Dimensions of Identity-Representing Belief

Martin Meyer (m.meyer@yale.edu)
Yale University, Department of Psychology, New Haven, CT 06510
Joshua Knobe (joshua.knobe@yale.edu)
Yale University, Program in Cognitive Science and Department of Philosophy, New Haven, CT 06510

Abstract

Recent work has proposed that there are symbolic beliefs. These beliefs do not serve primarily to track the facts in the world but rather to express the believer's own identity. On this view, several disparate features of belief – from whether a belief is important to identity to whether it is sensitive to evidence – would be related to an underlying "symbolicness" dimension. We converted the features potentially related to symbolicness into items and asked people to rate their own beliefs on them. Study 1 found that beliefs which were high on one feature (importance) were rated higher on all the items, except for insensitivity to evidence. Study 2 found that ratings of any beliefs on almost all the items loaded onto a single, symbolicness factor, except again for evidence insensitivity. Study 3 asked participants to rate their own beliefs on all the items in Study 2 and several additional items designed to measure whether a belief was subjective. We recovered the symbolicness factor, but found it was largely orthogonal to the subjectivity and evidence insensitivity items. These findings suggest that most of the features we tested relate to an underlying symbolicness factor, which corresponds to whether a belief represents identity. But, surprisingly, this factor was not related to items that get at whether a belief represents facts about the world. It would seem that the degree to which a belief aims to express the believer's identity and the degree to which a belief aims at accurately tracking facts in the world are two orthogonal dimensions, which can vary independently.

Keywords: cognitive attitudes, identity representation, dimensions of belief, symbolic belief

Introduction

Compare the beliefs "The Lord works in mysterious ways" and "There are three pieces of bread in the basket". Intuitively, these beliefs have wildly different psychological properties. Recent philosophical work suggests that they may even be *different kinds* of beliefs (Westra, 2023; Van Leeuwen, 2022). The proposal is, specifically, that there are fundamentally two kinds of belief. One sort are symbolic beliefs, which aim to represent the identity of the person who believes them. "The Lord works in mysterious ways" occupies this role. The other sort are non-symbolic beliefs, which aim to represent facts. For example, the belief "There are three pieces of bread in the basket" expresses something about the actual state of the world.

A critical example of symbolic beliefs are beliefs which seem to be about the world, but do not update to match the available evidence. For instance, the belief that "Donald Trump's inauguration was attended by more people than Obama's" is maintained in the face of blatant photographic evidence to the contrary (Ross & Levy, 2023). Much research has been done to identify the non-evidentiary reasons that justify such beliefs, e.g., identity-related concerns (Kahan, 2015) and ethical considerations (Metz et al. 2023). The theory of symbolic beliefs goes further, and suggests that beliefs like the Trump example belong to a different kind: They differ from evidence-sensitive beliefs not just in their justificatory conditions, but also on a broader set of other dimensions (e.g., if they are probabilistically represented). Importantly, unlike work on reflective beliefs (c.f., Mercier, 2021; Sperber, 1997), the point is not that they *can* vary on all these dimensions, but that they do.

Suppose then that there are two fundamentally different sorts of beliefs (symbolic and non-symbolic) which differ on several observable psychological dimensions. If this were true, then if we took these dimensions that differ between symbolic and non-symbolic beliefs, they should all load on a single, "symbolicness" factor. In this work, we explore whether there are symbolic beliefs by testing for a symbolicness factor.

We gathered several dimensions which prior work suggested would load on a symbolicness factor. The most important of these were the dimensions related to *identity centrality*. Given that the definition of symbolic beliefs is that they aim to represent identity, it seemed crucial to include them. They captured whether the belief is integral to who the person who believes it really is. For example, you might expect someone's moral beliefs to be core to who they are (Strohminger & Nichols, 2014). On the other hand, if someone told you they believed "There are 26 letters in the English alphabet", you probably would not have learned much about who they *really* are.

Symbolic beliefs are not only positively defined as being concerned with identity, but also negatively defined as *not* being concerned with facts. Therefore, it was essential that we captured this property with *evidence insensitivity*. Highly evidence sensitive beliefs readily update in light of countervailing evidence. Bob might, for example, believe Joey's car is blue; but he would immediately change that belief if he discovered Joey's car was green. Some beliefs, though, do not readily adjust to novel evidence. For example, it is hard to imagine what sort of evidence would change a mother's belief that her child is the most beautiful person in the world.

Another construct related to not capturing facts about the world is *subjectivity*. Indeed, it seems to be the very definition of subjectivity that it is not about facts in the world. Although there are several ways to get at the

In D. Barner, N.R. Bramley, A. Ruggeri and C.M. Walker (Eds.), *Proceedings of the 47th Annual Conference of the Cognitive Science Society*

dimension (Goodwin & Darley, 2008; Sarkissian et al., 2011; for a review, see Pölzler & Wright, 2019), consider just the difference between "Mercury is closer to the sun than Saturn" and "Elliott Smith's music is better than Billie Eilish's music." When it comes to the belief about the planets, it seems that there is a fact of the matter. Either the belief is true or it isn't. But when it comes to the belief about music, one might think that there is no fact of the matter. The whole issue might be thought to be a purely subjective one.

Symbolicness might even affect whether we describe certain beliefs using the word *believe* instead of *think*. Recent cross-cultural work suggests the phrase "someone believes" (or its linguistic equivalent) precedes religious beliefs at a much higher rate than "someone thinks". That is, one is more likely to say, "Someone *believes* Muhammad is the true prophet" than "Someone *thinks* Muhammad is the true prophet" (Van Leeuwen, Weisman, Luhrmann, 2021). Another study found that beliefs held for non-evidentiary reasons were more likely to be preceded by "I believe" (Vesga, Van Leeuwen, Lombrozo, 2024). This is especially intriguing given that evidence insensitivity could be a property of symbolic beliefs.

This brings us to the main question: do these dimensions all load positively onto a single symbolic factor? To answer this, in the following studies, we simply asked participants to share one of their beliefs and rate it on those dimensions. While past studies of symbolic beliefs have focused on the attribution of symbolicness to beliefs of others, we had participants rate their own beliefs. This gave us a more diverse set of participant beliefs to analyze.

Study 1 tested whether important beliefs scored higher than unimportant beliefs on the dimensions. Study 2 directly tested the factor structure of scores on the dimensions. Study 3 tested how subjective dimensions figured into factor structure.

Study 1

As a first study, we tested the relationship of many symbolic dimensions with just one – whether the belief was important to the participant. Specifically, we had participants share an important or unimportant belief and had them rate it on items based on the dimensions described in the introduction.

Methods

Preregistration Analyses were preregistered: https://researchbox.org/2552

Participants We recruited 300 adult participants on Prolific who spoke primarily English. 189 were included (54% female, 56% white, and on average 31 years old (SD = 10.7)) according to our preregistered exclusion criteria.

Procedure Participants in one condition were asked to share a belief that was important to them, i.e., "a fact or a value judgment you hold about something (or someone) other than yourself that is important to you. You place great value on holding it." Participants in another condition were asked to share a belief that was not important to them, i.e., "a fact or a value judgment you hold about something (or someone) other than yourself that is NOT important to you. You place little (if any) value on holding it."

They rated this belief on the nine items we thought could be related to symbolicness. Some of the items were about individual identity. These were: identity, difficulty of imagining otherwise, and importance. (Importance was our manipulation, but we asked participants to rate their given belief on it as a manipulation check.) Others were about social identity. These were: importance to reputation and intolerance of intolerance. And then there were a handful of dimensions that were not directly about identity, but nevertheless predicted by theories of symbolicness. Those items were: evidence insensitivity, believe instead of think, and no probabilistic representations. To see the definition and presentation of all these items, readers can follow this link: https://osf.io/bc4au/

Results

We first excluded all participants who failed one of our attention or comprehension checks. Then we used t-tests to compare whether scores on the items were different for important and unimportant beliefs. Important beliefs showed higher ratings on almost every item (p's < 0.0001). The exception was evidence insensitivity ($p = 0.07$).

We were cautious about these results given the high number of exclusions. We discovered that one of our manipulation checks (which required that participants in the important condition rate their belief above 50 on the importance slider and those in the unimportant condition rate their belief below 50 on the importance slider) excluded almost 20% of our participants. Therefore, in an exploratory analysis, we performed the t-tests without these exclusion. All significant effects remained so. However, evidence-insensitivity now also crossed the significance threshold, $t(244.57) = 2.37$, $p = 0.02$.

Discussion

Beliefs that participants considered important had higher ratings than unimportant beliefs on almost every item. This suggests that these items may be related. What's more, it suggests that these items might be related positively. A positive relationship between importance, no-probabilistic-representations, intolerance-of-intolerance, and so on is in line with the proposal by Van Leeuwen (2022) and Westra (2023) that these dimensions might correspond with whether a belief represents identity. In other words, this study provides preliminary evidence for the existence of a symbolicness factor.

Of course, these findings only regard the relationships between importance and the other items. A more thorough test of the existence of a symbolic factor requires examining the relationships between all the dimensions. Importantly, such a test would not only allow us to determine whether

evidence-insensitivity is orthogonal to importance, but also to determine whether it is orthogonal to the other items as well.

Study 2

To directly test the factor structure of our items, we asked participants to offer one of their beliefs and to rate it on the items from Study 1. If an exploratory factor analysis of these ratings revealed significant positive loadings of all items on a single factor, this would provide strong evidence for a single symbolic factor.

If we did uncover a symbolicness factor, we were curious to see how it was distributed. A bimodal distribution would suggest that there are two sorts of beliefs: those that serve an identity-representing function (like "The Lord works in mysterious ways") and those that serve a more fact-representing function (like "There are three pieces of bread in the basket"). A unimodal distribution would indicate that there are not two sorts of belief.

Methods

Preregistration Analyses were preregistered: https://researchbox.org/2552

Participants We recruited 375 adult participants on Prolific who spoke primarily English, and included 295 in the final analysis after performing our preregistered exclusions. These participants were 51% female, 58% white, and on average 33 years old (*SD* = 11.8).

Procedure

Participants shared a belief that was not about themselves and rated this belief on the nine items from Study 1. Along the way, they completed 3 comprehension checks and one attention check.

Results

We first ran Pearson correlations between scores on the nine items. Even after performing our preregistered Holm-Bonferroni correction, there were significant positive correlations between almost every item. The main exception was that evidence-insensitivity was uncorrelated with everything else.

The Kaiser-Meyer-Olkin factor adequacy of our data was strong (overall MSA = 0.82), so we proceeded with factor analysis. We used parallel analysis, with 10,000 bootstrapped iterations and principal axis extraction, and found that we should extract two factors.

We obtained the two-factor solution using principal axis extraction method, varimax rotation, and 10,000 bootstrapped iterations. The RMSEA was 0.064. According to our preregistered criterion, an item loads on a factor if the bootstrapped confidence intervals for its loading do not cross 0. According to this criterion, all items except for evidence-insensitivity loaded on at least one of the factors.

However, a one-factor solution also seemed plausible. The scree plot showed an elbow after the first factor; the eigenvalue for the second factor was well below the Kaiser criterion; and the items that loaded on different factors were intercorrelated. See Figure 1. Additionally, if we use a rotation method that allows for the factors to correlate, we find that the two factors in the two-factor solution are highly intercorrelated ($r(295) = 0.62$ using direct oblimin).

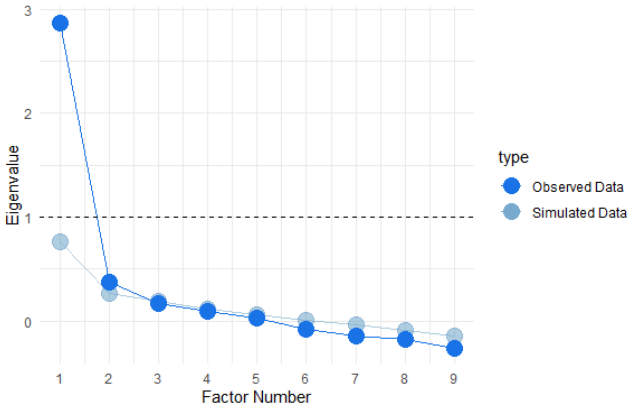

Figure 1: Scree Plot for Study 2. Eigenvalues of factors are in dark blue; factors simulated via parallel analysis are in light blue. The dotted y-intercept indicates the Kaiser Criterion for determining the quantity of factors one should extract. Based on code by Sakaluk & Short (2016).

Thus, in an exploratory analysis, we obtained a one-factor solution, again using the principal axis method and 10,000 bootstrapped iterations. Here, the RMSEA was still adequate: 0.097. Every item loaded positively on this factor except for evidence-insensitivity. See Table 1 for factor loadings. Therefore, a single factor explained the presence of eight of the nine symbolic items. We decided to call this the "symbolic" or "symbolicness" factor.

To test the distribution of the factors we extracted, we took the factor scores using the Thurstone Method and performed a Dip Test. Neither of the factors in the two factor solution showed evidence against unimodality ($D = 0.014$; $p = 0.99$, $D = 0.02$; $p = 0.88$). Nor did the symbolic factor show any evidence against unimodality ($D = 0.01$, $p = 0.99$).

Discussion

We found evidence for a single factor underlying several dimensions of belief – from the belief's relative importance to the believer, to whether the belief is described using "I believe" or "I think." This supports the suggestion (Westra, 2023; Van Leeuwen, 2022) that whether a belief represents identity is a stable, latent property that undergirds these dimensions: symbolicness. Interestingly, the distribution of the factors indicates that "symbolicness" is best understood as a continuum, rather than as a distinction between two discrete kinds. At the very least, people attribute symbolicness to their beliefs in a continuous rather than discrete fashion.

This finding is complicated by the fact that our preregistered parallel analysis suggested a two-factor solution was best. Intriguingly, these factors contained the same items as the single factor solution, and were highly intercorrelated. Given these results and our *a priori* theoretical reasons we have for expecting a single factor, we interpret these findings as pointing to a single symbolic factor with two-facets. However, we were curious to verify if these results – i.e., a plausible one factor solution, and a two factor solution with correlated facets – would replicate in another study.

Most strikingly, we found no evidence of evidence-insensitivity being related to the other items. You might think that a belief involved in representing identity would be less involved in representing facts. Thus, high symbolicness would imply high evidence insensitivity. But this was not so.

The question then arises: Does whether a belief represents identity (symbolicness) have anything to do with whether it represents facts? A way to test this would be to see (a) if evidence insensitivity is part of a larger construct, and (b) what relationship (if any) that larger construct has with symbolicness.

Study 3

Our main goal was to test whether a belief's symbolicness was related to whether it tracked facts about the world. To do this, we not only measured participant beliefs with the items we had used to measure symbolicness, but also several new items meant to capture whether a belief was subjective. Choosing this construct was advantageous for two reasons. Firstly, subjectivity, by definition, tracks whether something is about the world. It therefore seems to be a perfect candidate for assessing the construct of interest. Secondly, the emphasis on subjectivity clears up some of the ambiguity that accompanies self-reports of evidence insensitivity alone. For instance, it is unclear how a participant might respond if their belief is sensitive to evidence, but the available evidence is highly inaccessible (Sommer, 2024). Measures of subjectivity do not seem to have this ambiguous quality.

We also wanted to determine whether the two-facet interpretation of the factor structure in Study 2 held, and test again for any evidence against the unimodality of the factor scores.

Methods

Preregistration Analyses were preregistered: https://researchbox.org/2552

Participants We recruited 375 adult participants on Prolific who spoke primarily English. After our preregistered exclusions, we had 287 participants who were 52% female, 48% white, and on average 32 years old (SD = 10.3).

Procedure Participants were instructed to share a belief. They then rated this belief on the nine items from the first two studies, plus 7 more items meant to capture whether a belief was subjective. Along the way, they completed three comprehension checks and one attention check.

Results

We ran Pearson correlations between all 16 items. Broadly speaking, there were two sets of intercorrelations: a near replication of Study 2's correlations, and a correlations between the new subjectivity items.

Computation of the Kaiser-Meyer-Olkin factor adequacy indicated adequate MSA (0.77). We thus proceeded with parallel analysis (using principal axis factoring method and 10,000 iterations) which suggested three factors. We also looked at the Kaiser Criterion, which suggested two factors. In accordance with our preregistration, we explored both factor solutions.

The two-factor solution (using principal axis method, direct oblimin, and 10,000 bootstrapped iterations) yielded an RMSEA of 0.08. The two factors corresponded to the symbolic and subjective constructs. According to our preregistered loading criterion, every symbolic item loaded on the symbolic factor and so too for the subjective ones (crucially, evidence insensitivity loaded on the subjective factor). See Table 1 for factor loadings and Figure 2 for distribution of participant beliefs according to factor scores.

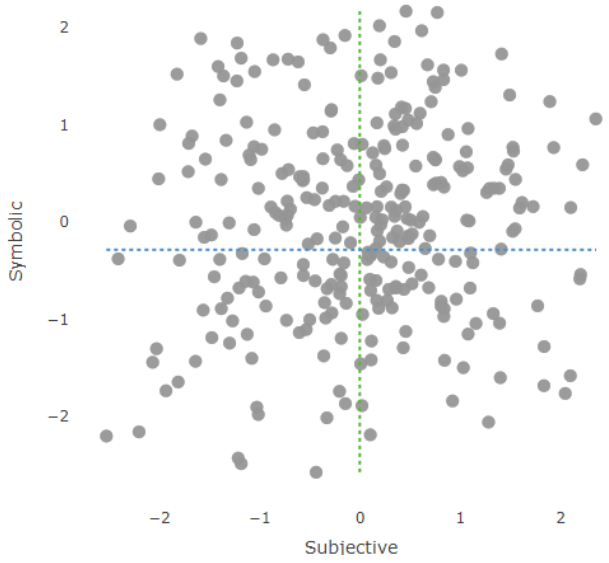

Figure 2: Scatterplot of factor scores from two-factor solution in Study 3. Rotation method was direct oblimin and factor scores were extracted using the tenBerge method. Axes represent approximate factor scores for a hypothetical participant who gave each item a rating of 50. Follow this link to view the participant belief which corresponds with each set of factor scores: https://plotly.com/~m.meyer/4.

Although there were several cross-loadings, there was no relationship between the factors themselves. Specifically, we used the tenBerge method to extract the scores from each factor and found the correlation was insignificant $r(287) = 0.08$, $p = 0.20$.

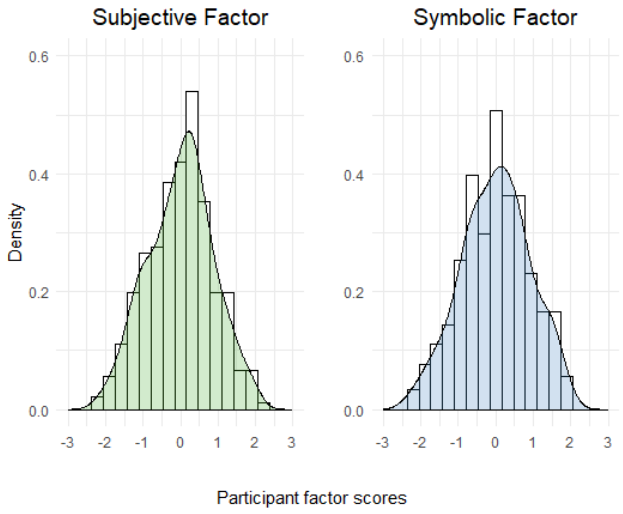

Figure 3: Continuous and binned density plot of the distributions of factor scores (extracted by the Thurstone method) for the two factor solution in Study 3. Dip Tests found no evidence against unimodality.

The three factor solution (again using principal axis method, direct oblimin, and 10,000 bootstrapped iterations) yielded an RMSEA of 0.06. There was a subjective factor, where only the items meant to capture subjectivity loaded (except for no-probabilistic-representations). Then there were two factors on which the symbolicness items loaded. They closely replicated the two-factor solution from Study 2, with the exception of a few subjectivity items that cross-loaded.

By extracting the factor scores via the tenBerge method, we ascertained that the symbolicness factors were significantly intercorrelated $r(287) = 0.48$, $p < 0.0001$. However, none of the other correlations were (p's > 0.11).

As before, we looked at the distribution of scores for each of these factors. Using either the Thurstone or tenBerge methods, we found no evidence against the unimodality of any of the factors (p's > 0.95). See Figure 3 for the distributions of the two factor solution.

Discussion

In this study, we found no evidence that higher symbolicness corresponded to higher subjectivity. We found a single symbolicness factor using the Kaiser Criterion and using parallel analysis two highly intercorrelated factors on which the symbolic items load. Thus, the exact results that led us to think symbolicness is a two-facet structure replicated. But now we see that evidence-insensitivity and seven new items meant to capture subjectivity load onto a separate, subjectivity factor. This subjectivity factor was uncorrelated to the symbolicness factor(s); if anything, its items loaded negatively on the symbolicness factor(s).

As a secondary point, we again found no evidence for the possibility that symbolic and non-symbolic beliefs are two discrete kinds of belief – at least as far as people conceptualize their own beliefs. So too for subjectivity.

Table 1: Factor loadings for Studies 2 and 3

	Study 2	Study 3	
Item	Symbolic Factor	Symbolic Factor	Subjective Factor
Importance	0.75 [0.67, 0.84]*	0.79 [0.71, 0.85]*	0.04 [-0.05, 0.13]
Surrounding emotions	0.65 [0.56, 0.74]*	0.65 [0.56, 0.74]*	0.12 [0.00, 0.23]
Believe instead of think	0.39 [0.27, 0.52]*	0.32 [0.18, 0.45]*	-0.01 [-0.16, 0.15]
Identity	0.71 [0.63, 0.78]*	0.66 [0.58, 0.75]*	0.03 [-0.08, 0.13]
Difficulty of imagining otherwise	0.56 [0.46, 0.65]*	0.57 [0.47, 0.67]*	0.01 [-0.10, 0.12]
No probabilistic representations	0.50 [0.38, 0.61]*	0.28 [0.15, 0.41]*	-0.27 [-0.42, -0.12]*
Importance to reputation	0.56 [0.47, 0.66]*	0.51 [0.39, 0.62]*	-0.08 [-0.22, 0.06]
Intolerance of intolerance	0.58 [0.49, 0.67]*	0.59 [0.48, 0.69]*	-0.16 [-0.27, -0.04]*
Evidence insensitivity	0.14 [-0.01, 0.28]	0.07 [-0.06, 0.20]	0.31 [0.15, 0.47]*
No contradiction	n/a	-0.24 [-0.36, -0.12]*	0.49 [0.37, 0.60]*

Value judgment	n/a	0.39 [0.27, 0.51]*	0.21 [0.07, 0.34]*
No matter of fact	n/a	-0.30 [-0.42, -0.19]*	0.47 [0.33, 0.60]*
Personal feelings normative	n/a	0.14 [0.04, 0.23]*	0.65 [0.55, 0.75]*
Unprovable	n/a	-0.18 [-0.30, -0.06]*	0.44 [0.30. 0.59]*
Personal feelings descriptive	n/a	0.11 [-0.02, 0.21]	0.68 [0.59, 0.77]*
No deference	n/a	0.09 [-0.01, 0.20]	0.47 [0.35, 0.60]*

Note. Numbers in brackets show 95% Confidence Intervals.
* indicates loadings whose confidence intervals do not cross 0

General Discussion

Recent work has suggested that there might be a single "symbolic" construct that underlies several diverse dimensions of belief. In these studies, we sought to determine how these hypothesized dimensions are related. Study 1 found that important beliefs scored higher on all the dimensions, except for evidence insensitivity. Study 2 found that all dimensions loaded on a single factor – except for evidence insensitivity. Study 3 pursued the null findings for evidence insensitivity. We found that this item loaded on a subjectivity factor. However, this factor was orthogonal to the symbolic factor. Across the latter two studies, we consistently recovered a two facet version of the symbolic factor, although its interpretation is obscure. We also found no evidence against the unimodality of the factor scores obtained in any study.

Here, we interpret our results as showing something about the dimensions on which peoples beliefs *actually vary*. On another view, our results show the way people *think of their own beliefs*. Our interpretation relies on the assumption that our participants are accurately reporting where their beliefs land on the dimensions we assessed. This seems like a pretty safe assumption for properties like whether a belief is important to the participant. But for a property like evidence insensitivity, it might be that participants are unaware or biased in their judgments of where the belief actually stands. Future work might test the assumption that participant ratings of belief dimensions are accurate by comparing them with third party ratings.

If we do interpret the results as saying something about the ontology of belief, they are a bit surprising. One might have expected that there would be two sorts of beliefs – symbolic beliefs that aim to represent identity and other beliefs that aim to represent facts about the world. In line with this view, we found a symbolicness factor that tracked whether a belief represented identity. However, two of our findings are a bit in tension with this view. For one, if there were two sorts of beliefs (symbolic and non-symbolic), we might expect symbolicness to be bimodally distributed. Instead, we found it was unimodally distributed. Second, if symbolic beliefs did not represent facts, then we might expect beliefs high in symbolicness to be unresponsive to evidence. Instead, judgements about whether a belief was evidence sensitive loaded on an uncorrelated, "subjectivity," factor.

Looking at Study 3, we see how fact representation and symbolicness come apart in two different ways. First, some beliefs shared by participants in the study were not symbolic but nonetheless subjective. Here's an example of one: "Silver is more beautiful than gold." Its subjective but not symbolic status is intuitive: there probably is not a fact of the matter about the aesthetic superiority of certain precious metals, and it is also hard to imagine how feelings about silver might be core to someone's being. Second, people shared beliefs that were symbolic but not subjective. Here's an example from one participant: "Global warming is human-caused." The participant probably thought this was an objective belief; but it is also plausible that this might be core to their identity.

These findings seem to suggest something larger about how to understand symbolicness. It does not seem that there are symbolic beliefs which represent identity and then other beliefs which represent facts about the world. Instead, it seems that beliefs vary on two continuous dimensions. At least broadly speaking, one dimension seems to be about the degree to which a belief aims to accurately represent the world, while the other seems to be about the degree to which the belief aims to express the believer's identity. Importantly, these two dimensions are independent. If a belief is high on one dimension, it does not necessarily have to be low on the other.

It might be helpful to interpret our finding through a metaphor. Different winter coats can serve different functions. One might primarily serve to keep its wearer warm, while another might make them appear more stylish. But, importantly, the same winter coat might serve multiple functions. Just because it keeps its wearer warm, this does not preclude it from making its wearer stylish. The present results suggest that beliefs are like this as well. Just because a belief represents identity, this does not preclude it from representing facts about the world. The same belief can serve multiple representational functions.

References

Goodwin, G. P., & Darley, J. M. (2008). The psychology of meta-ethics: Exploring objectivism. *Cognition*, *106*(3), 1339-1366.

Kahan, D. M. (2015). The expressive rationality of inaccurate perceptions. *Behavioral & Brain Sciences*, 40, 26-28.

Mercier, H. (2020). *Not born yesterday: The science of who we trust and what we believe*. Princeton University Press.

Mercier, H. (2021). How good are we at evaluating communicated information?. *Royal Institute of Philosophy Supplements*, 89, 257-272.

Metz, S. E., Liquin, E. G., & Lombrozo, T. (2023). Distinct profiles for beliefs about religion versus science. *Cognitive science*, 47(11), e13370.

Pölzler, T., & Wright, J. C. (2019). Empirical research on folk moral objectivism. *Philosophy Compass*, 14(5), e12589.

Ross, R. M., & Levy, N. (2023). Expressive responding in support of Donald Trump: An extended replication of Schaffner and Luks (2018). *Collabra: Psychology*, 9(1), 68054.

Sarkissian, H., Park, J., Tien, D., Wright, J. C., & Knobe, J. (2011). Folk moral relativism. *Mind & Language*, 26(4), 482-505.

Sakaluk, J. K., & Short, S. D. (2017). A methodological review of exploratory factor analysis in sexuality research: Used practices, best practices, and data analysis resources. *The Journal of Sex Research*, 54(1), 1-9.

Sommer, J. (2024). In the Mind or in the World? Types of Beliefs and the Locality of Evidence. PsyArXiv.

Sperber, D. (1997). Intuitive and reflective beliefs. *Mind & language*, 12(1), 67-83.

Strohminger, N., & Nichols, S. (2014). The essential moral self. *Cognition*, *131*(1), 159-171.

Van Leeuwen, N., Weisman, K., & Luhrmann, T. M. (2021). To believe is not to think: A cross-cultural finding. *Open Mind*, *5*, 91-99.

Van Leeuwen, N. (2022). Two concepts of belief strength: Epistemic confidence and identity centrality. *Frontiers in Psychology*, *13*, 939949.

Vesga, A., Van Leeuwen, N., & Lombrozo, T. (2024). Evidence for distinct cognitive attitudes of belief in theory of mind. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 46).

Westra, E. (2023). Symbolic belief in social cognition. *Philosophical perspectives*, *37*(1), 388-408.