

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Evidence for distinct cognitive attitudes of belief in theory of mind

Permalink

<https://escholarship.org/uc/item/7hm615q7>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 46(0)

Authors

Vesga, Alejandro

Van Leeuwen, Neil

Lombrozo, Tania

Publication Date

2024

Peer reviewed

Evidence for distinct cognitive attitudes of belief in theory of mind

Alejandro Vesga (av0277@princeton.edu)

Princeton University, University Center for Human Values/Cognitive Science, Princeton, NJ 08540

Neil Van Leeuwen (neilvl@gmail.com)

Georgia State University, Philosophy Department, Atlanta GA, 30303

Tania Lombrozo

Princeton University, Psychology Department, Princeton, NJ 08540

Abstract

Theory of mind is often referred to as “belief-desire” psychology, as these mental states (belief, desire) are accorded a central role. However, extant research has made it clear that defining the notion of belief or characterizing a consistent set of key characteristics is no trivial task. Across two studies (N=283, N=332), we explore the hypothesis that laypeople make more fine-grained distinctions among different kinds of “belief.” Specifically, we find evidence that beliefs with matching contents are judged differently depending on whether those beliefs are seen as playing predominantly epistemic roles (such as tracking evidence with the aim of forming accurate representations) versus non-epistemic roles (such as social signaling). Beliefs with epistemic aims, compared to those with non-epistemic aims, are more likely to be described with the term “thinks” (vs. “believes”), and to be redescribed in probabilistic (vs. binary) terms. These findings call for a refinement of the concepts posited to underlie theory of mind and offer indirect support for the idea that human psychology in fact features more than one kind of belief.

Keywords: belief formation; epistemic reasoning; belief change; belief resistance; theory of mind.

Introduction

The concept of “belief” plays a crucial explanatory role in scientific psychology (Schwitzgebel, 2021; Porot & Mandelbaum, 2021; Musolino et al., 2022; Van Leeuwen & Lombrozo 2023), and also in “intuitive” psychology, often referred to as “belief-desire” psychology or theory of mind (Wellman et al., 1990). As orthodoxy has it, belief is the basic ‘cognitive’ attitude in theory of mind, and desire the basic ‘conative’ attitude. That is, belief is the attitude that represents how the world *is*, and thus figures in our inferences about the past and the future, and in our deliberations about how to bring about our desires. This is why belief is the ‘representational’ component of the belief-desire pair in classic paradigms in theory of mind (Wellman & Wooley, 1990; Goldman 2006).

Yet consensus on the defining features of belief as a mental state has been elusive. Research in psychology suggests that beliefs play multiple—and sometimes seemingly conflicting—psychological roles. On the one hand, beliefs can play epistemic roles, such as tracking evidence and accurately representing the world. On the other hand, beliefs can also play non-epistemic roles, such as emotional regulation and social signaling (cf. Kunda, 1990; Bortolotti, 2010; Kahn & Stanovich, 2015). This variation in roles has

led some scholars to propose that there are different *kinds* of beliefs (Davoodi & Lombrozo 2022; Metz et al., 2023; Van Leeuwen 2022, 2023; Van Leeuwen & Lombrozo 2023). If so, belief as a single category does not capture the richness of our doxastic states.

Crucially, the difference at stake is not between the contents of different beliefs—say, between the belief that God exists and the belief that Jupiter has 12 moons. Rather, the suggestion is that there is a difference in the “cognitive attitudes” involved—i.e., in the way belief contents are processed. To illustrate this distinction between contents and attitudes, consider the following vignettes used in Heiphetz et al. (2021) and Van Leeuwen et al. (2021), which we develop and extend in Study 1:

Religious context: John goes to church every Sunday, and he reads the Bible every day. His preacher says that God punished the Egyptians for not freeing the Israelites from slavery. According to the preacher, God sent ten plagues to the Egyptians, including the death of every firstborn son.

Matter-of-Fact context: John goes to the library every week, and he reads history every day. One history expert says that sickness in Egypt was very common when the Israelites were there. According to the history expert, many plagues happened at that time, including one deadly sickness that disproportionately affected boys.

Participants read one of these two passages and were then asked to fill in the blank in the following sentence, choosing between “thinks” and “believes”: “John ____ that many Egyptian boys died in a plague.” In this judgment, ascribed belief content was held fixed across vignettes, but participants were invited to appropriately describe John’s *attitude* towards that content by choosing “thinks” or “believes.” Participants were significantly more likely to select “believes” over “thinks” in the religious context than in the matter-of-fact context, suggesting that they tended to view the attitudes as different in each case.

In this paper, we ask whether the theoretical distinction between beliefs that play predominantly epistemic vs. non-epistemic roles has a counterpart in intuitive theory of mind. That is, do laypeople systematically differentiate between different types of beliefs in how they attribute and evaluate

those beliefs? And is this differentiation driven by the belief's epistemic vs. non-epistemic features?

The findings from Heiphetz et al. (2021) and Van Leeuwen et al. (2021) offer initial support for this idea. In this paper, we go beyond this prior work in two ways. First, if people operate with two different notions of belief in their intuitive theory of mind, this should manifest in judgments beyond the differential use of “thinks” versus “believes.” In both Study 1 and Study 2, we test three additional dimensions: whether a belief is framed in binary or probabilistic terms, whether it is perceived as voluntary, and whether it is construed in more objective or subjective terms. In each case, we can generate a prediction about which feature would follow from an epistemic aim for belief.

With regard to *Binary vs. Probabilistic* framing, we start with the idea that any belief can be expressed or described as a binary commitment (p is true or p is false) or as a subjective probability (there is $x\%$ probability that p is true). Intuitively, if a cognitive attitude has the function of tracking and reflecting the available evidence, it should seem appropriate to translate it to a more fine-grained probabilistic judgment. For example, if someone believes that a politician is guilty on the basis of the balance of evidence, it seems natural to describe that person as believing there is some (high) probability, perhaps 90%, that the politician is guilty. However, a probabilistic restatement seems less appropriate if a belief's function is to signal religious or moral convictions (compare “I believe that goodness is rewarded in the afterlife” with “I believe there is a 99% chance that goodness is rewarded in the afterlife”). Our first prediction is thus that beliefs that primarily serve epistemic aims (such as tracking truth) should be more natural to describe in probabilistic terms than beliefs that primarily serve non-epistemic aims (such as social signaling).

The second dimension that we consider is the *Perceived Voluntariness* of a belief. Recent research has shown that people's judgments of whether a given belief is voluntary decrease when they attend to evidential considerations in favor of that belief (Cusimano & Goodwin, 2020). Moreover, evidence is seen as a constraint on the beliefs an individual can maintain or generate through reasoning (Cusimano et al., 2021, 2024). We therefore predict that beliefs based on relatively diagnostic evidence should be regarded as less voluntary than those that are not constrained by such evidence, such that beliefs with epistemic aims should be regarded as less voluntary than those with non-epistemic aims.

Finally, we consider *Objectivity vs. Subjectivity*. Research in moral psychology has shown that claims with cultural and moral import are often judged as less truth-apt than scientific or factual claims (Goodwin & Darley, 2008, 2012; Wright, 2018; Liquin et al., 2020; Heiphetz et al., 2014). If people differentiate between epistemic and non-epistemic features of others' beliefs, those beliefs that are formed with non-epistemic aims may be judged less objective than those formed with epistemic aims.

In Study 1 we test these dimensions using the religious and matter-of-fact contexts from Van Leeuwen et al. (2021). We do so on the assumption that beliefs formed in the religious contexts play importantly non-epistemic roles (including moral, social, and emotional roles), whereas those in matter-of-fact contexts have predominantly epistemic aims.

Further, if people differentiate religious and matter-of-fact contexts because these contexts differentially indicate epistemic and non-epistemic aims, as we assume in Study 1, then we should be able to elicit similar variation by manipulating these aims directly. We do this in Study 2 by testing our main DVs in a new task contrasting epistemic contexts (in which a belief is formed by a person who aims to accurately capture the truth) and contexts in which a person has non-epistemic aims, such as preserving a friendship, following her moral commitments, or cohering with her self-image.

Together, these studies begin to shed light on the fine-grained structure of intuitive evaluations of belief states. Further, they make the case for a revision in the assumption of unity in theory of mind construals of belief.

Study 1

In Study 1 we take first steps towards testing whether folk theory of mind distinguishes between cognitive attitudes based on the epistemic versus non-epistemic context for belief. To do so, we follow Heiphetz et al. (2021) and Van Leeuwen et al. (2021)'s method showing that people are more likely to use “*S believes that p*,” as opposed to “*S thinks that p*,” to describe beliefs formed in religious contexts. Our aim in this study is to complement these findings by testing for other relevant contrasts beyond the distinction between “thinking” and “believing” a proposition, namely *Binary vs. Probabilistic*, *Perceived Voluntariness*, and *Objectivity vs. Subjectivity* (as described above). The distinction between religious and matter-of-fact contexts of belief formation should be a good testing ground to evaluate whether ordinary judgments are sensitive to other epistemic and non-epistemic features of belief. As in some studies reported in Heiphetz et al. (2021) and Van Leeuwen et al. (2021), we match the contents of ascribed beliefs across paired vignettes, thus ensuring that differences across conditions are driven by participants' views of cognitive attitudes, and not by belief contents.

One concern is that beliefs could vary across the religious and matter-of-fact contexts not because of the cognitive attitude involved per se, but because the attitude is held more or less strongly. For example, participants might infer that characters in matter-of-fact contexts simply hold their beliefs *more strongly* than those in religious contexts (or vice versa), and this could potentially drive differences in attributions of “thinks” vs. “believes” and our other measures. For this reason, we adopt a rather stringent criterion: across measures we control for judgments of the certainty of the character's belief in the vignette. If participants' judgments along our main DVs indeed distinguish between different cognitive attitudes (where this difference is not itself highly correlated

with certainty), then these effects should survive after controlling for certainty.

Participants

Data were collected from 300 adults (150 male, 150 female) based in the U.S. and whose first language is English. Participants were recruited through Prolific and compensated \$0.50 for a 3-minute study. Participants who failed one or more of two basic attention checks were excluded, resulting in a final sample of 283 participants.

Procedure

The Study employed a fully between-subjects design. Participants were randomly assigned to one of ten vignettes based on the stimuli from Van Leeuwen et al. (2021). Each vignette described a character's belief that arose in either a religious context or a matter-of-fact context (Context: religious, matter-of fact). To match the content of beliefs across contexts, the ten vignettes were paired such that each pair involved a belief with the same ascribed belief content, as in the example from the introduction involving the belief that "many Egyptian boys died in a plague."

After reading their assigned vignette, participants were then asked to make several judgments pertaining to the character's belief (e.g., that "Many Egyptian boys died in a plague"). We had five primary DVs, described below.

Thinks/Believes Participants were asked: "What is a more natural way of completing the blank in the following sentence? 'John _____ that [claim].'" They selected between "thinks" and "believes." This DV was modified from Van Leeuwen et al. (2021).

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

Voluntariness Participants were asked to indicate the extent to which they agreed with two statements about voluntariness, which we will call *choose-not-believe* ("If John wanted to, he could choose to not believe that [claim]") and *choose-to-believe* ("John deliberately chose to believe that [claim]"). They responded on a 7-point scale from -3 ("Completely disagree") to 3 ("Completely agree"). These items were modified from Cusimano and Goodwin (2020).

Objective/Subjective Participants read: "Suppose that John meets Tim, who believes that it is not the case that [claim]. How do you think John is more likely to respond?" They had two response options: "By thinking that at least one of them (John or Tim) must be mistaken" (suggesting they thought John construed the issue as objective) or "By thinking it is possible that neither of them is mistaken" (suggesting they

thought John construed the issue as subjective). This item was modified from Heiphetz et al. (2013).

For control and exploratory purposes, participants answered further questions about the character's certainty that the claim is true ("How certain do you think [character] is that [claim]?" 1-7), the quality of the character's reasons ("To what extent do you think that [character] has good reasons to believe [claim]?" 1-7), and the importance of the belief to the character's identity ("To what extent do you think it is important to [character]'s identity to believe that [claim]?" 1-7).

Results

To analyze our binary DVs, we conducted logistic binomial regressions with context (religious, matter-of-fact) and vignette pair as predictors; for our voluntariness and Objectivity DVs, we performed equivalent linear regressions.

Participants were significantly more likely to indicate that the character "believes that [claim]", vs. "thinks that [claim]", in the religious context compared to the matter-of-fact context ($\beta = 2.182$, $SE = 0.429$, $p < .001$); see Figure 1, left panel. These results replicate those found by Heiphetz et al. (2021), and Van Leeuwen. (2021). Likewise, participants were significantly more likely to select a binary construal of the belief ("John simply believes that [claim]") over the probabilistic framing ("John believes that there is an x% chance that [claim]") in the religious context than the matter-of-fact context ($\beta = 1.151$, $SE = 0.290$, $p < .001$); see Figure 1, right panel. These differences remained significant when controlling for participants' judgments about the character's certainty (Think/Believe: $\beta = 2.157$, $SE = 0.431$, $p < .001$; Binary/Probabilistic: $\beta = 1.125$, $SE = 0.292$, $p < .001$), the character's reasons (Think/Believe: $\beta = 2.232$, $SE = 0.457$, $p < .001$; Binary/Probabilistic: $\beta = 1.011$, $SE = 0.308$, $p < .001$), and the character's identity (Think/Believe: $\beta = 2.045$, $SE = 0.452$, $p < .001$; Binary/Probabilistic: $\beta = 1.267$, $SE = 0.320$, $p < .001$).

The remaining DVs did not reveal significant effects of context (choose-not-believe: $\beta = 0.207$, $SE = 0.197$, $p = 0.295$; choose-to-believe: $\beta = 0.290$, $SE = 0.207$, $p = 0.161$; Objective/Subjective: $\beta = 0.0001$, $SE = 0.044$, $p = 0.996$). These effects also failed to reach significance when controlling for perceived character certainty (choose-not-believe: $\beta = 0.227$, $SE = 0.198$, $p = 0.252$; choose-to-believe: $\beta = 0.255$, $SE = 0.206$, $p = 0.217$; Objective/Subjective: $\beta = -0.017$, $SE = 0.043$, $p = 0.685$); character identity (choose-not-believe: $\beta = 0.242$, $SE = 0.218$, $p = 0.266$; choose-to-believe: $\beta = 0.041$, $SE = 0.225$, $p = 0.855$; Objective/Subjective: $\beta = -0.033$, $SE = 0.049$, $p = 0.491$); or the perceived quality of reasons (choose-not-believe: $\beta = 0.028$, $SE = 0.211$, $p = 0.892$; choose-to-believe: $\beta = 0.207$, $SE = 0.223$, $p = 0.252$; Objective/Subjective: $\beta = -0.018$, $SE = 0.048$, $p = 0.700$).




Figure 1: Mean proportions of participants in Study 1 choosing ‘believes’ for the Think/Believe DV and the binary framing for the Binary/Probabilistic DV. Error bars represent standard errors.

Discussion

Study 1 indicates that participants distinguish beliefs formed in religious and matter-of-fact contexts along two key dimensions. First, replicating Heiphetz et al. (2021) and Van Leeuwen (2021), we found that participants were more likely to select “believes” (vs. “thinks”) in religious compared to matter-of-fact contexts. Going beyond prior work, we also identified a novel dimension along which beliefs were differentiated: participants were significantly more likely to construe beliefs as binary (vs. probabilistic) in religious compared to matter-of-fact contexts. Crucially, these effects survived controlling for participants’ judgments concerning how certain the vignette character is. This indicates that the meta-cognitive distinctions tracked by judgments on these two dimensions are not simply a function of judgments of the perceived certainty with which a belief is held.

The absence of significant effects of context on perceptions of voluntariness and objectivity is also telling. In line with recent research on the perceived voluntariness of doxastic mental states (Cusimano & Goodwin, 2019, 2020; Turri et al. 2018), participants in both the religious and matter-of-fact conditions gave ratings above the scale midpoint on our Voluntariness measures, suggesting that they attributed control over beliefs in both contexts. And although prior work has found that religious beliefs tend to be regarded as less objective than matter-of-fact beliefs (Heiphetz et al., 2013), it is plausible that these effects were driven more by the content of the claims in question than the attitude with which they were held.

Study 2

In Study 1 we assessed whether paradigmatic epistemic and non-epistemic contexts of belief formation (namely, religious and matter-of-fact) yielded differences along four different dimensions of belief. Notably, our Think/Believe and Binary/Probabilistic measures showed significant effects. However, the question remains as to what it is about these two contexts that gave rise to different judgments. If our hypothesis is correct, participants were responding to the

epistemic and non-epistemic roles of belief. Study 2 tests this hypothesis directly.

In Study 2, we use the paradigm from Study 1 to contrast beliefs that are held with different aims (epistemic vs. non-epistemic). In the Epistemic condition, the character was described as holding their belief for epistemic reasons, would have changed their belief in light of contrary evidence, and had the goal of believing whatever was true. In the Non-Epistemic condition, the character was described as holding the same belief for non-epistemic reasons (religious, moral, intuitive, or affiliative), would not have changed their belief in the face of contrary evidence, and had a goal aligned with non-epistemic considerations (e.g., being as loyal as possible). If folk judgments about belief in Study 1 indeed tracked epistemic and non-epistemic features of belief, we should observe analogous effects when contrasting these epistemic and non-epistemic features more directly.

Participants

Data were collected from 350 adults (157 male, 175 female) based within the U.S., and whose first language is English. Participants were recruited through Prolific and compensated \$0.65 for a 4-minute study. Participants who failed one or more of two basic attention checks were excluded, resulting in a final sample of 332 participants.

Procedure

As in Study 1, Study 2 employed a fully between-subjects design. Participants were randomly assigned to one of eight vignettes in which a character holds a belief. As in Study 1, these were paired to match belief content. Instead of varying belief *context*, however, we manipulated the belief’s *aim*. Below are two sample vignettes:

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.

In both cases, participants were then asked to make judgments concerning the belief that "John did not steal money from the organization he was working for."

We used the same DVs used in Study 1. The only change was in the Think/Believe DV, where we asked "What is a more natural way of completing the blank in the sentence above?" ('thinks'/'believes'). Once participants answered this question, they were shown the same vignette with the blank replaced by their choice. Participants also answered questions about the character's certainty, the quality of the character's reasons for belief, and the importance of the belief to the character's identity.

Results

To analyze our binary DVs, we conducted logistic binomial regressions with context (religious, matter-of-fact) and vignette pair as predictors; for our Voluntariness and Objectivity DVs, we performed equivalent linear regressions. This revealed that participants in the Non-epistemic condition were more likely than those on the Epistemic condition to select 'believes' over 'thinks' ($\beta = 1.082$, $SE = 0.336$, $p < .01$); see Figure 2, left panel. Likewise, participants in the Non-epistemic condition were more likely than those in the Epistemic condition to choose the binary framing over the probabilistic framing ($\beta = 0.594$, $SE = 0.244$, $p < .05$); see Figure 2, right panel. As in Study 1, these differences remained significant when controlling for character certainty (Think/Believe: $\beta = 1.080$, $SE = 0.336$, $p < .01$; Binary/Probabilistic: $\beta = 0.604$, $SE = 0.248$, $p < .05$), and the quality of the character's reasons (Think/Believe: $\beta = 0.989$, $SE = 0.347$, $p < .01$; Binary/Probabilistic: $\beta = 1.011$, $SE = 0.308$, $p < .001$). However, when controlling for the character's identity, only the effects of Think/Believe were significant (Think/Believe: $\beta = 1.229$, $SE = 0.359$, $p < .001$; Binary/Probabilistic: $\beta = 0.436$, $SE = 0.260$, $p = 0.093$).




Figure 2: Mean proportions of participants in Study 2 choosing 'believes' for the Thinks/Believes DV and the

binary framing for the Binary/Probabilistic DV. Error bars represent standard errors.

As in Study 1, a logistic binomial regression revealed no significant effect of condition on objectivity ratings (Objective/Subjective: $\beta = -0.05792$, $SE = 0.31426$, $p = 0.854$). However, there was a significant effect of condition on our voluntariness DVs. Participants were more likely to judge that the vignette character could choose *not* to believe the claim in question in the Epistemic condition (choose-not-believe: $\beta = -0.334$, $SE = 0.165$, $p < 0.05$). In contrast participants were more likely to judge that the vignette character chose to believe the claim in question in the Non-Epistemic condition (choose-to-believe: $\beta = 0.615$, $SE = 0.172$, $p < 0.001$). These effects also survived controlling for character certainty (choose-not-believe: $\beta = -0.342$, $SE = 0.165$, $p < 0.05$; choose-to-believe: $\beta = 0.605$, $SE = 0.172$, $p < 0.001$), and the perceived importance of the belief for the character's identity (choose-not-believe: $\beta = -0.487$, $SE = 0.176$, $p < 0.01$; choose-to-believe: $\beta = 0.399$, $SE = 0.182$, $p < 0.05$). However only choose-to-believe survived controlling for the perceived quality of the character's reasons (choose-not-believe: $\beta = -0.317$, $SE = 0.173$, $p = 0.067$; choose-to-believe: $\beta = 0.675$, $SE = 0.180$, $p < 0.001$).




Figure 3: Mean ratings in Study 2 for *choose-not-believe* and *choose-to-believe*. Error bars indicate 95% confidence intervals. Points represent individual participant responses.

Discussion

Study 2 indicates that the distinctions tracked by the Think/Believe and Binary/Probabilistic variables go beyond the differences in general context probed in Study 1. When we manipulated a belief's epistemic vs. non-epistemic aims more directly, we found similar effects.

Intriguingly, although both studies found that beliefs across both conditions were generally considered voluntary (i.e., ratings were above the scale midpoint), there was a significant effect of condition on voluntariness in Study 2: Participants more strongly agreed that vignette characters deliberately chose to hold their belief in the Non-Epistemic condition than in the Epistemic condition, consistent with our predictions. However, participants also more strongly agreed that the vignette characters could choose *not* to hold their

belief in the Epistemic condition, compared to the Non-epistemic condition. Given this pattern, more research is needed before it is appropriate to interpret results as evidence for a general difference in voluntariness, and we refrain from drawing substantive conclusions.

General Discussion

The findings from Studies 1 and 2 offer critical insights into the fine-grained distinctions in laypeople's judgments about belief. Rather than having a unified representation of belief, our results suggest that people make subtle distinctions about the kinds of doxastic attitudes that others have with respect to a proposition. Holding the content of the beliefs fixed, participants made systematic distinctions between beliefs that were formed in epistemic contexts (or with epistemic aims) and beliefs that were formed in non-epistemic contexts (or with non-epistemic aims). In particular, our results strongly suggest that beliefs formed in epistemic contexts are more likely than beliefs formed in non-epistemic contexts to prompt descriptions using the word 'thinks' and probabilistic framing—as opposed to the word 'believes' and binary framing. Notably, these distinctions are not a function of the inferred certainty of the character's belief, which might otherwise be a confound.

Study 1 replicated and built on the results from Heiphetz et al. (2021) and Van Leeuwen et al. (2021). The distinction between 'thinks' and 'believes' is not the only doxastic distinction elicited by religious and matter-of-fact contexts. Moreover, Study 2 indicates that these judgements are not bound to the contrast between religious and matter-of-fact contexts. Rather, the differences found across our studies seem to correspond to more general clusters of epistemic and non-epistemic roles. More precisely, the results in Study 2 indicate that the judgements in question are prompted by differences in the perceived epistemic and non-epistemic *aims* that a belief has in a person's psychology. If a belief is perceived to aim for truth and accuracy, it is more likely to be redescribed in probabilistic and 'think' terms, compared to when it is not. In contrast, when a belief is perceived to play non-epistemic (moral, religious, affiliative, or identity-preserving) roles, it is more likely to be redescribed in binary and 'believe' terms. That we find these differential patterns in linguistic use beyond religious contexts gives weight to the idea that they track different kinds of folk evaluations of beliefs in themselves, beyond the contextual cues of a specific thematic domain.

An aspect of our research that leaves open questions concerns the perceived voluntariness of belief. Our results corroborate previous research indicating that people generally perceive others' beliefs as voluntary (Cusimano & Goodwin, 2020; Cusimano et al., 2021). However, only in Study 2 was there a difference between conditions, and the observed differences only partially matched our predictions. Additional research will be necessary to expand on the implications of these results, and to understand the role of motivations in belief formation more broadly. For instance, under a classic construal, motivated beliefs are somehow

inappropriate or defective (cf. Cusimano & Lombrozo, 2021b). However, if there is a distinctive kind of belief that takes non-epistemic motivations as input, it would be of interest whether such beliefs are deemed inappropriately formed or defective, and whether these judgments vary depending on whether motivation was involved in forming a belief or in preventing a belief from being formed. Alternatively, it could be that these judgments are just reserved for the epistemic variety of belief. Further research should also investigate whether people distinguish between epistemic and non-epistemic features of their own beliefs. Another largely open question concerns whether people make the reverse inference to our present results. That is, are people more likely to assign epistemic features to beliefs expressed as 'I think that p' and 'There is x% probability that p' than to beliefs expressed as 'I believe that p' and simply 'p is true'? Ongoing research suggests this is indeed the case, further supporting the conclusion that laypeople track the posited distinction between kinds of beliefs.

Finally, we take our results to be compatible with what has been called the *Varieties of Belief* thesis (Van Leeuwen & Lombrozo, 2023), which proposes that our best psychological and philosophical theories ought to distinguish between kinds of belief that play different roles in our psychology. If our folk theory of mind incorporates systematic distinctions between beliefs with epistemic and non-epistemic roles, perhaps this is because people are attuned to genuinely different natural kinds, and our best psychological and philosophical theories should incorporate such a distinction, too. Nevertheless, more research is necessary to decisively favor this thesis. In particular, it should be distinguished from theories that posit a single kind of belief whose epistemic features are often "masked" by non-doxastic psychological processes (cf. Flores, forthcoming; Helton, 2021).

References

- Bortolotti, L. (2010). *Delusions and Other Irrational Beliefs*. Oxford University Press.
- 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.
- Cusimano, C., & Lombrozo, T. (2021a). Morality justifies motivated reasoning in the folk ethics of belief. *Cognition*, 209, 104513.
- Cusimano, C., & Lombrozo, T. (2021b). Reconciling scientific and commonsense values to improve reasoning. *Trends in Cognitive Sciences*, 25(11), 937-949.
- Cusimano, C., Zorrilla, N., Danks, D., & Lombrozo, T. (2024). Psychological freedom, rationality, and the naïve theory of reasoning. *Journal of Experimental Psychology: General*, 153(3), 837.
- Davoodi, T., & Lombrozo, T. (2022). Explaining the existential: Scientific and religious explanations play different functional roles. *Journal of Experimental Psychology: General*, 151(5), 1199.

- Flores, C. (forthcoming). Resistant Beliefs, Responsive Believers. *Journal of Philosophy*.
- Goldman, A. I. (2006). *Simulating Minds: The Philosophy, Psychology, and Neuroscience of Mindreading*. Oxford University Press.
- Goodwin, G. P., & Darley, J. M. (2008). The psychology of meta-ethics: Exploring objectivism. *Cognition*, 106(3), 1339-1366.
- Goodwin, G. P., & Darley, J. M. (2012). Why are some moral beliefs perceived to be more objective than others?. *Journal of Experimental Social Psychology*, 48(1), 250-256.
- 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.
- Heiphetz, L., Spelke, E. S., Harris, P. L., & Banaji, M. R. (2013). The development of reasoning about beliefs: Fact, preference, and ideology. *Journal of experimental social psychology*, 49(3), 559-565.
- Heiphetz, L., Spelke, E. S., Harris, P. L., & Banaji, M. R. (2014). What do different beliefs tell us? An examination of factual, opinion-based, and religious beliefs. *Cognitive development*, 30, 15-29.
- Helton, G. (2018). If You Can't Change What You Believe, You Don't Believe It. *Noûs*, 54(3), 501–526.
- Kahan, D. M., & Stanovich, K. (2016). Rationality and belief in human evolution. Annenberg Public Policy Center Working Paper, (5).
- Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, 108(3), 480–498.
- Liquin, E. G., Metz, S. E., & Lombrozo, T. (2020). Science demands explanation, religion tolerates mystery. *Cognition*, 204, 104398.
- Metz, S. E., Liquin, E. G., & Lombrozo, T. (2023). Distinct Profiles for Beliefs About Religion Versus Science. *Cognitive science*, 47(11), e13370.
- Musolino, J., Sommer, J., & Hemmer, P. (Eds.). (2022). *The Cognitive Science of Belief: A Multidisciplinary Approach*. Cambridge University Press.
- Porot, N., & Mandelbaum, E. (2021). The science of belief: A progress report. *Wiley Interdisciplinary Reviews: Cognitive Science*, 12(2), e1539.
- Schwitzgebel, E. (2021). In E. N. Zalta (Ed.), "Belief," *The Stanford Encyclopedia of Philosophy* (Winter 2021 Edition) URL= <<https://plato.stanford.edu/archives/win2021/entries/belief/>>.
- Turri, J., Rose, D., & Buckwalter, W. (2018). Choosing and refusing: doxastic voluntarism and folk psychology. *Philosophical Studies*, 175(10), 2507-2537.
- Van Leeuwen, N. (2022). Two Concepts of Belief Strength: Epistemic Confidence and Identity Centrality. *Frontiers in Psychology*, 13, Article 939949.
- Van Leeuwen, N. (2023). *Religion as Make-Believe: A Theory of Belief, Imagination, and Group Identity*. Harvard University Press.
- Van Leeuwen, N. & Lombrozo, T. (2023). The puzzle of belief. *Cognitive Science*, 47(2), e13245.
- Van Leeuwen, N., Weisman, K., & Luhrmann, T. M. (2021). To believe is not to think: A cross-cultural finding. *Open Mind*, 5, 91-99.
- Wellman, H. M., & Woolley, J. D. (1990). From simple desires to ordinary beliefs: The early development of everyday psychology. *Cognition*, 35(3), 245–275.
- Wright, J. C. (2018). The fact and function of meta-ethical pluralism: Exploring the evidence. *Oxford studies in experimental philosophy*, 2, 119-150.