaneously endorsed inconsistent forms of evolution and creationism. While that appears to be a flagrant inconsistency, one participant reported that, due to his upbringing, “I *believe* we were created by a God,” but as a biologist, “I *think* [human origins are] based in evolutionary principle” (emphasis added). Thus, with pluralism in cognitive attitudes, perhaps individuals can maintain some *beliefs* while simultaneously *thinking* alternatives are true (at least with some probability). Although maintaining inconsistent “beliefs” may well come with risks, it could also support more flexible reasoning and behavior. And importantly, if non-epistemic beliefs are compartmentalized in relation to epistemic beliefs, the latter may typically constitute a more rationally consistent “web” than it might appear at first brush. Understanding the relationships between epistemic and non-epistemic beliefs could thus have significant implications for human rationality more generally.

Another relevant question concerns the nature of what we call non-epistemic beliefs. Ours is a general contrastive claim holding that people recognize beliefs with epistemic functions and beliefs that lack these epistemic functions. In this sense, our hypothesis is more general than Westra’s (2023) hypothesis that people’s theories of mind recognize non-epistemic beliefs with specifically symbolic functions. While we admit that symbolic functions are plausibly important, we are open to the possibility that non-epistemic beliefs fulfill other functions, such as emotional regulation and identity preservation. Hence, it is worth exploring in future research whether

people's intuitive theory of mind has a unified category of non-epistemic beliefs, or, alternatively, whether this category is systematically subdivided into different but stable kinds of non-epistemic beliefs. The possibility of multiple belief-like cognitive attitudes also raises new questions for subdisciplines of psychology that have been traditionally concerned with theory of mind. For example, developmental psychologists have charted the emergence of theory of mind using paradigms like the false belief task (e.g., Wellman, 2018), which implicitly involve an epistemic notion of belief. One possibility is that as soon as children demonstrate some understanding of belief, they differentiate epistemic from non-epistemic believing. Another possibility, which we judge more likely, is that the differentiation between epistemic and non-epistemic belief we see in our adult sample emerges over the course of development. Indirect support for this possibility comes from research on children's trust in others' testimony: for instance, 3-year-olds do not reliably favor a more knowledgeable informant (an epistemic characteristic) over one who is nice or prosocial (a non-epistemic characteristic), but 4- and 5-year-olds typically do (Tong et al., 2020). More generally, research on selective trust suggests that over the course of development, children become sensitive to differences between informants along both epistemic dimensions (such as past accuracy) and non-epistemic dimensions (such as social group membership), and learn to coordinate these sources of information (for relevant discussion, see Koenig et al., 2022). This suggests that with age, children develop a stronger differentiation between the epistemic and non-epistemic, including the relevance of each kind of information for different judgments.

## **Conclusions**

To sum up, we document an important nuance in “belief-desire” psychology: in representing and reasoning about the beliefs of others, people track distinctions in cognitive attitudes, systematically differentiating beliefs that play predominantly epistemic roles from those that play important non-epistemic roles. This supports the Belief Pluralism hypothesis that theory of mind represents more than one variety of believing. It also offers indirect support for the Varieties of Belief hypothesis, that the mind itself supports more than one flavor of believing.

## **Constraints on Generality**

Our sample was restricted to U.S.-based adults participating in online experiments. That said, we drew from a participant population with considerable variation in age, income, and education. Hence, it is plausible that the patterns in our data would generalize to the broader population of North American adults. But we are not yet licensed to draw conclusions (i) about whether those patterns would also surface in other cultures or (ii) about the developmental timeline on which the patterns of differentiation we document emerge. On the first issue (cross-cultural comparison), Van Leeuwen et al. (2021) give reason for optimism, since participants in the Fante, Thai, Mandarin, and Bislama versions of their studies displayed sensitivity to analogues of our thinks-believes signature in their respective languages. But experimental cross-cultural studies focusing on our other signatures and on action type have yet to be done, so it is worth exercising caution. As for the developmental question (at what point does the differentiation we document emerge?), the existing literature on theory of mind development may furnish some guidance as to which hypotheses in this space are plausible, but any specific views at this point would still be speculative.

## Acknowledgments

The authors would like to thank Dan Weiskopf and Verónica Sánchez Gómez for helpful suggestions in conversation, audiences at the 50<sup>th</sup> Conference for the Society for Philosophy and Psychology 2024, and reviewers for helpful comments.

## References

- Abelson, R. P. (1986). Beliefs Are Like Possessions. *Journal for the Theory of Social Behaviour*, 16(3), 223–250. <https://doi.org/10.1111/j.1468-5914.1986.tb00078.x>
- Alston, W. P. (1988). The Deontological Conception of Epistemic Justification. *Philosophical Perspectives*, 2, 257. <https://doi.org/10.2307/2214077>
- Apperly, I. A., & Butterfill, S. A. (2009). Do humans have two systems to track beliefs and belief-like states?. *Psychological review*, 116(4), 953.
- Arpaly, N., & Brinkerhoff, A. (2018). Why Epistemic Partiality Is Overrated. *Philosophical Topics*, 46(1), 37–52.
- Audi, R. (2001). Doxastic Voluntarism and the Ethics of Belief. In M. Steup (Ed.), *Knowledge, Truth, and Duty* (1<sup>st</sup> ed., pp. 93–112). Oxford University Press New York.  
<https://doi.org/10.1093/0195128923.003.0007>
- Bendaña, J., & Mandelbaum, E. (2021). The fragmentation of belief. In *The fragmented mind* (pp. 78–107). Oxford University Press. <https://doi.org/10.1093/oso/9780198850670.001.0001>
- Bergamaschi Ganapini, M. (2020). Confabulating Reasons. *Topoi*, 39(1), 189–201.  
<https://doi.org/10.1007/s11245-018-09629-y>
- Bortolotti, L. (2010). *Delusions and Other Irrational Beliefs*. Oxford University Press.
- Buckwalter, W., Rose, D., & Turri, J. (2015). Belief through Thick and Thin. *Noûs*, 49(4), 748–775.  
<https://doi.org/10.1111/nous.12048>

- Campbell, T. H., & Kay, A. C. (2014). Solution aversion: On the relation between ideology and motivated disbelief. *Journal of Personality and Social Psychology, 107*(5), 809–824.  
<https://doi.org/10.1037/a0037963>
- Chinn, C. A., Rinehart, R. W., & Buckland, L. A. (2014). Epistemic cognition and evaluating information: Applying the AIR model of epistemic cognition. In *Processing inaccurate information: Theoretical and applied perspectives from cognitive science and the educational sciences* (pp. 425–453). The MIT Press. <https://doi.org/10.7551/mitpress/9737.001.0001>
- Cohen, G. L. (2003). Party Over Policy: The Dominating Impact of Group Influence on Political Beliefs. *Journal of Personality and Social Psychology, 85*(5), 808–822. <https://doi.org/10.1037/0022-3514.85.5.808>
- Cusimano, C., & Goodwin, G. P. (2019). Lay beliefs about the controllability of everyday mental states. *Journal of Experimental Psychology: General, 148*(10), 1701–1732.  
<https://doi.org/10.1037/xge0000547>
- Cusimano, C., & Goodwin, G. P. (2020). People judge others to have more voluntary control over beliefs than they themselves do. *Journal of Personality and Social Psychology, 119*(5), 999–1029.  
<https://doi.org/10.1037/pspa0000198>
- Cusimano, C., & Lombrozo, T. (2021a). Morality justifies motivated reasoning in the folk ethics of belief. *Cognition, 209*, 104513. <https://doi.org/10.1016/j.cognition.2020.104513>
- Cusimano, C., & Lombrozo, T. (2021b). Reconciling scientific and commonsense values to improve reasoning. *Trends in Cognitive Sciences, 25*(11), 937–949.  
<https://doi.org/10.1016/j.tics.2021.06.004>
- Cusimano, C., & Lombrozo, T. (2023). People recognize and condone their own morally motivated reasoning. *Cognition, 234*, 105379. <https://doi.org/10.1016/j.cognition.2023.105379>
- Davoodi, T., Jamshidi-Sianaki, M., Abedi, F., Payir, A., Cui, Y. K., Harris, P. L., & Corriveau, K. H. (2019). Beliefs about religious and scientific entities among parents and children in iran. *Social*

- Psychological and Personality Science*, 10(7), 847–855.  
<https://doi.org/10.1177/1948550618806057>
- Davoodi, T., & Lombrozo, T. (2022). Explaining the existential: Scientific and religious explanations play different functional roles. *Journal of Experimental Psychology: General*, 151(5), 1199–1218.  
<https://doi.org/10.1037/xge0001129>
- Douglas, B. D., Ewell, P. J., & Brauer, M. (2023). Data quality in online human-subjects research: Comparisons between MTurk, Prolific, CloudResearch, Qualtrics, and SONA. *PLoS ONE*, 18(3), e0279720.
- Evans, J. S. B. T. (2008). Dual-processing accounts of reasoning, judgment, and social cognition. *Annual Review of Psychology*, 59, 255–278. <https://doi.org/10.1146/annurev.psych.59.103006.093629>
- Flores, C. (n.d.). Resistant Beliefs, Responsive Believers. *Journal of Philosophy*.
- Gilovich, T., & Regan, D. T. (1986). The actor and the experiencer: Divergent patterns of causal attribution. *Social Cognition*, 4(3), 342–352. <https://doi.org/10.1521/soco.1986.4.3.342>
- Goldman, A. I. (2006). *Simulating minds: The philosophy, psychology, and neuroscience of mindreading*. Oxford University press.
- Heiphetz, L., Landers, C. L., & Van Leeuwen, N. (2021). Does think mean the same thing as believe? Linguistic insights into religious cognition. *Psychology of Religion and Spirituality*, 13(3), 287–297. <https://doi.org/10.1037/rel0000238>
- Helton, G. (2018). If You Can't Change What You Believe, You Don't Believe It. *Noûs*, 54(3), 501–526.  
<https://doi.org/10.1111/nous.12265>
- Hieronymi, P. (2006). Controlling Attitudes. *Pacific Philosophical Quarterly*, 87(1), 45–74.  
<https://doi.org/10.1111/j.1468-0114.2006.00247.x>
- Ho, M. K., Saxe, R., & Cushman, F. (2022). Planning with Theory of Mind. *Trends in Cognitive Sciences*, 26(11), 959–971. <https://doi.org/10.1016/j.tics.2022.08.003>

- Hofer, B. K., & Pintrich, P. R. (1997). The development of epistemological theories: Beliefs about knowledge and knowing and their relation to learning. *Review of Educational Research*, 67(1), 88–140. <https://doi.org/10.2307/1170620>
- Jong, J. (2018). Beliefs are Object-Attribute Associations of Varying Strength. *Contemporary Pragmatism*, 15(3), 284–301. <https://doi.org/10.1163/18758185-01503002>
- Kahan, D. M. (2015). Climate-Science Communication and the Measurement Problem. *Political Psychology*, 36(S1), 1–43. <https://doi.org/10.1111/pops.12244>
- Kelly, T. (2003). Epistemic Rationality as Instrumental Rationality: A Critique. *Philosophy and Phenomenological Research*, 66(3), 612–640. <https://doi.org/10.1111/j.1933-1592.2003.tb00281.x>
- Koenig, M. A., Li, P. H., & McMyler, B. (2022). Interpersonal trust in children’s testimonial learning. *Mind & Language*, 37(5), 955–974. <https://doi.org/10.1111/mila.12361>
- Luhrmann, T. M. (2012). *When God Talks Back: Understanding the American Evangelical Relationship with God* (Reprint edition). Vintage.
- Malle, B. F., & Knobe, J. (1997). The folk concept of intentionality. *Journal of Experimental Social Psychology*, 33(2), 101–121. <https://doi.org/10.1006/jesp.1996.1314>
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information* (S. Ullman, Ed.). MIT Press.
- Metz, S. E., Liquin, E. G., & Lombrozo, T. (2023). Distinct Profiles for Beliefs About Religion Versus Science. *Cognitive Science*, 47(11), e13370. <https://doi.org/10.1111/cogs.13370>
- Navarro, E. (2022). What is theory of mind? A psychometric study of theory of mind and intelligence. *Cognitive Psychology*, 136, 101495. <https://doi.org/10.1016/j.cogpsych.2022.101495>
- Ossipov, M. H., Dussor, G. O., & Porreca, F. (2010). Central modulation of pain. *The Journal of Clinical Investigation*, 120(11), 3779–3787. <https://doi.org/10.1172/JCI43766>



- Porot, N., & Mandelbaum, E. (2021). The science of belief: A progress report. *WIREs Cognitive Science*, 12(2), e1539. <https://doi.org/10.1002/wcs.1539>
- Pronin, E. (2007). Perception and misperception of bias in human judgment. *Trends in Cognitive Sciences*, 11(1), 37–43. <https://doi.org/10.1016/j.tics.2006.11.001>
- Pronin, E., & Hazel, L. (2023). Humans' bias blind spot and its societal significance. *Current Directions in Psychological Science*, 32(5), 402–409. <https://doi.org/10.1177/09637214231178745>
- Quilty-Dunn, J., & Mandelbaum, E. (2018). Against dispositionalism: Belief in cognitive science. *Philosophical Studies*, 175(9), 2353–2372. <https://doi.org/10.1007/s11098-017-0962-x>
- Ranney, M. A., & Clark, D. (2016). Climate Change Conceptual Change: Scientific Information Can Transform Attitudes. *Topics in Cognitive Science*, 8(1), 49–75. <https://doi.org/10.1111/tops.12187>
- Rivers, M. L. (2021). Metacognition About Practice Testing: A Review of Learners' Beliefs, Monitoring, and Control of Test-Enhanced Learning. *Educational Psychology Review*, 33(3), 823–862. <https://doi.org/10.1007/s10648-020-09578-2>
- Rutjens, B. T., & Preston, J. L. (2020). Science and religion: A rocky relationship shaped by shared psychological functions. In *The science of religion, spirituality, and existentialism* (pp. 373–385). Elsevier Academic Press. <https://doi.org/10.1016/B978-0-12-817204-9.00027-5>
- Schaafsma, S. M., Pfaff, D. W., Spunt, R. P., & Adolphs, R. (2015). Deconstructing and reconstructing theory of mind. *Trends in Cognitive Sciences*, 19(2), 65–72. <https://doi.org/10.1016/j.tics.2014.11.007>
- Skitka, L. J., Hanson, B. E., Morgan, G. S., & Wisneski, D. C. (2021). The psychology of moral conviction. *Annual Review of Psychology*, 72, 347–366. <https://doi.org/10.1146/annurev-psych-063020-030612>
- Slooman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological bulletin*, 119(1), 3.
- Sommer, J., Musolino, J., & Hemmer, P. (2022). Introduction: Toward a Cognitive Science of Belief. In J. Musolino, J. Sommer, & P. Hemmer (Eds.), *The Cognitive Science of Belief: A Multidisciplinary Approach* (pp. 1–28). chapter, Cambridge: Cambridge University Press.

- Spaulding, S. (2018). Mindreading beyond belief: A more comprehensive conception of how we understand others. *Philosophy Compass*, 13(11), e12526. <https://doi.org/10.1111/phc3.12526>
- Sperber, D. (1997). Intuitive and reflective beliefs. *Mind & Language*, 12(1), 67–83. <https://doi.org/10.1111/1468-0017.00036>
- Stanovich, K. E., & West, R. F. (2000). Individual differences in reasoning: Implications for the rationality debate? *Behavioral and Brain Sciences*, 23(5), 645–665. <https://doi.org/10.1017/S0140525X00003435>
- Shtulman, A., & Lombrozo, T. (2016). Bundles of Contradiction. In D. Barner & A. S. Baron (Eds.), *Core Knowledge and Conceptual Change* (pp. 53–72). Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780190467630.003.0004>
- Tong, Y., Wang, F., & Danovitch, J. (2020). The role of epistemic and social characteristics in children's selective trust: Three meta-analyses. *Developmental Science*, 23(2), e12895. <https://doi.org/10.1111/desc.12895>
- Turri, J., Rose, D., & Buckwalter, W. (2018). Choosing and refusing: Doxastic voluntarism and folk psychology. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 175(10), 2507–2537.
- Tversky, A., & Kahneman, D. (1981). The Framing of Decisions and the Psychology of Choice. *Science*, 211(4481), 453–458. <https://doi.org/10.1126/science.7455683>
- Van Leeuwen, N. (2014). Religious credence is not factual belief. *Cognition*, 133(3), 698–715. <https://doi.org/10.1016/j.cognition.2014.08.015>
- Van Leeuwen, N. (2022). Two Concepts of Belief Strength: Epistemic Confidence and Identity Centrality. *Frontiers in Psychology*, 13. <https://doi.org/10.3389/fpsyg.2022.939949>
- Van Leeuwen, N., & Lombrozo, T. (2023). The Puzzle of Belief. *Cognitive Science*, 47(2), e13245. <https://doi.org/10.1111/cogs.13245>

- Van Leeuwen, N., Weisman, K., & Luhrmann, T. M. (2021). To Believe Is Not to Think: A Cross-Cultural Finding. *Open Mind*, 5, 91–99. [https://doi.org/10.1162/opmi\\_a\\_00044](https://doi.org/10.1162/opmi_a_00044)
- Vesga, A., Van Leeuwen, N., & Lombrozo, T. (2024). Evidence for distinct cognitive attitudes of belief in theory of mind. In Proceedings of the 46th Annual Meeting of the Cognitive Science Society (pp. 2527-2536).
- Watson, A., Power, A., Brown, C., El-Deredy, W., & Jones, A. (2012). Placebo analgesia: Cognitive influences on therapeutic outcome. *Arthritis Research & Therapy*, 14(3), 206. <https://doi.org/10.1186/ar3783>
- Wellman, H. M. (2018). Theory of mind: The state of the art. *European Journal of Developmental Psychology*, 15(6), 728–755. <https://doi.org/10.1080/17405629.2018.1435413>
- Wellman, H. M., & Woolley, J. D. (1990). From simple desires to ordinary beliefs: The early development of everyday psychology. *Cognition*, 35(3), 245–275. [https://doi.org/10.1016/0010-0277\(90\)90024-E](https://doi.org/10.1016/0010-0277(90)90024-E)
- Westra, E. (2018). Character and theory of mind: An integrative approach. *Philosophical Studies*, 175(5), 1217–1241. <https://doi.org/10.1007/s11098-017-0908-3>
- Westra, E. (2023). Symbolic Belief in Social Cognition. *Philosophical Perspectives*, 37(1), 388–408. <https://doi.org/10.1111/phpe.12196>
- Yan, V. X., Thai, K.-P., & Bjork, R. A. (2014). Habits and beliefs that guide self-regulated learning: Do they vary with mindset? *Journal of Applied Research in Memory and Cognition*, 3(3), 140–152. <https://doi.org/10.1037/h0101799>