

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Mechanisms of Belief Persistence in the Face of Societal Disagreement

Permalink

<https://escholarship.org/uc/item/3380n01h>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 44(44)

Authors

Oktar, Kerem
Lombrozo, Tania

Publication Date

2022

Peer reviewed

Mechanisms of Belief Persistence in the Face of Societal Disagreement

Kerem Oktar (oktar@princeton.edu)

Tania Lombrozo (lombrozo@princeton.edu)

Department of Psychology, Princeton University
Princeton, NJ 08540 USA

Abstract

People have a remarkable ability to remain steadfast in their beliefs in the face of large-scale disagreement. This has important consequences (e.g., societal polarization), yet its psychological underpinnings are poorly understood. In this paper, we answer foundational questions regarding belief persistence, from its prevalence to variability. Across two Experiments ($N = 356$, $N = 354$), we find that participants are aware of societal disagreement about controversial issues, yet overwhelmingly (~85%) do not question their views if asked to reflect on this disagreement. Both studies provide evidence that explanations for persistence vary across domains, with epistemic and meta-epistemic explanations among the most prevalent.

Keywords: disagreement; polarization; domain; judgment; persistence; conciliation; suspension; controversial

Introduction

For most of our important beliefs – from abortion to vaccination – we know that *millions* disagree with us, yet we manage to remain steadfast in our views. Surprisingly, our understanding of the psychological mechanisms that support this persistence is quite patchy and siloed (Oktar & Lombrozo, 2022)—and there is little empirical evidence that grounds current theory.

In this paper, we present two experiments that are among the first to investigate whether, when, and how people persist in their beliefs amid controversy. These experiments lead to the following novel contributions. First, we find that people are broadly cognizant of mass disagreement about controversial issues, yet do not question their own views when they reflect on this disagreement. Second, we show that explanations for persistence can be clustered into four main categories: rejection, epistemic, non-epistemic, and meta-epistemic, though these are not equally prevalent. Third, we find that some explanations for persistence vary in a systematic and interpretable manner across domains. In the General Discussion, we consider the implications of this work for belief revision and polarization.

Note that our work is not about why large-scale disagreement occurs (see the references in the next section for cross-disciplinary theory and evidence on this question). Instead, this paper is about what enables people to persist in their beliefs amid large-scale controversies—the mechanisms that allow people to privilege their opinions when the issues are obviously split from a societal perspective.

Persistent Disagreement

Disciplines agree that people disagree. There are vast literatures in social (e.g., Valdesolo & Graham, 2016), developmental (e.g., Heiphetz et al., 2013), and cognitive (e.g., Kruglanski, 2004) psychology, public opinion research (e.g., Zaller, 1992), epistemology (e.g., Frances, 2014; Hardwig, 1985), political science (e.g., Huckfeldt et al., 2014; Iyengar & Westwood, 2015), economics (e.g., Enke & Zimmerman, 2019; Golman et al., 2016), and computational social science (e.g., Axelrod, 1997; Flache et al., 2017) that are broadly pertinent to the question of persistent disagreement. Though a comprehensive overview of this work is beyond the scope of the current paper, we can highlight common themes. For instance, much of this work is concerned with what ought to be done when others have beliefs that differ from our own—and the clearest picture of potential responses to disagreement comes from recent work in philosophy.

Responses to Disagreement

Focusing on disagreement between peers, epistemologists have identified three kinds of responses to disagreement: one can suspend judgment about the issue, move one's beliefs towards those of others, or persist in one's prior views (Frances & Matheson, 2019). Experimental philosophers have begun investigating how people respond to such disagreement. Heinzelmann, Hölzgen, and Tran (2021) found that when participants attended an informational workshop in which they also participated in small-group discussions on ethical issues with disagreeing others, the majority persisted in their views (or became even more confident), echoing the findings of work on polarization in social psychology (Valdesolo & Graham, 2016).

Though epistemologists have not converged on what the rational response to disagreement should be, the large political (Pew Research, 2014), moral (Gallup, 2007), and scientific (Pew Research, 2015) divisions within the U.S. offer preliminary evidence that people often sustain their views, even when this controversy occurs “at scale” (vs. between peers). What explains this widespread persistence?

Explaining Persistence

Oktar and Lombrozo (2022) propose a four-factor model for understanding persistence in the face of controversy. The first

factor, *rejection*, captures the possibility that belief persistence may result from insufficient awareness or consideration of disagreement (Bendana & Mandelbaum, 2021; Shamir & Shamir, 1997). Yet even in the face of *active* disagreement, people often persist in their beliefs (at least in small-group settings; Heinzlmann et al., 2021).

Recent work in philosophy highlights the importance of the second factor, the *epistemic*, when evaluating such active disagreement (see Frances, 2014). If others have weaker evidence than oneself, for example, it can be rational to sustain one's views (Ross & Ward, 1996; Steele & Stefansson, 2015).

Beyond epistemic considerations, there are many costs and benefits associated with one's beliefs (Golman et al., 2016). Many of these come from the social functions of beliefs, such as signaling group membership (Rokeach, 1960). The third factor, the *non-epistemic*, captures the idea that such costs and benefits can independently guide one's beliefs (Cusimano & Lombrozo, 2021; Davoodi & Lombrozo, 2021; Iyengar & Westwood, 2015; Mason & Wronski, 2018).

Finally, there are some issues for which others' views are fundamentally irrelevant, perhaps because they are a matter of opinion (e.g., what the best color is) or because they are unknowable (e.g., whether there is an afterlife). The fourth factor, the *meta-epistemic*, covers such cases (Goodwin & Darley, 2012; Gollwitzer & Oettingen, 2019).

These four factors (rejection, epistemic, non-epistemic, and meta-epistemic) offer a potentially exhaustive classification scheme for persistence amid controversy. Yet the lack of relevant empirical data makes it impossible to answer even the most basic questions surrounding persistence: for instance, are any of these paths utilized more frequently than others? Do the paths that people pursue when evaluating disagreement depend on the content of the issue, as suggested by epistemologists (Lackey, 2010) and work on cross-domain variation in judgment and decision-making (Oktar & Lombrozo, 2022; Pachur & Spaar, 2015)? Are some components of these factors more frequently utilized or important than others (e.g., subjectivity vs. knowability for meta-epistemic explanations)?

Overview of Experiments

We present two pre-registered experiments that answer these questions by testing the four-factor model (pre-registrations, materials, and data are available in our open-access repository at <https://osf.io/qfz7m/>). In Experiment 1, we elicit open-ended responses to disagreement concerning one of four highly controversial statements from each of four domains (science, politics, morality, and religion). We examine whether participants tend to persist in their views, and we document the sorts of explanations they generate.

In Experiment 2, we use the same stimuli in a closed-ended design. Beyond serving as a replication, this allows us to test whether the prevalence and cross-domain variation in explanations persists when participants must evaluate explanations (vs. generate explanations themselves).

Experiment 1

In Experiment 1, we presented participants with controversial statements from science, religion, politics, and morality. Each participant received one of four statements from a given domain and indicated whether they agreed or disagreed with the statement. Participants then estimated the societal distribution of opinion regarding that statement, and finally explained whether the disagreement regarding the issue made them question their own views.

Our primary aim was to investigate three foundational questions regarding persistent disagreement: First, do people in fact overwhelmingly persist in their beliefs in the face of massive disagreement, or do they express doubt upon reflection? Second, do the ways in which people sustain their beliefs accord with Oktar and Lombrozo's (2022) four-factor model? Finally, do people rely on different mechanisms to persist in their beliefs across domains? Given that the explanandum (i.e., the disagreement) is the same across domains, one might expect people's explanations to be relatively similar. On the other hand, people typically reason in systematically different ways across domains (Goodwin & Darley, 2012; Inbar, Cone, & Gilovich, 2010), so they may make different attributions for disagreement as well.

These three questions are foundational in the sense that they can ground and organize much subsequent philosophical and psychological inquiry. Without knowing whether, when, and how people persist in their beliefs, it is difficult to gain empirical footing on further questions about disagreement in particular, and perhaps the nature of controversial belief in general. To the best of our knowledge, these experiments are the first to try to establish such a foundation.

Methods

Participants Participants were 356 adults (164 male, 188 female, 4 other, mean age = 33) recruited on Prolific in exchange for monetary compensation (\$0.50 for a 4-minute study). Participation across all studies was restricted to users currently residing in the United States with an approval rating $\geq 98\%$ on at least 100 tasks. Repeat participation within and across studies was restricted using the Prolific platform.

Materials and Procedure Participants were first assigned to one of 16 statements across four domains (see Table 1).

On the first screen, participants indicated their own opinions about their assigned statement by responding to the question "Do you personally agree that X?" with 'Yes,' 'No,' or 'I have not made up my mind about this claim.' Afterwards, they rated how confident they were in their responses on a five-point scale from 'Very Confident' to 'Not at all Confident.' In the interest of space, we do not analyze data concerning confidence here.

Participants next estimated the distribution of opinions in the US regarding the statement. They were asked to drag three sliders, all initialized at zero, to indicate a population percentage for three groups: the percentage of the US population that agrees with the claim, disagrees with the

claim, and has not made up their mind yet. The three estimates were required to sum to 100.

Participants then encountered our key prompt and measure:

In the previous question, you indicated that [N]% of people in the US share your view about whether X. This means that [100-N]% of people in the US do not share your view.

Does the fact that [100-N]% of people in the US do not share your view about whether X make you question your own view? Why or why not? Please carefully explain your thoughts in a few sentences.

There are no right or wrong answers to this question, we are simply interested in what you think.

After providing a response in a text box, participants answered demographic questions (age, sex, educational background, level of religiosity, political affiliation) and received debriefing information.

Results

Explanation coding Two independent coders coded all open-ended justifications for persistence for the presence of four pre-defined categories (taken from Oktar & Lombrozo, 2022), as well as new categories that were identified from the responses themselves. The pre-defined categories included: rejection of premise (denying the presence or significance of disagreement), appeal to non-epistemic factors (intrapersonal, interpersonal), appeal to epistemic factors (evidence, processing of evidence), and appeal to meta-epistemic factors (objectivity, knowability). The coders also included a fifth category for “other,” which initially included responses that were blank, simply claimed ignorance (“I don’t know”), or were irrelevant (e.g., non-sensical, did not address the prompt, etc.). Anything sensible not covered by the rubric was initially coded as “not covered.”

Interrater reliability on these initial five codes (collapsing “other” and “not covered”) suggested substantial agreement between our raters, $\kappa = .62$, 95% CI [.54, .69] (see Landis & Koch, 1977). To reach our final coding categories, the coders discussed discrepancies in their coding (fewer than 5% of judgments), as well as responses marked as “not covered” (21.3% of responses). These could be resolved by either changing prior assignments or creating new categories. Two

new categories emerged from this discussion. First, some responses were coded as “opinion” if they claimed something to the effect that ‘everyone is entitled to their opinions’ (e.g., “No it does not make me question my choice because everyone has their opinions and feelings.”) – these responses were plausibly meta-epistemic, but were sufficiently vague that they were segregated from unambiguously meta-epistemic responses. Second, some responses were coded as “unclear” - participants provided a response to the question, but not one that unambiguously fell into existing categories nor suggested a coherent novel category (e.g., “Most people think the rich just get richer and never pay their share”). Explanations were coded as “uncodable” if, after this process, they were not classified as rejections, epistemic, non-epistemic, meta-epistemic, opinion, or unclear.

Perceived Consensus & Persistence Given that our main hypotheses concern responses to large-scale disagreement, we first checked whether participants perceived our items as being controversial. Across all participants, the perceived population agreement with one’s own position was nearly half of the population ($M = 46.3\%$, $SD = 22.9$), which is the maximal population disagreement for a binary issue. The extent of perceived agreement was relatively stable across domains, with a minimum of 40.8% for Politics, and 52.3% for Science. Further, both overall and domain-level perceived agreement were normally distributed.

Supporting our first hypothesis, an overwhelming majority of participants reported not being influenced by mass disagreement. Overall, 86.0% (95% CI [81.8, 89.3]; test that the proportion was observed by chance: $\chi^2(1, N = 356) = 182.7, p < .001$) of participants did not report questioning their views in the face of disagreement (i.e., persisted in their views). This proportion was consistent across domains and ranged from 84% for religion to 89% for morality. Our subsequent analyses of persistence use data from this 86% ($N = 306$) who in fact persisted in their views.

Explanations for Persistence By a large margin, the most common type of explanation was epistemic (57.3% of codable responses), followed by meta-epistemic (13.2%), rejection (10.5%), and non-epistemic (5.6%). Table 2 lists all codes, with examples from actual participant responses, and explanations for why they exemplify the relevant code.

Table 1: List of Statements Used in Experiments 1 and 2

Science	Religion	Morality	Politics
Human activity is a major contributor to climate change.	Holy scripture is the literal word of God.	Doctor-assisted suicide is immoral.	Upper-income people in the US pay too much in taxes.
Genetically modified crops are safe for human consumption.	Humans have souls.	Eating meat is immoral.	Gun laws in the us are not strict enough.
Power lines do not cause cancer.	All animals have souls.	Having a baby outside of marriage is immoral.	The federal government has too much power.
Vaccines do not cause autism.	There is an afterlife.	Testing products on animals is immoral.	The US has too much immigration.

Note. Statements were chosen to be as controversial as possible (i.e., close to an even split in opinions in the US).

Table 2: Coding Categories and Example Responses from Experiment 1

Main Code	Subcodes	Ratio	Example Participant Response	Explanation of Example
Epistemic	Evidential Competence	27.1%	“I honestly think that people believing this are unintelligent. There’s no evidence of this at all (...)”	Appeals to asymmetries in both competence and evidence to justify persistence.
Meta-epistemic	Subjective Unknowable Difficult	7.5%	“I don’t think this is something that can be proved or disproved by living people. This is just a personal choice.”	The unknowability of the claim is taken to justify subjectivity and persistence.
Non-epistemic	Intrapersonal Interpersonal	2.6%	“I also think people are inclined to believe in it because the thought of no longer existing is scary. It scares me too.”	The intrapersonal cost of changing one’s view about the issue justifies persistence.
Rejection	Awareness Majority	4.9%	“I have never heard anyone claim this view, not even in joking or conspiracy theory (...). Thus, I assume it is a very niche view if it exists at all.”	The participant does not recognize the presence of large-scale disagreement, and hence persists in their view.
Opinion		16.3%	“Everyone is allowed their own opinion and can do what they want, as long as they don’t try to force me (...)”	Some participants consider people’s opinions to be fundamentally isolated from each other.
Unclear		36.9%	“I think and believe that every animals have soul. And it should have.”	It is difficult to interpret these responses clearly.
Uncodable		8.4%	“Nah, if anything historically the majority have not necessarily always aligned with what was objectively ethical (...)”	These responses are interpretable but difficult to cluster with each other or any of the existing items.

Note. Ratio refers to the percentage of responses from participants who persisted that were classified as belonging to that main code by both coders post discussion. The codes were used mostly exclusively (3.7% of responses fit multiple codes).

Cross-domain Variation Our data also supported the hypothesis that explanations for persistence differ across domains. We conducted mixed effects logistic regression analyses, predicting the classification of our main categories from domain, with random intercepts for statements.¹ These analyses revealed that *epistemic* explanations were significantly more likely to be used to justify persistence in *science*, whereas *meta-epistemic* explanations were more likely to be used for *religion* (see Figure 1).




Figure 1: Explanations of Persistence Across Domains in Experiment 1. Error bars show bootstrapped 95% confidence intervals. Asterisk denotes $p < .001$ for deviation-coded domain contrasts.

¹ These were four separate regressions for each explanation (with binary dependent variables for explanation type, and binary independent variables for statement domain). We used the most complex random effects structure that fit for each model. For the epistemic and meta-epistemic, this included nested random

Discussion

The results of Experiment 1 provide clear evidence in favor of our initial hypotheses: (i) people report that they persist in their beliefs in the face of mass disagreement, (ii) self-generated explanations for persistence can be captured by some of the theoretically motivated paths to persistence; and (iii) the use of these explanation shows systematic variation across domains, with epistemic explanations most common overall and especially for science, and meta-epistemic explanations most common for religion.

Experiment 2

In Experiment 2, we develop a closed-ended analogue of our first study and investigate whether the patterns observed in our first study replicate under more controlled circumstances. Instead of answering an open-ended question about their reactions to disagreement, participants provided ratings for eleven Likert-type items that corresponded to the pre-defined coding categories used in the prior study. Beyond serving as a replication, this allows us to test whether the prevalence of

intercepts of prompt. For the other explanations, no random effect models converged. Note that restricting analyses to participants who initially did not suspend judgment does not change these results, nor does including the participants’ perceived distributions of opinion for the issue (i.e., percent who agree, disagree, suspend).

explanation types and their variation across domains persists when participants must explicitly evaluate a set of explanations, vs. generate explanations themselves. By having participants rate items corresponding to factors that were infrequently generated in Study 1 (such as non-epistemic and rejection), we can also gain greater insight into whether all four factors of the four-factor model capture relevant variation in judgments.

Methods

Participants Participants were 354 adults (153 male, 196 female, 4 other, 1 unidentified, mean age = 34) recruited as in Study 1. An additional two participants were excluded for failing a pre-registered attention check.

Materials and Procedure Participants were assigned to one of 16 statements from Experiment 1, and shown the same opinion, confidence, and societal distribution measures. The content of the studies diverged after this point. Instead of an open-ended question, participants first gave a binary (yes / no) response to whether “the fact that (100 - N) % of people in the US do not share your view about whether X make you question your own view?” After responding to this question, participants answered 11 Likert-type agreement items, from Strongly Disagree (1) to Strongly Agree (7), in random order. For all items, participants rated how much they agreed with each statement as an explanation for why the disagreement did or did not make them question their own views.

Slightly abridged versions of these eleven items are shown below, with labels that indicate the corresponding coding categories from Experiment 1.

Epistemic Items

Processing - Competence: People who disagree with me about Y are not as good as I am at evaluating this issue.

Processing - Bias: People who disagree with me about Y are more biased than I am on this issue.

Evidential - Quality: People who disagree with me about Y have evidence (...) that is less good than my own.

Evidential - Quantity: People who disagree with me about Y have less evidence about this issue than I do.

Meta-Epistemic Items

Subjective: Whether Y is more a matter of opinion than fact; there are no right or wrong answers about it.

Hard: Knowing what to believe about Y is very difficult.

Unknowable: People will never know whether Y; it is an issue that is fundamentally unknowable.

Non-Epistemic Items

Interpersonal: (...) changing my view about Y could damage my relationships with important people in my life.

Intrapersonal: (...) changing my view about Y would threaten who I am as a person.

Rejection of Premise Items

Awareness: I have deeply considered why people disagree about whether Y.

Majority: People disagreeing about Y does not matter to me, given that enough people actually agree with me.

Finally, participants answered the same demographic questions used in Experiment 1.

Results

Perceived Consensus & Persistence As in Experiment 1, we first confirmed that issues were perceived to be controversial (they were: across all participants, the percentage of perceived agreement with one’s position in the population was nearly half, at 47.6%). We also verified whether, with our close-ended measure, most participants reported that they did not question their views on the basis of large-scale disagreement (they overwhelmingly did not: 87.0%, 95% CI = [83.8, 90.8]). As in Experiment 1, these proportions were fairly consistent across domains, with disagreement ranging from 52% in science to 45% for politics, and persistence from ~80% for science to 93% for morality.

Explanations for Persistence To reduce the dimensionality of our 11-item measure of explanations for persistence, and in line with our pre-registered analysis plan, we conducted the following factor analysis (in accordance with recent recommendations; see Jackson, Gillaspay, & Purc-Stephenson, 2009). The analysis was based on data from the 309 participants who reported persistence.

After verifying that our data were suitable for factor analysis, we implemented a factor analysis with a minimal residual factoring method and an oblique rotation (oblimin). We determined the number of factors to include by conducting a parallel analysis, which suggested a three- or four-factor solution. We retained four factors.

The four-factor solution explained 52% of total variance in ratings, and the four factors respectively accounted for 23%, 13%, 12%, and 4% of variation in the data. Importantly, the factor loadings were exactly in line with our a priori expectations (see Figure 2). These findings help validate the coding categories employed in Experiment 1, and additionally replicate the relative prevalence of response types, with higher agreement on the epistemic items ($M = 3.84$) followed by the meta-epistemic ($M = 3.74$), rejection ($M = 3.66$), and the non-epistemic ($M = 2.77$).




Figure 2: Factor analysis on responses from Experiment 2. Factor loadings greater than .35 are shown in white text. Blue indicates positive loadings, and red negative.

Cross-domain Variation Having reduced the dimensionality of our data, we proceeded to investigate whether we replicated the domain-dependence observed in Experiment 1. To do so, we first extracted factor scores from the four-factor solution (using Thurstone's method, though the results presented below also hold for the Bartlett method among others; see Grice, 2001). We then conducted separate one-way ANOVAs with domain as the independent variable for each of our factors. This analysis revealed significant domain variation for the epistemic, $F(3,305) = 4.22, p < .01$, non-epistemic, $F(3,305) = 3.34, p < 0.05$, and meta-epistemic factors, $F(3,305) = 25.5, p < .001$, but not for the rejection factor, $F(3,305) = 1.65, p = 0.18$.

As predicted, we found that the epistemic factor was utilized more frequently for scientific controversies than the average ($M = 0.3; t(73) = 2.33, p < .05$), and that the meta-epistemic factor was more common for religious controversies than the average, ($M = 0.54; t(74) = 5.28, p < .001$). We also observed a significant effect in the opposite direction for science ($M = -0.58; t(73) = -6.28, p < .001$); and a weaker positive effect for morality ($M = 0.18; t(79) = 2.15, p < .05$). For non-epistemic, we did not have predictions; we found weak effects for morality ($M = 0.2; t(79) = 2.12, p < .05$) and religion ($M = -0.24; t(74) = -2.71, p < .01$).

Discussion

Experiment 2's findings echo those of Experiment 1. Once again, we find that participants report sustaining their beliefs in the face of societal disagreement and doing so in a domain-dependent manner. Importantly, the results obtained in this study were derived using a bottom-up approach on closed-ended responses (vs. human coding as in Experiment 1). The fact that this method yielded categories that accord with the four-factor framework thus offers support for the framework beyond that found in Experiment 1, and gives us finer-grained insight into the mechanisms of persistence. For instance, the factor loadings we obtained suggest that amount of evidence and competence may be more characteristic of epistemic persistence (above perceptions of bias and evidential quality), whereas knowability may be a more characteristic component of meta-epistemic persistence (above subjectivity). The consistent domain variation observed in Experiments 1 and 2 further indicates that these mechanisms might be especially important for understanding persistent disagreement in science and religion.

General Discussion

Modern life is full of controversy, disagreement, and disdain. Americans, for instance, are increasingly polarized on political (Pew Research, 2014), moral (Gallup, 2017), religious, and scientific (Pew Research, 2015) matters. This polarization carries significant consequences: widening rifts between liberals and conservatives across the globe are corroding the social and institutional foundations of deliberative democracies (Carothers & O'Donohue, 2019; Svobik, 2019). The ubiquity of such large-scale disagreement

across highly influential domains (e.g., religion or science) raises important descriptive and prescriptive questions: What are the factors that determine whether, when, and why people remain steadfast in their beliefs, and should they?

As reviewed in the introduction, many literatures across the social sciences bear on these questions. However, there have been few attempts to integrate these findings into a coherent theoretical framework for understanding the mechanisms that enable persistent large-scale disagreement (see Oktar & Lombrozo, 2022). Perhaps as a consequence, there has been little empirical work documenting basic facts concerning persistence, from its prevalence to its variability.

The two experiments in this paper provide preliminary answers to these foundational issues. Across both studies, we find strong evidence of persistence: people are simultaneously aware of societal disagreement about controversial issues, yet overwhelmingly (~85%) do not question their views if asked to reflect on this disagreement. Across both experiments, we additionally find evidence for four factors explaining persistence, as well as variation in their application across domains: rejection, epistemic, non-epistemic, and meta-epistemic. This variation is important because interventions designed to promote open-mindedness or conciliation are unlikely to be effective if they fail to address the source of persistence: for instance, changing meta-epistemic assumptions is unlikely to promote conciliation about science, and claims of epistemic peerhood are unlikely to promote open-mindedness about religion.

Several important questions regarding persistence remain unexplored. For example, it is currently unknown whether the explanations people provide when faced with disagreement are causally linked with their tendency to persist. It is possible (though unlikely) that they are merely post-hoc rationalizations, and that a different set of drivers underlie persistence. In ongoing work we are addressing this by experimentally manipulating explanations of disagreement.

Furthermore, we do not know whether the tendency to rely on particular explanations is related to other important behaviors, such as preferential information seeking, mind-changing, or social judgment. Future work would also benefit from using a broader range of measures (including implicit and behavioral), as well as attending to cultural and contextual variation.

Despite these limitations, the results presented in this paper constitute the first pieces of empirical evidence obtained in the service of comprehensively understanding the individual mechanisms that support large-scale, persistent disagreement. We hope that these findings will set the stage for future inquiry that further elucidates why, when, and how the seeds of societal polarization are sown. Such understanding can (hopefully) guide the design of interventions that bridge the ever-widening rifts in our societies.

Acknowledgments

We thank Tom Griffiths, Eldar Shafir, Corey Cusimano, Thiago Varella, and other members of the Concepts and Cognition Lab for their valuable feedback.

References

- Axelrod, R. (1997). The dissemination of culture: A model with local convergence and global polarization. *Journal of Conflict Resolution*, 41(2), 203-226.
- Bendaña, J., & Mandelbaum, E. (2021). *The Fragmentation of Belief*. In: *The Fragmented Mind*. Editor: Borgoni, C., Kindermann, D., & Onofri, A. Oxford University Press.
- Carothers, T., & O'Donohue, A. (Eds.). (2019). *Democracies divided: The global challenge of political polarization*. Washington, D.C.: The Brookings Institution.
- Cusimano, C., & Lombrozo, T. (2021). Morality justifies motivated reasoning in the folk ethics of belief. *Cognition*.
- Davoodi, T., & Lombrozo, T. (2021). Explaining the existential: Scientific and religious explanations play different functional roles. *Journal of Experimental Psychology: General*.
- Enke, B., & Zimmermann, F. (2019). Correlation Neglect in Belief Formation. *The Review of Economic Studies*, 86(1), 313-332.
- Flache, A., Mäs, M., Feliciani, T., Chattoe-Brown, E., Deffuant, G., Huet, S., & Lorenz, J. (2017). Models of social influence: Towards the next frontiers. *Journal of Artificial Societies and Social Simulation*, 20(4).
- Frances, B. (2014). *Disagreement*. John Wiley & Sons.
- Frances, B., & Matheson, J. (2019). *Disagreement*. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University.
- Gallup, Inc. (2007, June 4). Americans Rate the Morality of 16 Social Issues. Gallup. <https://news.gallup.com/poll/27757/Americans-Rate-Morality-Social-Issues.aspx>
- Gollwitzer, P. M., & Oettingen, G. (2019). Goal attainment. *The Oxford handbook of human motivation*, 245-268.
- Golman, R., Loewenstein, G., Moene, K. O., & Zarri, L. (2016). The Preference for Belief Consonance. *Journal of Economic Perspectives*, 30(3), 165-188.
- 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.
- Grice, J. W. (2001). Computing and evaluating factor scores. *Psychological Methods*, 6(4), 430.
- Hardwig, J. (1985). Epistemic Dependence. *The Journal of Philosophy*, 82(7), 335-349.
- Heinzelmann, N., Hölzgen, B. T., & Tran, V. (2021). Moral discourse boosts confidence in moral judgments. *Philosophical Psychology*, 34(8), 1192-1216.
- 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.
- Huckfeldt, R., Johnson, P. E., & Sprague, J. (2004). *Political Disagreement: The Survival of Diverse Opinions within Communication Networks*. Cambridge University Press.
- Inbar, Y., Cone, J., & Gilovich, T. (2010). People's intuitions about intuitive insight and intuitive choice. *Journal of Personality and Social Psychology*, 99(2), 232.
- Iyengar, S., & Westwood, S. J. (2015). Fear and Loathing across Party Lines: New Evidence on Group Polarization. *American Journal of Political Science*, 59(3), 690-707.
- Iyengar, S., & Westwood, S. J. (2015). Fear and loathing across party lines: New evidence on group polarization. *American Journal of Political Science*, 59(3), 690-707.
- Jackson, D. L., Gillaspay Jr, J. A., & Purc-Stephenson, R. (2009). Reporting practices in confirmatory factor analysis: an overview and some recommendations. *Psychological Methods*, 14(1), 6.
- Kruglanski, A. W. (2004). *The Psychology of Closed Mindedness*. Taylor & Francis Group.
- Lackey, J. (2010). *Learning from words: Testimony as a source of knowledge*. OUP Oxford.
- Landis, J. R., & Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics*, 33(1), 159-174.
- Mason, L., & Wronski, J. (2018). One tribe to bind them all: How our social group attachments strengthen partisanship. *Political Psychology*, 39, 257-277.
- Oktar, K. and Lombrozo, T. (2022). Deciding to be Authentic: Intuition is Favored Over Deliberation When Authenticity Matters. *Cognition*. 223, 105021.
- Oktar, K., & Lombrozo, T. (2022). Sustaining Belief in a Contentious World. (Manuscript in Prep).
- Pachur, T., & Spaar, M. (2015). Domain-specific preferences for intuition and deliberation in decision making. *Journal of Applied Research in Memory and Cognition*, 4(3), 303-311.
- Pew Research. (2014, June 12). Political Polarization in the American Public. Pew Research Center - U.S. Politics & Policy. <https://www.pewresearch.org/politics/2014/06/12/political-polarization-in-the-american-public/>
- Pew Research. (2015, October 22). What U.S. Religious Groups Think About Science Issues. Pew Research Center Science & Society. <https://www.pewresearch.org/science/2015/10/22/science-and-religion/>
- Rokeach, M. (1960). *The open and closed mind* (p. 447). Basic Books.
- Ross, L., & Ward, A. (1996). Naive Realism in Everyday Life: Implications for Social Conflict and Misunderstanding. In: *Values and Knowledge*. Psychology Press.
- Shamir, J., & Shamir, M. (1997). Pluralistic Ignorance Across Issues and Over Time: Information Cues and Biases. *The Public Opinion Quarterly*, 61(2), 227-260.
- Svolik, M. W. (2019). Polarization versus democracy. *Journal of Democracy*, 30(3), 20-32.
- Valdesolo, P., & Graham, J. (Eds.). (2016). *Social Psychology of Political Polarization* (1st ed.). Routledge.
- Zaller, J. R. (1992). *The nature and origins of mass opinion*. Cambridge University Press.