Measuring the Belief System of a Person

Mark J. Brandt

Michigan State University

Acknowledgements: Thanks to Dan Wisneski for useful comments on a prior draft of this manuscript. This manuscript is part of a project that received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 759320).

Address correspondence to brandt15@msu.edu

Draft date: January 4, 2022

The paper has been accepted for publication at the *Journal of Personality and Social Psychology*.

© 2022, American Psychological Association. This paper is not the copy of record and may not exactly replicate the final, authoritative version of the article. Please do not copy or cite without authors' permission. The final article will be available, upon publication, via its DOI: 10.1037/pspp0000416

Data and code available here:

https://osf.io/u9bew/?view_only=5d0bbefcb9434147a8c953f06cfc41dd

Abstract

Theories of belief system structure and dynamics assume that belief systems are a person-level construct. However, measures of belief system structure do not measure the structure of person-level belief systems and instead measure aggregated belief system structure (e.g., the belief system in a particular country). In this paper, I show that a measure of conceptual similarity between attitudes and identities of a belief system works as a valid, reliable, flexible, and efficient measure of person-level belief system structure. In Studies 1 (N = 387), 2 (N = 389), and 3 (N = 598), I show conceptual similarity judgments are reliable and are related to measures of political engagement, political knowledge, attitude consistency, and preference congruence as predicted by computational models of belief system dynamics. In Studies 4 (N = 981) and 5 (N = 983), I show that conceptual similarity judgments are affected by partisan frames and that changes in conceptual similarity judgments are associated with attitude change as predicted by computational models of belief system dynamics. Conceptual similarity judgments can be used with a variety of attitudes and identities in easy to administer studies. It provides a tool to fill an empirical gap identified by theories of belief system dynamics.

Measuring the Belief System of a Person

Belief systems shape our perceptions of our political world. They are networks of political attitudes and identities, the structure of which is thought to play a role in how attitudes shift and change, as well as how attitudes sometimes stay the same (Brandt & Sleegers, 2021). Belief systems also reflect what we see in our political world. They are shaped by the attitudes we see politicians express and the types of information we see in our news and social networks (Goldberg & Stein, 2018). The study of the structure of belief systems, however, is plagued by methodological shortcomings with standard methods only indirectly assessing how belief systems might be structured. These shortcomings prevent a rigorous evaluation and extension of theories of belief system structure. In this manuscript, I develop and validate a method for assessing person-level belief system structure.

Two theoretical models of belief systems have demonstrated how conceptualizing and modeling belief systems as networks of political attitudes and identities can account for a number of phenomenon in the literature (Brandt & Sleegers, 2021; Goldberg & Stein, 2018; see also Boutyline & Vaisey, 2017; Brandt, Sibley, & Osborne, 2019; Fishman & Davis, in press). Although the precise mathematics underlying these two models is different¹, both models assume that the political attitudes and identities that make up the belief system mutually influence each other. These mutual influences are what then drives belief system stability and change, as well as the stability and change of individual attitudes embedded in the belief system. For example, Brandt and Sleegers (2021) used simulations to show that (simulated) people with belief systems with many strong bi-directional relationships between attitudes and identities are the most likely to have stable belief systems over time. Goldberg

¹ Brandt and Sleegers' (2021) model is computationally instantiated with an Ising model, whereas Goldberg and Stein's (2018) model uses a parallel constraint satisfaction model. Ising models and parallel constraint satisfaction models are conceptually similar and share similar conceptual and theoretical foundations (e.g., Gurney, 1997; Hopfield, 1982).

and Stein (2018) examined a different phenomenon in their simulations. They showed that when (simulated) people learn that two attitudes of a belief system are connected (e.g., by seeing people in the population hold both of the attitudes), this can cause attitude change and polarization in the population. These two models are theoretical models that have been implemented using simulation studies and account for key empirical results in the literature.

Both the Brandt-Sleegers and Goldberg-Stein models assume that each person potentially has their own unique belief system. That is, the belief system is the property of the individual. This is consistent with how belief systems are often defined. For example, when Kinder and Kalmoe (2017) talk about ideology, which is a particularly sophisticated belief system, they describe it as being “in the mind” (p. 12; see also Morgan & Wisneski, 2017). It is also consistent with paradigmatic social psychological theories which locates attitudes and identities within people (Albarracin & Vargas, 2010; Banaji & Heiphetz, 2010; Eagly & Chaiken, 1993; Leach et al., 2008). Although it is certainly the case that attitudes and identities are influenced by factors outside the person (e.g., social context, position in social networks etc. Reicher, 2004; Visser & Mirabile, 2004), the attitudes and identities themselves are within persons.

It is also certainly the case that people will have similarities in their belief systems because they learn about how attitudes and identities co-occur through similar processes (e.g., experiences in the same political system, exposure to political messaging; Converse, 1964; Zaller, 1992). However, these outside influences and potential similarities are both thought to operate by changing connections between attitudes at the level of the individual (Goldberg & Stein, 2018). Consistent with the idea that belief systems operate within-persons, Brandt and Morgan (in press) analyzed data from longitudinal studies conducted in the United States and the Netherlands and found that in these samples average within-person correlations between political attitudes were dissimilar from between-person correlations between political

attitudes. Moreover, in supplementary analyses these authors found that there were large differences between a person's within-person correlations between attitudes and the average within-person correlation between the same attitudes in the sample. This suggests that the structure of a belief system may differ from person to person. In short, a person has a belief system and this belief system may not be the same as another person's belief system. This point is embedded in the Brandt-Sleegers and Goldberg-Stein theoretical models which assume that belief systems are within people.

The location of belief systems within people may seem like an obvious theoretical decision; however, it has an unsolved empirical implication. That is, there are no methods for measuring the belief system of a person. Currently, theoretical models of belief systems (Brandt & Sleegers, 2021; Goldberg & Stein, 2018) make assumptions about the structure of belief systems that can be used to conduct simulation studies. The outcomes of these simulation studies can then be compared with the outcomes of traditional empirical studies to assess whether or not the theoretical model made accurate predictions (see Brandt & Sleegers, 2021). What these theoretical models highlight is that a measure that can assess belief system structure is missing from our collective methodological toolbox.

Existing Approaches for Measuring Belief System Networks

Some methods have been used that purport to measure the structure of belief systems and build on similar network approaches as the Brandt-Sleegers and Goldberg-Stein models. The precise methods differ, but all of these approaches use cross-sectional survey data consisting of a variety of political attitudes (e.g., attitudes about abortion, taxes, gay rights etc) and political identities (e.g., partisan identification). The approaches then use this cross-sectional data to estimate the connections between the attitudes and identities as an estimate of the structure of the belief system. For example, using data from a cross-sectional survey, Boutyline and Vaisey (2017) estimated the connections between the attitudes and identities in

the belief system using squared correlations. Building on this work, other scholars (Brandt et al., 2019; Fishman & Davis, in press) estimated the connections between the attitudes and identities in the belief system using partial correlations, which index the unique associations between any pair of attitudes or identities in the belief system. Other scholars, without explicitly drawing on network approaches, also use cross-sectional correlations between attitudes and identities to make inferences about the structure of belief systems (e.g., Converse, 1964; Kalmoe, 2020; Malka et al., 2014).

Although the precise assumptions made by each of these approaches differ, they all draw on correlations from cross-sectional data. Such analyses cannot be used to make conclusions about individual-level phenomena except under stringent assumptions (Epskamp, Waldorp, Möttus, & Borsboom, 2018; Hamaker, 2012; Molenaar, 2004; Molenaar & Campbell, 2009). This is because cross-sectional correlations contain a mix of *both* between-person variance *and* within-person variance (Hamaker, 2012).² Whether a cross-sectional correlation will better represent the between-person association or the within-person association depends on the amount of variance in the constructs is due to within- or between-person variance. Even in cases where there is a small amount of between-person variance, cross-sectional correlations will not necessarily provide an accurate estimate of the average within-person association in the sample. This shortcoming is compounded when people have more idiosyncratic within-person associations. In short, cross-sectional correlations – the basis for nearly all discussions about the structure of belief systems – cannot accurately estimate belief systems for individuals (see also Brandt & In press; Martin, 2000).

² This is similar to the emphasis in some of the multilevel literature on the different associations at Level 1 (e.g., within-persons) and Level 2 (e.g., between-persons).

Existing Approaches for Measuring Within-Person Psychological Systems

Some methods have been proposed to measure within-person psychological systems for other psychological constructs. For example, within-person networks of psychopathological symptoms and personality traits have been measured using intensive longitudinal studies with about 50 observations per person (e.g., Beck & Jackson, 2020). These methods are resource intensive. More perplexingly, methodological work shows that they need to be *even more* resource intensive to work properly (Mansueto, Wiers, van Weert, Schouten, & Epskamp, 2020). Even with extensive resources, these models are unlikely to be useful for scholars of belief systems because many of the attitudes and identities that make up political belief systems have high temporal stability (Brandt & In press; Kiley & Vaisey, 2020). This makes it difficult to use methods that rely on within-person variation over time to estimate model parameters (e.g., as is the case with Beck & Jackson, 2020). Although the work on attitude stability also shows some individual-level changes over time (cf. Brandt & In press; Kiley & Vaisey, 2020), the amount of variance is not enough to accurately model the structure of a within-person belief system. Said another way, intensive longitudinal design rely on variance over time to make accurate estimates and the stability of political attitudes and identities means that these methods are unlikely to be appropriate for assessing belief systems. This suggests that the typical approach for measuring within-person psychological systems is likely not a solution for political belief systems. What are researchers to do?

Conceptually Similarity Judgments

Here I propose that conceptual similarity judgments can help fill the methodological gap. Participants are asked to make judgments about pairs of political attitudes and the extent those political attitudes are related. A similar method has been used to assess the conceptual similarity of traits (Stolier, Hehman, & Freeman, 2020; Stolier, Hehman, Keller, Walker, &

Freeman, 2018; Xie, Flake, Stoler, Freeman, & Hehman, in press) and mental models of emotion transitions (Thornton & Tamir, 2017).³ In these studies of traits, the conceptual similarity of traits as rated by participants was associated with the correlation between those traits in other domains (e.g., when rating faces). This work shows that conceptual knowledge about trait associations manifests in perceptions of traits in other domains.

I borrow this conceptual similarity technique and apply it to the political domain. The conceptual similarity task is summarized in Figure 1. In my case, participants are asked to make conceptual similarity judgments about attitudes and identities that are (potentially) in the political belief system. This will give me information for how individual participants think about the relationships between the attitudes and identities of the belief system and allow me to estimate the belief system structure for individuals. Across trials, participants report how likely it is to support both attitudes. The responses to each trial represent a link between attitudes in the belief system. The order of the items can be counterbalanced, so that participants respond to both the stimuli_x–stimuli_y pairing as well as the stimuli_y–stimuli_x pairing. In such a case, the response for both orders can be averaged to represent a single link between attitudes in the belief system (and be used to estimate within-task reliability).

The links in the belief system can be either averaged across participants to create a sample average belief system, or be treated as indicators of a person-specific belief system. Both ways of treating the data can be useful depending on the research question. For example, the sample average belief system can be correlated with other sample level variables (e.g., country of the sample, correlations between other variables in that sample). The person-specific belief systems, however, are best suited for filling the methodological gap identified earlier. Researchers can correlate data from these estimates with other

³ A related and concurrently developed method has been used to assess how people think about political parties working together in a multi-party political system (Ertan, Çarkoğlu, & Aytac, 2022). This approach, however, has not been validated by testing it against theoretical predictions.

individual differences (e.g., partisanship, political engagement) or outcomes of belief system dynamics (e.g., ambivalence). They can also be used as dependent variables in experiments when, for example, researchers are interested in how a manipulation affects the structure of people's belief systems.

An important open question is the best way to word the questions in the similarity judgment task. Figure 1 shows the conceptual similarity judgment task when it is worded to focus on other people. This builds on the intent of work in other subfields, such as stereotyping, which assumes that people may be more willing to divulge their opinions when asked about judgments about someone else. In the stereotyping literature, this might be the stereotypes held by American society, for example (Fiske et al., 2002). I test this other-focused version. In addition, I test a version where I ask people to make these judgments about themselves. This wording is a more face valid measurement of a person's belief system. Indeed, it asks about the relationship between their own attitudes. I have no a priori reason to prefer one version over the other and so test both.

Figure 1. Summary of conceptual similarity method for estimating person-specific belief systems

Benefits

There are several benefits to this approach. First, it is substantially less resource intensive than other methods, like intensive longitudinal studies, that have been used to estimate networks of other psychological constructs. Second, it can flexibly handle different attitudes and identities; it is just a matter of adding them to the measure. Third, it does not require there to be observable variation over time and instead can be estimated at just one time point. Fourth, similar methods have been successfully used. This includes their use to assess the conceptual structure of traits (Stolier et al., 2020). Another study has applied a similar method to understand how people think about political parties working together in a multi-party political system (Ertan et al., 2022). Fifth, and most importantly, it has the potential to estimate the structure of individual belief systems without relying on problematic cross-sectional correlations.

Limitations

It is also important to be upfront about limitations of this approach. First, it is a self-report method and so requires some degree of insight. Although I share the skepticism that people have deep insight into the causes of their attitudes, the conceptual similarity task does not require such deep insight and instead only requires that people are able to approximately assess the likelihood the two attitudes or identities are likely to cooccur. This is a simpler task than identifying specific causal mechanisms. Second, conceptual similarity judgments do not capture causal direction and instead, at best, capture something akin to a correlation between two attitudes or identities. This may be sufficient at this moment in time as the theoretical models of belief systems also do not specify causal direction (see General Discussion for more discussion of this issue).

Like any methodological choice, there are trade-offs. Given the inability for between-person correlations to assess within-person belief systems and the challenges of intensive

longitudinal designs for assessing belief systems (e.g., resource intensive, lack of necessary variation), conceptual similarity judgments may be the preferable tool for assessing belief systems within-persons.

Validity

When developing a new method a key step is assessing its validity. Although this process can go by several names and can assess multiple types of validity, the general premise of this step is to test if a new method maps onto theoretically expected results. For example, when developing a new measure of a motivation, it's validity is determined – in part – by its association with a behavior that would be theoretically caused by the motivations (e.g., Gross & John, 2003; Price, Ottati, Wilson, & Kim, 2015). Drew and Rosenthal (2003, p. 609) summarize it by saying, “A statement about the validity of an instrument is a statement about the extent to which its observed associations with measures of other variables match theoretical predictions about how it should be associated with those variables.”

To test the validity of the conceptual similarity judgment method of estimating individual belief system structure, I will compare data collected with this method with the theoretical predictions of the two computational models of belief systems mentioned above, the Brandt-Sleegers and Goldberg-Stein models of belief systems (Brandt & Sleegers, 2021; Goldberg & Stein, 2018). Although this strategy might deviate from prototypical validity studies in social and personality psychology, the underlying logic is precisely the same. I am testing how a new measure is associated with other variables and the extent those variables match theoretical predictions from precisely specified theoretical models.

The Brandt-Sleegers model of belief system dynamics focuses on how attitudes and identities that are part of the belief systems mutually influence each other (Brandt & Sleegers, 2021). These mutual influences, combined with the strength and the direction of the associations between the attitudes and identities for a person are thought to cause a number of

phenomena in the belief systems literature. I will test one key result from this model's simulations in Studies 1-3. Brandt-Sleegers model's simulations show that (simulated) people with belief systems with stronger connections between attitudes and identities have more consistent (e.g., all left-wing) positions on the attitudes and identities in the belief system than (simulated) people with belief systems with weaker connections between the attitudes and identities. Brandt and Sleegers (2021) conceptually link this result with prior work suggesting that people who are politically engaged and more politically knowledgeable have stronger correlations between attitudes and identities than people who are not politically engaged (e.g., Baldassarri & Goldberg, 2014; Boonen et al., 2017; Converse, 1964; Kinder & Kalmoe, 2017; Lupton et al., 2015).

I test this idea with conceptual similarity judgments. I test if the perception that attitudes and identities with a consistent ideological direction (e.g., liberal attitude with liberal attitude) are more similar is correlated with more political engagement, political knowledge, partisan extremity, and ideological extremity. Conversely, I test if the perception that attitudes and identities with an inconsistent ideological direction (e.g., liberal attitude with conservative attitude) are more dissimilar is correlated with more political engagement, political knowledge, partisan extremity, and ideological extremity. As a more direct test of the Brandt-Sleegers results, I similarly test if the perception that attitudes and identities with a consistent ideological direction are more similar and an inconsistent ideological direction are more dissimilar is correlated with whether people hold more consistently liberal or conservative attitudes themselves. These findings would be consistent with the Brandt-Sleegers model prediction that the political attitudes and identities that people hold are determined, in part, by the structure of their belief systems.

The Goldberg-Stein model of belief systems focuses on how exposure to attitudes in the environment (e.g., the media, via friends) affects how attitudes and identities are

connected (Goldberg & Stein, 2018). For example, by observing that attitude_x and attitude_y are often held by the same people, a person may infer that there is a strong positive (i.e. reinforcing) link between the two attitudes. The person then may update their own position on these attitudes due to the newly created strong positive link between the two attitudes (e.g., changing attitude_x to maximize constraint satisfaction in the belief system). In contrast, by observing that it is difficult to predict if a person will hold attitude_y based on attitude_z, a person may infer that there is no link between these attitudes. The person then may update their own position on these attitudes due to the removal of the link between the two attitudes. The Goldberg-Stein model examines these dynamics and then, in turn, how those dynamics affect the distribution of attitudes in the population. I will test one key result from this model's simulations in Studies 1-3.

One empirical implication of Goldberg-Stein's model – further confirmed in their simulations – is that people who have similar belief system structures (i.e. a similar set of connections between attitudes in the belief system) are more likely to have a similar pattern to the attitudes that they hold. This is an important finding for their model because it shows how people can be exposed to similar information and yet still become polarized. That idea is that although people with similar belief system structures may disagree, there may be similarities in how different attitudes are thought to go together. For example, an extreme liberal and an extreme conservative may share a belief system structure with strong links between the various attitudes and identities (i.e. similar conceptual similarity judgments) and they may vigorously disagree with each other (i.e. different preferences). They may agree that pro-choice attitudes go with high taxes on corporations attitudes and that anti-immigrant attitudes go with anti-trans attitudes. At the same time, they are likely to disagree on all of these attitudes. This shared patterning is called preference congruence. The idea is that the belief system structure causes people to adopt particular patterns of attitudes and so if two

people have similar belief system structure they should also exhibit preference congruence. That is, they should also exhibit a similar pattern of attitudes. This pattern has been found in a simulation study but not in empirical data.

I test the idea that people with more similar belief systems as assessed with the conceptual similarity judgments will have more similar preference congruence on a measure of their support for the attitudes and identities. I will also test if this association between belief system similarity and preference congruence holds when controlling for preference similarity (i.e. agreement on the attitudes).

A key assumption of the Goldberg-Stein model is that when people observe attitudes or identities together in the environment, they update the structure of their belief system (Goldberg & Stein, 2018). Consistent with this, results from the Brandt-Sleegers model's simulations suggest that a mechanism behind partisan cueing experiments is that partisan cues create a link in the belief system between a (typically) novel attitude and a person's partisan identity. In partisan cuing experiments (e.g., Bakker et al., 2020; Brader et al., 2020; Bullock, 2011; Cohen, 2003; Kam, 2005), participants are randomly assigned to read text linking an attitude with their inparty or their outparty, with the typical finding that people are more likely to support the attitude following the inparty vs. the outparty cue. Both the Brandt-Sleegers and Goldberg-Stein model predict that such changes are due to changes in the belief system structure.

In Studies 4 and 5, I test this implication of both models using partisan cuing experiments. I test if partisan cues cause a change in conceptual similarity judgments and then I test if these changes in conceptual similarity judgments account for the effects of the experimental conditions on attitude change.

In short, across five studies I aim to test the validity of conceptual similarity judgments as measures of a person's belief system. I do this by comparing the predictions of

two different theoretical models of belief systems to the empirical results from my studies.

All of the studies are preregistered and non-preregistered analyses are noted.⁴ The supplementary materials contains additional preregistered and exploratory analyses.

Studies 1-3

The shared preregistered purpose of Studies 1-3 is threefold. First, as a basic test to assess if conceptual similarity judgments are working as they have in prior work, I aim to test if aggregated political conceptual similarity judgments are associated with the observed correlations between political attitudes in the sample. This is conceptually similar to prior work on conceptual similarity judgments that has found that conceptual similarity judgments of trait pairs is correlated with the correlations between trait pairs (Stolier et al., 2020). This will help me understand if conceptual similarity ratings for belief systems work like past implementations of conceptual similarity judgments and give me insight into whether the judgments are tracking similarity in the world.

Second, as a test of the predictions of the Brandt-Sleegers model, I aim to test if more extreme political identities, political engagement, political knowledge (Study 3 only), or attitude consistency will have stronger conceptual associations for attitudes with a consistent ideological direction (e.g., liberal-liberal combinations) vs attitudes with an inconsistent ideological direction (e.g., liberal-conservative combinations). For short, I call these “ideologically consistent attitudes” and “ideologically inconsistent attitudes”. This will help me understand if conceptual similarity ratings can be used to conceptually replicate past work

⁴ Study 1: https://osf.io/34xqs/?view_only=74bff229ee1741c79b12cb1881aa922f. Study 2: https://osf.io/7r65h/?view_only=1f51c910dbab448aa8548a85c1de1ca3. Study 3: https://osf.io/zq2yu/?view_only=35a982e8873d4b729cc2bad239f054ce. Study 4: https://osf.io/g3u9t/?view_only=0e6c311b4af949758b18c23553988e2b. Study 5: https://osf.io/kqj8p/?view_only=98df483de0f74738b0ea381232843103. Note that Study 2 was originally preregistered as Study 3. The original Study 2 tried a different question wording that had a fundamental methodological flaw (higher scores could mean either highly similar or highly dissimilar) and is not reported for that reason. Study 3 was conducted after all the other studies

suggesting people with more constrained and interconnected belief systems are more involved in politics and have more consistent attitudes.

Third, I aim to test if people who share similar political conceptual associations are likely to share a similar pattern in their support of political attitudes (preference congruence). This will allow me to empirically test a pattern of results from the Goldberg-Stein model.

In addition, the studies will allow me to see if different types of conceptual similarity judgments influence the results. Participants will be randomly assigned to complete conceptual similarity judgments for other people or for themselves. Each version of the task has its own trade-offs. The *own-task* directly assesses the participants own conceptual associations; however, to ask about these associations it is necessary for them to imagine that they support one of the policies; this may be cognitively taxing and prone to error. The *other-task* does not directly assess the participants own associations, but it may be easier to answer because it only require reporting what the participants observes in their life. Because of these trade-offs, I use both versions of the task in all of the studies reported in this paper and test for differences between the tasks.

An additional goal that Study 3 helps me address is test-retest reliability. This study included a second wave collected one week after the first wave. The test-retest reliability of the conceptual similarity measure is benchmarked to the test-retest reliability of other measures of political attitudes, identities, and engagement included in both waves.

Method

Studies 1-3 shared a design and so are described simultaneously, with differences highlighted. The key differences are the political attitudes used in each study and Study 3 included two waves of data.

Ethics Statement

Studies 1, 2, 4, and 5 were approved by the Ethics Review Board at XXXX University (“Individual Belief Systems”, RP54). Study 3 was determined exempt by the Institutional Review Board at YYYYYY University (“Measuring the belief system of a person”, STUDY00006894).

Participants

Demographic information is in Table 1. Sample sizes were determined based on funding considerations. For Studies 1 and 2, the task was opened for 400 people with approval ratings 95 or higher, who were born in the US, live in the US, and have US nationality. Participants in Study 2 could not have participated in Study 1. Because of the verbal nature of the task, I further restricted my analyses to people who reported that English was their primary language (asked at the end of the study). For Study 3, the task was opened for 600 participants at Time 1 with approval ratings 95 or higher, who were born in the US, live in the US, have US nationality, and who reported being fluent in English (according to Prolific’s prescreening tool). Participants who participated in studies from my lab using similar measures were restricted from participating in Study 3.⁵ Studies 1 and 2 had 98% power to detect $r = .20$. Study 3 had 99% power to detect $r = .20$.

Table 1. *Demographic information for each study.*

	Study 1	Study 2	Study 3 Time 1	Study 3 Time 2
<i>N</i>	387	389	598	536
Response Rate				89.7%
Own condition <i>n</i>	197	206	302	266
Other condition <i>n</i>	190	183	296	268
<i>M</i> age	35.2	36.2	37.9	38.5
<i>SD</i> age	12.6	13.5	13.9	14.1
Gender				
Female	—	—	291	252
Male	—	—	294	272

⁵ Because I moved institutions and changed Prolific accounts between the completion of all of the studies and Study 3, I did not have the information restrict participants from all of the other studies. However, Study 3 was conducted more than 1 year after the other studies reported here and so a large degree of overlap in participants is unlikely.

Sex & Gender Diverse	—	—	11	10
Missing	—	—	2	2
Nationality	United States	United States		United States
Payment	£4.00	£2.20	\$2.42	\$2.20

Procedure and Measures

Participants were randomly assigned to complete the conceptual similarity measures or the measures of their own attitudes first. For the measures of their own attitudes, participants reported whether they supported or opposed 12 different political positions (see attitudes listed in Figures 2-4) on a scale ranging from 1 (strongly oppose) to 7 (strongly support). The attitudes in Study 1 were chosen to cover a number of salient attitudes in American politics and to have items with a liberal (e.g., “taxes on the rich”) or a conservative (e.g., “military spending”) direction. I aimed to choose attitudes that covered a breadth of possible topics, including both economic and cultural attitudes, two key attitude dimensions in American politics (Johnston & Ollerenshaw, 2020).

In Study 2, I replaced seven of the items used in Study 1 with seven political items that were less clearly partisan (e.g., social security, legalizing marijuana).⁶ I identified these attitudes by finding attitudes with small partisan and ideological differences in contemporary polling data.

In Study 3, I replaced all but two items in an effort to ensure that the findings represent a broad set of current attitudes relevant to American belief systems. To select the new items, I selected attitudes from a database of all of the unique policy attitude questions that were fielded in 2020 and included in the Roper iPoll database (Clifford & Kennedy, 2021). Although this is likely not an exhaustive list of all possible political attitudes, it does represent attitudes that polling experts saw fit to poll. Policy attitude questions were first

⁶ Although there isn’t consensus on topics like social security and legalizing marijuana in the United States, the partisan differences among citizens is not as large as it is for other attitudes.

sorted into topics (e.g., Supreme Court, COVID-19) and then I selected one attitude per topic. To fit the format I use here, the wording of the items were changed and in some cases the items were further updated based on current events.⁷

After recoding the items in each study so that higher scores indicated more conservative attitudes, I combined the items (Study 1 $\alpha = .91$; Study 2 $\alpha = .78$, Study 3 $\alpha = .91$) to create a measure of attitude positions. I then centered this measure of attitude positions at the midpoint and took the absolute value to create a measure of attitude consistency.^{8,9} In Study 3, at Time 2, I again asked about three of the attitudes (renewable energy, war with Iran, increased Supreme Court justices). I formed these three attitudes into a short attitude position scale (Study 3 Time 1 $\alpha = .55$, Study 3 Time 2 $\alpha = .58$) to use as a benchmark for test-retest reliability (see below). After measuring their own attitudes, I then measured political engagement with two items (based on Malka et al., 2014). They read, “How interested are you in politics?” (1 = very uninterested, 7 = very interested) and “How important are politics to you?” (1 = very important, 7 = very important). The items were correlated (Study 1 $r = .86$; Study 2 $r = .86$; Study 3 Time 1 $r = .85$; Study 3 Time 2 $r = .85$) and so were averaged to form a scale.

In Study 3, after measuring political engagement I measured political knowledge. I used 10 multiple choice questions that asked factual questions about current events (e.g., rate of inflation), political power (e.g., who control the U.S. House of Representatives), government procedures (e.g., number of votes necessary to override a presidential veto).

⁷ In particular, the original COVID-19 attitude was out of date.

⁸ I did not preregister to use the attitude items in this way for Studies 1 and 2, however, I realized that it allows me to directly test the prediction that people with more attitude consistency will have fewer cross-pressures in their belief systems (Brandt & Sleegers, 2021). If this is the case, attitude consistency will be positively associated with the conceptual similarity to ideologically consistent attitudes and negatively associated with the conceptual similarity to ideologically inconsistent attitudes. This was preregistered for Study 3.

⁹ Sometimes this type of measure is referred to as a measure of extremism (e.g., Brandt, Evans, & Crawford, 2015), however, it is more accurately described as a measure of consistency (Broockman, 2016).

Response options for each question were presented in a random order. All questions at four response options with the exception of one question asking about the political party who controls the U.S. House of Representatives. To reduce cheating, I used Clifford and Jerit's (2016) recommendations. Specifically, we told participants, "It is important to us that you do NOT use outside sources like the Internet to search for the correct answer. Will you answer the following questions without help from outside sources?" Consistent with Clifford and Jerit, we included responses for people who both said yes and no to this question (yes $n = 583$, no $n = 14$, missing $n = 1$). The items were coded as 1 if the response was correct and 0 otherwise. The items were averaged to indicate the proportion of correct responses ($M = .62$).

Conceptual similarity between the political positions was assessed in all three studies by adopting the procedure Stoller et al (2020) used to measure the conceptual similarity of traits. A summary of the task is in Figure 1. In the other-task, I asked "Given that a person supports one attitude, how likely is it that they support the other?" (1 = not at all likely, 7 = very likely) for each political position pair. In the own-task, I asked "Imagine that you support one attitude, how likely is it that you will support the other?" for each political position pair. Both versions of the task were preceded by additional instructions explaining the task and an example pair (see supplemental materials).¹⁰ In Study 3, participants could be randomly assigned to either condition at both time points.

In Study 1, attitude pairs were rated twice in each task, once in each order (e.g., $\text{stimuli}_x - \text{stimuli}_y$, $\text{stimuli}_y - \text{stimuli}_x$) for a total of 156 total ratings. The two ratings for each stimuli pair were averaged for all participants (not preregistered, $M_r = .76$, $Mdn_r = .86$).

Based on simulations of the data from Study 1, Studies 2 and 3 randomly assigned

¹⁰ After each task in Studies 1 and 2, participants completed items assessing their perceptions of the task. Overall, participants found the task slightly easy, they were certain about their answers, they concentrated on the task, and they took the task seriously. These questions were not asked in Study 3.

participants to complete 90 of the 156 total possible ratings.¹¹ This was done to shorten the study and reduce participant fatigue. Because the data is missing completely at random (by design), it is possible to ignore this missing data. For the cases where participants responded to both orders of a pair (Study 2 M /participant = 25.9; Study 3 Time 1 M /participant = 25.8; Study 3 Time 2 M /participant = 25.2), the two ratings were highly correlated within-participants (Study 2 not preregistered, $M_r = .70$, $Mdn_r = .77$; Study 3 Time 1 preregistered, $M_r = .75$, $Mdn_r = .83$; Study 3 Time 2 preregistered, $M_r = .77$, $Mdn_r = .85$).

The last section of the study assessed demographic information, including age, income, education, ethnicity/race, English language (Studies 1 and 2), state, religious affiliation, and gender (Study 3). Most importantly, I included measures of ideological and partisan identification. Ideological identification was measured with the item, “When it comes to politics, do you think of yourself as a liberal, conservative, moderate, or haven't you thought much about this?” (1 = strongly liberal, 7 = strongly conservative). Participants who selected “don’t know” or “I haven’t thought much about it” response options were recoded to the moderate midpoint (4). I used this to assess ideological identification extremity by “folding over” the measure, such that higher scores indicate strong identification as a liberal or a conservative and lower scores indicate less identification with either ideological label. Partisan identification was measured in a similar way with an item that read, “Do you think of yourself as a Republican, a Democrat, an Independent, or haven’t you thought much about this?” (1 = strongly Democrat, 7 = strongly Republican). Participants who selected “don’t know” or “I haven’t thought much about it” response options were recoded to the independent midpoint (4). I used this to assess partisan identification extremity by “folding

¹¹ In these simulations, I reran the analyses from Study 1 with random subsets of 30, 60, 90, or 120 stimuli pairs (out of 156). Across the different analyses, these simulations suggested that 90 stimuli pairs would be adequate (i.e. parameter estimates with partial data were similar to parameter estimates using full data).

over” the measure, such that higher scores indicate strong identification as a Democrat or a Republican and lower scores indicate less identification with either party.

Time 2 of Study 3 included the conceptual similarity task, the two political engagement items, the measures of ideological and partisan identification, and the single items assessing support for funding more research into renewable energy sources, going to war with Iran, and increasing the number of Supreme Court justices. All of these items were collected to assess and compare test-retest reliability.

Figure 2. Study 1 mean conceptual association in the other condition (Top Left) and the own condition (Top Right). The correlation between each of the political attitudes (Bottom Left).

The relationship between conceptual associations and correlations, with the lines of linear best fit (Bottom Right).

Figure 3. Study 2 mean conceptual association in the other condition (Top Left) and the own condition (Top Right). The correlation between each of the political attitudes (Bottom Left). The relationship between conceptual associations and correlations, with the lines of linear best fit (Bottom Right).

Figure 4. Study 3 mean conceptual association in the other condition (Top Left) and the own condition (Top Right). The correlation between each of the political attitudes (Bottom Left). The relationship between conceptual associations and correlations, with the lines of linear best fit (Bottom Right).

Results

Preregistered Analyses

The analyses are based on representational similarity analyses (RSA; Kriegeskorte et al., 2008), which are useful for testing the correspondence between outputs from different types of analyses (e.g., brain-activity measures and computational models). I follow the approach of Stolier and colleagues (2020) who applied RSA to understand the conceptual structure of traits.

Aggregate Analyses

First, I tested if aggregated political conceptual similarity judgments were associated with the observed correlations between political attitudes in the sample. To do this, I calculated the correlation between each of the political attitudes in the sample to represent political attitude similarity (Figures 2-4, Bottom Left Panel). Then, I calculated the mean conceptual similarity for each item pair in each experimental condition (Figures 2-4, Top Row). Finally, I ranked both the conceptual similarity judgments and the political attitude similarities to compute a Spearman correlation between the two. This helps address possible non-linear relationships as recommended for RSA (Kriegeskorte et al., 2008), although Pearson correlations result in the same conclusions.

I found that in both the own condition (Study 1: $r_s = .87$, 95% CI [.81, .92], $p < .001$; Study 2: $r_s = .91$, 95% CI [.87, .94], $p < .001$; Study 3: $r_s = .89$, 95% CI [.84, .93], $p < .001$) and the other condition (Study 1: $r_s = .88$, 95% CI [.82, .92], $p < .001$; Study 2: $r_s = .94$, 95% CI [.91, .96], $p < .001$; Study 3: $r_s = .92$, 95% CI [.87, .95], $p < .001$) conceptual similarity was related to the political attitudes' similarity (Figures 2-4, Bottom Right Panel). This suggests that conceptual similarity ratings are tracking similarity in the world and conceptually replicate results of conceptual similarity analyses in other domains (Stolier et

al., 2020).¹² Additional analyses using permutation tests showed that the correlations in both conditions were substantially greater than would be expected by chance. A linear regression analysis using ranked conceptual similarity, dummy coded condition, and their interaction to predict ranked political attitude similarity showed that the two conditions did not differ from one another (Study 1: interaction $p = .93$; Study 2: interaction $p = .66$; Study 3: interaction $p = .74$). This is also evident in Figures 2-4 (Bottom Right Panel) where the lines for linear fit overlap substantially.

Individual Differences

Second, I tested if more extreme political identities, political engagement, or political knowledge (Study 3 only) will have stronger conceptual associations for ideologically consistent attitudes (e.g., liberal-liberal combinations) vs ideologically inconsistent attitudes (e.g., liberal-conservative combinations). I also tested if attitude positions (whether a person holds relatively liberal or conservative attitude positions) and attitude consistency (whether a person holds consistently liberal or conservative attitude positions) will have stronger conceptual associations for ideologically consistent attitudes vs ideologically inconsistent attitudes (not preregistered for Studies 1 and 2).

The first step in this analysis was to categorize each of the item pairs in the conceptual association task as ideologically consistent and ideologically inconsistent. When an item that is a typical liberal (e.g., “gay rights”) or conservative (“severe criminal sentences”) position was paired with another typical liberal (e.g., “environmental protection”) or conservative (“restrictions on immigration”) position it was coded as ideologically consistent. When a typical liberal or conservative position was paired with a typical

¹² An alternative method for testing this same idea that preserves individual level variation is presented in the supplemental materials.

conservative or liberal position it was coded as ideologically inconsistent.¹³ Then, I estimated five multilevel models where conceptual association ratings were nested within participant and both item pairs and participant were treated as random effects. In Model 1, I regressed each person's conceptual association ratings on political engagement (mean centered), experimental condition (own vs. other, contrast coded), and pair type (i.e. ideologically consistent vs. ideologically inconsistent, contrast coded). Pair type was included as a random slope. Model 2 was like Model 1, except that I removed political engagement and included political knowledge (mean centered) (this was only estimated for Study 3). Model 3 was similar to Model 1, except that I removed political engagement and instead included midpoint-centered ideological identification and mean-centered party extremity. Model 4 was just like Model 3, except that I swapped out the ideology-related measures for party identification and party extremity. Model 5 was just like Models 3 and 4, except that I used midpoint-centered attitude positions and mean-centered attitude consistency.

The full results of each model are in Tables S1 to S13 in the supplemental materials. The data and fit lines from these models are plotted in Figures 5-7. By design, across all models the item pairs that were ideologically consistent had higher conceptual similarity than did item pairs that were ideologically inconsistent.

Political Engagement. Model 1 showed that the relationship between political engagement and conceptual similarity depended on the item type (Study 1: interaction $b = .26$, $SE = .04$, $p < .001$; Study 2: interaction $b = .17$, $SE = .03$, $p < .001$; Study 3: interaction $b = .30$, $SE = .03$, $p < .001$), such that political engagement was associated with more conceptual similarity for ideologically consistent pairs (Study 1: $b = .14$, $SE = .02$, $p < .001$; Study 2: $b = .13$, $SE = .02$, $p < .001$; Study 3: $b = .16$, $SE = .02$, $p < .001$) and less conceptual

¹³ An item was considered typically liberal if contemporary polling data showed that liberals or Democrats were more supportive of the attitude than conservatives or Republicans (vice versa for determining if an item was typically conservative).

similarity for ideologically inconsistent pairs (Study 1: $b = -.12$, $SE = .03$, $p < .001$; Study 2: $b = -.04$, $SE = .03$, $p = .12$; Study 3: $b = -.15$, $SE = .02$, $p < .001$; Figures 5-7). This last effect was not significant in Study 2. In other words, people higher in political engagement had more strongly connected belief systems than people lower in political engagement. This was expected.

Political Knowledge. Model 2 (Study 3 only) showed that the relationship between political knowledge and conceptual similarity depended on the item type (Study 3: interaction $b = 2.14$, $SE = .32$, $p < .001$), such that political knowledge was associated with more conceptual similarity for ideologically consistent pairs (Study 3: $b = 1.21$, $SE = .21$, $p < .001$) and less conceptual similarity for ideologically inconsistent pairs (Study 3: $b = -.93$, $SE = .19$, $p < .001$; Figure 7). In other words, people higher in political knowledge had more strongly connected belief systems than people lower in political knowledge. This was expected.

Ideological Identity. Model 3 showed that the relationship between ideological extremity and conceptual similarity depended on the item type in Studies 1 and 3 (Study 1: interaction $b = .37$, $SE = .07$, $p < .001$; Study 3: interaction $b = .25$, $SE = .06$, $p < .001$), such that ideological extremity (vs. moderation) was associated with more conceptual similarity for ideologically consistent pairs (Study 1: $b = .17$, $SE = .04$, $p < .001$; Study 3: $b = .16$, $SE = .04$, $p < .001$) and less conceptual similarity for ideologically inconsistent pairs (Study 1: $b = -.19$, $SE = .04$, $p < .001$; Study 3: $b = -.09$, $SE = .03$, $p = .007$; Figures 5 & 7). In other words, people higher in ideological extremity had more strongly connected belief systems than people lower in ideological extremity. This was not replicated in Study 2. In Study 2, ideological extremity was positively related to conceptual similarity (Study 2: $b = .09$, $SE = .03$, $p = .007$) and this was not moderated by item type (Study 2: interaction $b = .08$, $SE = .06$, $p = .18$; Figure 6).

In all three studies there was an unexpected interaction between item type and ideological identification (Study 1: interaction $b = -.23$, $SE = .05$, $p < .001$; Study 2: interaction $b = -.26$, $SE = .03$, $p < .001$; Study 3: interaction $b = -.24$, $SE = .03$, $p < .001$), such that conservative political identification (vs. liberal identification) was associated with less conceptual similarity for ideologically consistent pairs (Study 1: $b = -.10$, $SE = .03$, $p < .001$; Study 2: $b = -.09$, $SE = .02$, $p < .001$; Study 3: $b = -.11$, $SE = .02$, $p < .001$) and more conceptual similarity for ideologically inconsistent pairs (Study 1: $b = .13$, $SE = .03$, $p < .001$; Study 2: $b = .17$, $SE = .03$, $p < .001$; Study 3: $b = .13$, $SE = .02$, $p < .001$; Figures 5-7). In other words, conservatives tended to see item pairs thought to go together as less similar than did liberals. Conversely, conservatives tended to see item pairs thought to not go together as more similar than did liberals. This was unexpected, but might be a result of conservatives' tendency to endorse policies that are inconsistent with their ideological identity (i.e. people who identify as conservative often endorse liberal policies, Ellis & Stimson, 2012).

Partisan Identity. Model 4 showed conceptually similar results to Model 3. It showed that the relationship between partisan extremity and conceptual similarity depended on the item type in Studies 1 & 3 (Study 1: interaction $b = .21$, $SE = .08$, $p = .007$; Study 3: interaction $b = .14$, $SE = .06$, $p = .01$), such that partisan extremity (vs. neutrality) was associated with more conceptual similarity for ideologically consistent pairs (Study 1: $b = .09$, $SE = .04$, $p = .04$; Study 3: $b = .10$, $SE = .04$, $p = .01$) and less conceptual similarity for ideologically inconsistent pairs (Study 1: $b = -.12$, $SE = .05$, $p = .007$; Study 3: $b = -.04$, $SE = .03$, $p = .19$; Figures 5 & 7). Although, this simple slope was not significant in Study 3. In other words, people higher in partisan extremity had more strongly connected belief systems than people lower in partisan extremity. This was not replicated in Study 2. In Study 2, partisan extremity was not moderated by item type (Study 2: interaction $b = .08$, $SE = .05$, $p = .14$; Figure 6).

There was also an unexpected was the interaction between item type and party identification (Study 1: interaction $b = -.22$, $SE = .05$, $p < .001$; Study 2: interaction $b = -.25$, $SE = .04$, $p < .001$; Study 3: interaction $b = -.29$, $SE = .03$, $p < .001$), such that Republican political identification (vs. Democratic identification) was associated with less conceptual similarity for ideologically consistent pairs (Study 1: $b = -.11$, $SE = .03$, $p < .001$; Study 2: $b = -.10$, $SE = .03$, $p < .001$; Study 3: $b = -.16$, $SE = .02$, $p < .001$) and more conceptual similarity for ideologically inconsistent pairs (Study 1: $b = .10$, $SE = .03$, $p < .001$; Study 2: $b = .15$, $SE = .03$, $p < .001$; Study 3: $b = .14$, $SE = .02$, $p < .001$; Figures 5-7). In other words, Republicans tended to see item pairs thought to go together as less similar than did Democrats. Conversely, Republicans tended to see item pairs thought to not go together as more similar than did Democrats. This was unexpected.

Attitude Positions. Model 5 showed that the relationship between attitude consistency and conceptual similarity depended on the item type (Study 1: interaction $b = 1.00$, $SE = .11$, $p < .001$; Study 2: interaction $b = .96$, $SE = .18$, $p < .001$; Study 3: interaction $b = .98$, $SE = .08$, $p < .001$). This suggests that people with more attitude consistency had more conceptual similarity for ideologically consistent pairs (Study 1: $b = .50$, $SE = .06$, $p < .001$; Study 2: $b = .42$, $SE = .14$, $p = .003$; Study 3: $b = .38$, $SE = .05$, $p < .001$) and less conceptual similarity for ideologically inconsistent pairs (Study 1: $b = -.49$, $SE = .07$, $p < .001$; Study 2: $b = -.54$, $SE = .16$, $p < .001$; Study 3: $b = -.60$, $SE = .05$, $p < .001$; Figures 5-7). In other words, people higher in attitude consistency had more strongly connected belief systems than people lower in attitude consistency. This was expected. In contrast to Models 2 and 3 whereby ideology and partisanship unexpectedly interacted with item type, there were no significant two-way interactions between people's positions on the attitudes and item type.

Qualifying Interactions. I did not find any significant interactions that qualified the main conclusions described above. There were other significant coefficients (see

supplemental tables). Most importantly, there is a three-way interaction in Study 3 between engagement, own vs. other experimental condition, and item type ($p = .03$; first column, Figure 7). This interaction did not emerge in any other study and it did not hold when controlling for political knowledge (see below). Therefore, I did not interpret it further. There is also a three-way interaction in Study 3 between people's attitude positions, own vs. other experimental condition, and attitude type ($p = .001$; last column, Figure 7). This interaction did not emerge in any other study; however, it is a similar pattern to the same non-significant interaction in Study 2. The pattern suggests that people with more conservative attitude positions had less conceptual similarity for ideologically consistent pairs in the own condition, but more conceptual similarity in the other condition. Conversely, in the other condition people with more conservative attitude positions appeared to see more conceptual similarity for ideological consistent pairs. Interestingly, a similar pattern did not emerge for partisan or ideological identification, suggesting that if this pattern is robust it is specific to people's attitude positions.

Political Knowledge as Control. One objection to the conceptual similarity measure of belief systems is that it is just a proxy for political knowledge. Although political knowledge was associated with conceptual similarity, these associations were not so large as to suggest they are measuring the same construct.¹⁴ Moreover, I re-ran (not preregistered) the engagement, ideological identification, party identification, and attitude models in Study 3 again controlling for political knowledge and political knowledge's interaction with both experimental condition and attitude type (Tables S14-S17). The key two-way interaction between attitude type and engagement, ideological extremity, party extremity, or attitude consistency remained significant when controlling for political knowledge. Similarly, the

¹⁴ When collapsing the consistent conceptual similarity judgments and the inconsistent conceptual similarity judgments separately for each condition, the correlation with political knowledge was always $r < |.19|$.

two-way interaction between attitude type and political knowledge remained significant in these models. This suggests that the measure of conceptual similarity does not reduce to a measure of political knowledge.

Summary. Across these five models we see that conceptual similarity ratings map onto expectations from the literature. People who are more politically engaged, are more knowledgeable about politics, are more committed to their political identities, or have more consistent political attitudes have stronger associations between ideologically consistent attitudes than ideologically inconsistent attitudes. This shows that conceptual similarity ratings can be used to replicate past work suggesting that people who are involved in politics likely belief systems with stronger connections.

Figure 5. Study 1 data and multilevel estimates of the association between individual differences and conceptual similarity for ideologically consistent (liberal-liberal, conservative-conservative) and ideologically inconsistent (liberal-conservative) attitudes.

Figure 6. Study 2 data and multilevel estimates of the association between individual differences and conceptual similarity for ideologically consistent (liberal-liberal, conservative-conservative) and ideologically inconsistent (liberal-conservative) attitudes.

Figure 7. Study 3 data and multilevel estimates of the association between individual differences and conceptual similarity for ideologically consistent (liberal-liberal, conservative-conservative) and ideologically inconsistent (liberal-conservative) attitudes.

Preference Congruence

Third, I tested if people who share similar political conceptual associations are likely to share a similar pattern in their support of political attitudes. This allowed me to empirically test a pattern of results found in the Goldberg-Stein model. I used the same methods to calculate the necessary indices for these tests as did Goldberg and Stein (2018). For each participant in the sample, I calculated the similarity between their conceptual similarity ratings and the conceptual similarity ratings of every other participant in their experimental condition by finding the absolute value of the difference between each conceptual similarity

rating. This was reversed by multiplying it by -1 to create a measure of coherence (larger values indicate more coherence). Then, for each participant in the sample, I calculated their level of agreement with each other participant in their experimental condition by correlating their responses to the political attitude items (positive correlations indicate more agreement). Finally, I used the absolute value of the level of agreement to assess if participants share a similar pattern for understanding the items (larger values indicate more of a shared pattern). This is called preference congruence.¹⁵ According to the Goldberg-Stein model, conceptual coherence (i.e. similarity in conceptual structure) should be correlated with preference congruence (i.e. sharing a similar logic in the endorsement of attitudes).

To test this, I regressed preference congruence on conceptual coherence. To ensure any association is not due to agreement, I also included agreement in the model. Each row of the dataset used to estimate this model is a pair of participants, representing all possible pairs of participants. In the preregistered analyses, I did not account for this nested structure; however, in the process of analyzing Study 3 I realized that a more appropriate model is a multilevel model with two random intercepts, one for each participant in each participant pair. Therefore, the final model is a multilevel model with preference congruence regressed on conceptual coherence and agreement, with random intercepts for each participant in each participant pair and a random slope for conceptual coherence for each participant in each participant pair. A conceptually similar model that uses agreement as the outcome variable and controls for preference congruence was also conducted for comparison purposes. Two additional models also tested if experimental condition was a significant moderator of the association between preference congruence and conceptual coherence.

¹⁵ As a toy example, if Participant A responded to 5 attitude items with responses of [1, 3, 4, 5, 7] and Participant B responded to the same 5 attitude items with the responses of [5, 7, 4, 3, 1] then their level of agreement would be -1 because the two sets of responses are perfectly negatively correlated. At the same time, Participant A and Participant B both share the same logic to their responses (e.g., the first and the last item should have opposite responses) and so their preference congruence is 1.

Consistent with expectations, I found that conceptual coherence across participant pairs was associated with preference congruence (Study 1: $b = .06$, $SE = .006$, $p < .001$; Study 2: $b = .03$, $SE = .004$, $p < .001$; Study 3: $b = .05$, $SE = .004$, $p < .001$) and this was not consistently moderated by experimental condition (Study 1 interaction $p = .01$, Studies 2 & 3 p 's $> .12$). I also found that conceptual coherence across all possible participant pairs was associated with attitude agreement (Study 1: $b = .03$, $SE = .001$, $p = .01$; Study 2: $b = .06$, $SE = .007$, $p < .001$), except in Study 3 (Study 3: $b = .005$, $SE = .007$, $p = .53$) and this was not moderated by experimental condition (all p 's $> .27$). The raw data and the fitted estimates from these models are in Figures 8 to 9. The key is that in both conditions conceptual coherence is positively associated with preference congruence and this is empirical confirmation of the Goldberg-Stein model.

Figure 8. Study 1 association between conceptual coherence, agreement, and preference congruence for the own and other conditions of the task. Each point is the raw data from a pair of participants. Lines of best fit when predicting agreement control for preference congruence. Lines of best fit when predicting preference congruence control for agreement.

Figure 9. Study 2 association between conceptual coherence, agreement, and preference congruence for the own and other conditions of the task. Each point is the raw data from a pair of participants. Lines of best fit when predicting agreement control for preference congruence. Lines of best fit when predicting preference congruence control for agreement.

Figure 10. Study 3 association between conceptual coherence, agreement, and preference congruence for the own and other conditions of the task. Each point is the raw data from a pair of participants. Lines of best fit when predicting agreement control for preference congruence. Lines of best fit when predicting preference congruence control for agreement.

Test-Retest Reliability

Study 3 contained a second wave that I used to estimate test-retest reliability over 1 week. As reported in the Method's section above, all three studies showed good test-retest

reliability within a time point (i.e. Study 1: $M_r = .76$, $Mdn_r = .86$; Study 2: $M_r = .70$, $Mdn_r = .77$; Study 3 Time 1 $M_r = .75$, $Mdn_r = .83$; Study 3 Time 2 $M_r = .77$, $Mdn_r = .85$). The question I answer here is whether they are also stable over the course of a week. To estimate test-retest reliability, for each participant I correlated all of their conceptual similarity judgments at Time 1 with all of their conceptual similarity judgments at Time 2. Overall, I found levels of test-retest reliability similar to the reliability within time points ($M_r = .72$, $Mdn_r = .83$). The test-retest reliability for each participant and the test-retest reliability for the other measures included at both time points are in Figure 11. The median test-retest reliability of the conceptual similarity measure is lower than the test-retest reliability of people's political identities (party identification $r = .95$, ideological identification $r = .96$) or engagement ($r = .9$), but it is similar or higher to test-retest reliability of people's attitude positions (short attitude scale $r = .87$, renewable energy $r = .81$, war with Iran $r = .58$, increase SCOTUS $r = .86$).

Figure 11. Left panel: Participant test-retest reliability for the conceptual similarity judgments for each participant in each condition at Time 1 and Time 2. Right panel: Test-retest reliability for comparison measures. X-axis is ordered by the size of the estimate of

test-retest reliability. The dashed gray horizontal line is the median participant test-retest reliability of the conceptual similarity judgments.

Discussion

In all three studies, I found that aggregated conceptual similarity judgments between attitudes and identities closely corresponds to the correlations between the same attitudes and identities. This conceptually replicates work on conceptual similarity judgments of traits (Stolier et al., 2020) and suggest that the belief system version is working as expected. More importantly, the conceptual similarity judgments were associated with other variables as would be expected from theoretical models of belief systems (Brandt & Sleegers, 2021; Goldberg & Stein, 2018). I found that more extreme political identities,¹⁶ political engagement, political knowledge, and attitude consistency had stronger conceptual associations for ideologically consistent attitudes (e.g., liberal-liberal combinations) vs ideologically inconsistent attitudes (e.g., liberal-conservative combinations). I also found that participant pairs with more conceptual coherence (i.e. similar conceptual similarity judgments) had more preference congruence than those with less conceptual coherence. This empirically validates the Goldberg-Stein simulations. Overall, the results suggest that the conceptual similarity judgments are a viable and valid method for assessing individual belief systems.

A consistent unexpected result was that conservatives and Republicans reported that inconsistent attitude pairs (i.e. those thought to not go together) were more similar than did liberals and Democrats. Similarly, conservatives and Republicans reported that consistent attitude pairs (i.e. those thought to go together) were less similar than did liberals and

¹⁶ In Study 2, political identity extremity did not interact with item type. One possibility is that the less polarized attitudes in Study 2 clouded the relationship. To test this idea, I reran the multilevel models for Study 2 only using the items that Study 2 shared with Study 1. The results of these analyses resulted in similar conclusions for political identity extremity as did the analyses with all of the attitudes in Study 2. At this stage, I believe that the lack of replication may be due to random chance.

Democrats. This was unexpected, but might be a result of conservatives' tendency to endorse policies that are inconsistent with their ideological identity (i.e. people who identify as conservative often endorse liberal policies, Ellis & Stimson, 2012) and Gidron's (2022) observation that people who identify as right-wing tend to endorse a wider variety of possible attitudes. The conceptual similarity measure suggests that these differences may be reflected in the structure of conservative and Republican belief systems in general. That is, the differences observed by these prior scholars may be due to different connections between attitudes in the belief systems of people who identify as conservative and Republican.

Each study used a different set of political attitudes, including more and less polarized attitudes. Notably, the effects were largely consistent across all three studies. Studies 2 and 3 also used an abbreviated version of the task (i.e. a random subset of trials) and the results were consistent with the full version. This suggests that the measures can be used in a more efficient fashion and increases the feasibility of the measure. Taken together, the results suggest that the conceptual similarity judgments are a viable method for assessing individual belief systems.

Studies 4 & 5

The purpose of Study 4 and 5 was to test the theoretical prediction that partisan framing changes the conceptual structure of the belief system (Brandt & Sleegers, 2021). This is also consistent with Goldberg and Stein's (2018) model which predicts that people adjust the extent attitudes are connected by observing how the attitudes are connected in the world. The idea is that attitude change comes about by – in part – the change in the structure of the belief system. I directly test this by replicating a partisan framing experiment and additionally testing if an in-party frame increases the conceptual similarity between the target attitude and in-party support, as well as if an out-party frame decreases the conceptual similarity between the target attitude and out-party support. Changes in the conceptual

structure should then mediate between the partisan frames and the change in attitude. For the purpose of exploration, I also tested how the frames shift the similarity between the targeted attitude and support for other attitudes. This tests the possibility that partisan frames not only shift the conceptual similarity between the targeted attitude and party support, but also the conceptual similarity with other attitudes. In Study 4 I tested this using farm subsidies and in Study 5 I test this using food irradiation, both building on prior studies using these same materials (Bakker et al., 2020; Kam, 2005; Malka & Lelkes, 2010).

Method

Studies 4 and 5 shared a design and so are described simultaneously, with differences highlighted. The key differences are the target attitude used in each study.

Participants

Demographic information is in Table 2. With the goal of having approximately 300 participants per condition, the task was opened for 1000 people with approval ratings 95 or higher, who were born in the US, live in the US, have US nationality, and who did not participate in prior studies (Studies 1 and 2 for Study 4, and Studies 1, 2, and 4 for Study 5). For Study 5 only I tried to recruit an equal number of Republicans and Democrats and reduce the number of necessary exclusions (see below). Therefore, for Study 5, I opened two tasks for a total of 1000 participants. One task was opened to people who were self-reported Republicans and one task was opened to people who were self-reported Democrats according to Prolific's prescreening tool. Because of the verbal nature of the task, I further restricted my analyses to people who reported that English was their primary language (asked at the end of the study).

Table 2. *Demographic and sample size information for Studies 4 and 5.*

	Study 4	Study 5
<i>N</i>	981	983
Excluded <i>n</i>	105	11
<i>Post-exclusions Frame Conditions</i>		

In-group frame condition <i>n</i>	318	326
Out-group frame condition <i>n</i>	261	324
Control condition <i>n</i>	297	322
<i>Post-exclusions Conceptual Similarity Conditions</i>		
Own condition <i>n</i>	442	487
Other condition <i>n</i>	443	485
<i>M</i> age	34.2	34.6
<i>SD</i> age	12.8	13.4
Gender		
Female	—	476
Male	—	501
Something else	—	6
Nationality	United States	United States
Payment	£0.90	£0.90

Procedures and Measures

I built on Bakker and colleagues (Studies 3 and 3 of Bakker et al., 2020). Participants completed a measure of their partisanship using a branching method. They were asked if “Generally speaking, do you usually think of yourself as a Republican, a Democrat, an Independent, or what?”. Participants who selected Republican or Democrat were then asked “Would you call yourself a strong or not very strong [Republican/Democrat]?” Participants who selected Independent or something else were asked “Do you think of yourself as closer to the Republican or Democratic party?” For this last question, they could also select “neither”. I used the first and last questions of the branching items to identify which frame was an in-party or out-party frame.

Next, participants completed the same two-item measure of political engagement as in Studies 1-3 (Study 4 only) and a brief 11-item measure of the need for cognition. Following Bakker et al (2020), I included need for cognition in my analyses by combining the items using a single factor latent variable model and extracting factor scores. Items all highly loaded on the single factor (Study 4: all *B*’s > .73; CFI = .90; Study 5: all *B*’s > .72; CFI = .86). The measure was recoded to range from 0 to 1.

Participants were then randomly assigned to one of three partisan frame conditions that were used in Bakker and colleagues (2020). Study 4 used a frame about a farm subsidy (see also Malka & Lelkes, 2010) and Study 5 used a frame about food irradiation policy (see also Kam, 2005). In all conditions, participants read about either a farm subsidy policy or a food irradiation policy. In the Republican frame condition (Study 4 $n = 305$; Study 5 $n = 319$), participants learned that Republicans supported the policy and Democrats opposed the policy. In the Democratic frame condition (Study 4 $n = 341$; Study 5 $n = 337$), participants learned that Democrats supported the policy and the Republicans opposed it. In the control condition (Study 4 $n = 335$; Study 5 $n = 327$), participants learned that “various groups” supported the policy and that “various other groups” opposed the policy (Study 4) or participants learned opponents and supporters of the plan without any reference to political party (Study 5). For analyses, I recoded the Republican and Democratic frame conditions to represent in-group or out-group frames depending on participants previously reported political identity. Participants who did not report a Republican or a Democratic identity at any stage of the branching question were not included in the analyses (Study 4 $n = 105$ not included; Study 5 $n = 11$ not included). Total sample sizes for in-party, out-party, and control conditions are in Table 2. This resulting sample size gives me 94% (Study 4) or 97% (Study 5) power to detect a difference of $d = .3$ between the two conditions with the fewest participants.

In Study 4, following the frame, all participants answered one item asking, “Do you support or oppose the U.S. government policy of giving money to American farmers?”. Responses ranged from strongly support (1) to strongly oppose (7). This item was reverse scored so that higher numbers indicate more support. It was recoded to range from 0 to 1. The item was used to assess if the partisan frame caused a change in attitude and serves to

replicate the main effect of partisan frame on attitude change found in prior work using this same manipulation.

In Study 5, following the frame, all participants answered three items asking about their position on the food irradiation policy. These items read, “Do you support or oppose a ban on food irradiation?”, “All things considered, food irradiation is a bad thing”, and “The costs of food irradiation outweigh the benefits.” Responses ranged from Strong support for a ban on food irradiation (1) to Strong opposition for a ban on food irradiation (5) or from Strongly agree (1) to Strongly disagree (5). Items were scored so that higher scores indicated more support for the food irradiation ban, averaged ($\alpha = .77$), and recoded to range from 0 to 1.

My other key dependent variable was the conceptual association between participants’ in-party support and farm subsidies or food irradiation, as well as the conceptual association between participants’ out-party support and farm subsidies or food irradiation. I measured this using a similar task as Studies 1 and 2. Participants were randomized into own and other conditions and received the same instructions as in the first three studies. Then they completed conceptual associations for 28 attitude pairs (for list of attitudes, see supplemental materials). Randomly interspersed in this task were the pairs, “the Republican party – government money to American farmers/food irradiation ban” and “the Democratic party – government money to American farmers/food irradiation ban”. These two items were recoded to represent in-party conceptual association and out-party conceptual association using participants previously reported political identity. They were recoded to range from 0 to 1.

I also analyzed a subset of the additional conceptual association measures that assessed the conceptual similarity between the targeted attitudes and 5 other political attitudes. This helps me explore if the frames cause changes in conceptual association

between the explicitly targeted attitude pair. Four attitudes were common across studies (taxes on the rich, restrictions on immigration, gun ownership, and government run healthcare). Study 4 included “COVID-19 quarantines” and Study 4 included “Black Lives Matter protests.”

Lastly, participants completed demographic measures.

Results

Preregistered

Consistent with prior research (e.g., Bakker et al., 2020; Kam, 2005) and the Brandt-Sleegers and Goldberg-Stein models, I expected that people would be more supportive of the targeted attitude (i.e., farm subsidies, food irradiation policies) following an in-party frame than following an out-party frame. To test this, I regressed support for the targeted attitude on dummy codes indicating in-party frame condition and out-party frame condition, party strength, need for cognition, age, race (1 = non-White, 0 = White), education (three categories, dummy coded; high school education or less, some college or bachelors degree, post-graduate degree), and Republican party identification (1 = Republican identification, 0 = Democratic identification). This was an identical model to that used in Bakker et al. (2020) and full results are in Table S 19 and Table S 25. Most importantly, the in-party cue boosted support (Study 4: $b = .07$, $SE = .02$, $p = .002$; Study 5: $b = .05$, $SE = .02$, $p = .03$) and the out-party cue reduced support (Study 4: $b = -.08$, $SE = .03$, $p = .001$; Study 5: $b = -.07$, $SE = .02$, $p = .002$) compared to the control condition. See Figures 12 and 13 (Left Panels) for the unadjusted means, confidence intervals, and raw data for each condition. See supplemental materials for additional analyses testing if need for cognition or partisan strength are moderators and models that do not adjust for demographic factors and partisanship; results are consistent across all models and conditions. This replicates past work and shows that the framing conditions worked as expected.

Figure 12. Study 4 unadjusted means, confidence intervals, and data points for the effect of cues on policy support, in-party conceptual similarity, and out-party conceptual similarity.

Figure 13. Study 5 unadjusted means, confidence intervals, and data points for the effect of cues on policy support, in-party conceptual similarity, and out-party conceptual similarity.

Next, I tested if the experimental conditions affected in-party and out-party conceptual similarity. I used the same model as in the prior paragraph, but swapped out support for the targeted attitude for the measures of in-party and out-party conceptual similarity (see Table S 20, Table S 21, Table S 26, and Table S 27 for full models). Consistent with expectations, the in-party cue boosted in-party similarity (Study 4: $b = .11$, $SE = .02$, $p < .001$; Study 5: $b = .19$, $SE = .03$, $p < .001$) and the out-party cue reduced in-party similarity (Study 4: $b = -.17$, $SE = .03$, $p < .001$; Study 5: $b = -.15$, $SE = .03$, $p < .001$)

compared to the control condition. Similarly, the in-party cue reduced out-party similarity (Study 4: $b = -.21$, $SE = .03$, $p < .001$; Study 5: $b = -.17$, $SE = .03$, $p < .001$) and the out-party cue boosted out-party similarity (Study 4: $b = .15$, $SE = .03$, $p < .001$; Study 5: $b = .19$, $SE = .03$, $p < .001$). See Figures 8 and 9 (Center and Right Panels) for means, confidence intervals, and raw data for each condition. See supplemental materials for additional analyses testing if need for cognition or partisan strength are moderators, models that do not adjust for demographic factors and partisanship, and models testing for differences between the own and other conditions; results are consistent across all models and conditions. Overall, this shows that partisan frames affect the conceptual structure of people's belief systems.

Although only formally preregistered for Study 5, the two theoretical models I am using to inform my predictions both suggest that sufficient changes to the conceptual structure will be associated with changes in the expressed attitude. To assess if conceptual similarity was associated with support for the targeted attitude, I regressed support for the targeted attitude on the condition dummy codes, in-party similarity, out-party similarity, and the control variables. The two condition dummy codes were no longer significant predictors of support for the targeted attitude (Study 4: in-party frame $b = .03$, $SE = .02$, $p = .21$; Study 5: in-party frame $b = .005$, $SE = .02$, $p = .84$; Study 4: out-party frame $b = -.01$, $SE = .02$, $p = .58$; Study 5: out-party frame $b = -.04$, $SE = .02$, $p = .12$), whereas in-party similarity was a strong predictor support for the targeted attitude (Study 4: $b = .39$, $SE = .04$, $p < .001$; Study 5: $b = .24$, $SE = .03$, $p < .001$). Out-party similarity did not have an effect (Study 4: $b = -.02$, $SE = .04$, $p = .65$; Study 5: $b = -.01$, $SE = .03$, $p = .83$).¹⁷

¹⁷ In exploratory analyses, I also tested if this differed by own vs. other experimental condition. In Study 4, own vs. other condition did not moderate the effect of in-party similarity (Study 4: $b = .12$, $SE = .08$, $p = .15$), but it did moderate the effect of out-party similarity at the $p < .05$ level (Study 4: $b = .15$, $SE = .07$, $p = .03$). Probing this interaction showed that the effect of out-party similarity on farm subsidies was non-significant in the own condition (Study 4: $b = .05$, $SE = .05$, $p = .34$) and it was negative and marginally significant in the other condition (Study 4: $b = -.11$, $SE = .06$, $p = .052$). In Study 5, there was an interaction between condition and out-party similarity ($b = .13$, $SE = .06$, $p = .03$). In the other condition, out-party similarity was non-significantly

Exploratory Mediation

I found that the experimental conditions affected support for the targeted attitude, that the experimental conditions affected in-party similarity, and that in-party similarity is associated with support for the targeted attitude. This is consistent with a mediation whereby the framing conditions change the conceptual structure which causes a change in attitude. This is tested statistically using a causal mediation analysis from the mediation package (Tingley et al., 2014). I conducted one test for the in-party frame and one test for the out-party frame and tested the indirect effects with 1000 resampled bias corrected and accelerated bootstrapped confidence intervals. The indirect effect of the in-party frame (Study 4: indirect effect = .04, 95% CI [.025, .060], $p < .001$, proportions mediated = .59; Study 5: indirect effect = .04, 95% CI [.03, .06], $p < .001$, proportions mediated = .95) and the out-party frame (Study 4: indirect effect = -.06, 95% CI [-.094, -.040], $p < .001$, proportions mediated = .83; Study 5: indirect effect = -.03, 95% CI [-.06, -.02], $p < .001$, proportions mediated = .49) through in-party similarity was significant. Robustness tests show that the error terms of the mediator and the dependent variable (indicating confounding) would need to be correlated at .40 (Study 4) or .30 (Study 5) in both mediation models to bring the indirect effect to zero (Imai, Keele, & Tingley, 2010; Imai, Keele, & Yamamoto, 2010).

Exploratory Effects on Similarity with Other Attitudes

To explore how partisan frames might affect the conceptual similarity of the belief system more broadly, I also tested if the frame conditions changes the conceptual similarity between the targeted attitude and the other attitudes included in each study. One possible expectation is that the partisan frames also cause the targeted attitude to be seen as more

negatively related to the food irradiation measure ($b = -.08$, $SE = .05$, $p = .07$). In the own condition, out-party similarity was non-significantly positively related to the food irradiation measure ($b = .05$, $SE = .04$, $p = .25$). This is a similar result to Study 4, suggesting that this interaction may be weak, but replicable. It implies that perceiving changes in *others* conceptual structures might be more important for attitude change than perceiving changes in your *own* conceptual structure.

similar to other attitudes that are consistent with the partisan frame. That is, the targeted attitude will be seen as more similar with typically Democratic supported attitudes (COVID-19 quarantines, Black Lives Matters protests, government run healthcare, taxes on the rich) in the Democratic party frame condition. However, the targeted attitude will be seen as more similar with typically Republican supported attitudes (gun ownership, restrictions on immigration) in the Republican party frame condition. To test this, in separate regression models, I regressed the similarity measures on dummy codes indicating Democratic frame condition and Republican frame condition, party strength, need for cognition, age, race (1 = non-White, 0 = White), education (three categories, dummy coded; high school education or less, some college or bachelor's degree, post-graduate degree), and Republican party identification (1 = Republican identification, 0 = Democratic identification). The coefficients for the key comparisons are in Table 3 (full model results are on Tables S 30 to S 39). Raw means and confidence intervals are in Figures 14 and 15.

In both studies, across the typically Democratic supported attitudes, I found that the Democratic frame increased the perceived similarity with the targeted attitude (all p 's < .002 with one exception at $p = .07$) and the Republican frame decreased the perceived similarity with the targeted attitude (all p 's < .001 with one exception at $p = .07$). Conversely, across the typically Republican supported attitudes, I found that the Democratic frame decreased the perceived similarity with the targeted attitude (all p 's < .006) and the Republican frame increased the perceived similarity with the targeted attitude (all p 's < .03). This suggests that partisan frames may re-organize a person's belief system beyond the targeted attitude.

Table 3. *Estimates of the difference between Democratic frames, Republican frames, and control conditions on the similarity between the targeted attitudes and additional attitudes for Studies 4 and 5.*

Study 4			Study 5		
<i>b</i>	<i>SE</i>	<i>p</i> -value	<i>b</i>	<i>SE</i>	<i>p</i> -value

Democratic Frame vs. Control

Targeted-Black Lives Matter Protests	-	-	-	0.108	0.028	<0.001
Targeted-COVID-19 Quarantines	0.121	0.025	<0.001	-	-	-
Targeted-Gun Ownership	-0.139	0.028	<0.001	-0.079	0.028	0.005
Targeted-Government Run Healthcare	0.106	0.026	<0.001	0.091	0.029	0.002
Targeted-Restrictions on Immigration	-0.133	0.028	<0.001	-0.088	0.029	0.003
Targeted-Taxes on the Rich	0.125	0.029	<0.001	0.052	0.029	0.074

Republican Frame vs. Control

Targeted-Black Lives Matter Protests	-	-	-	-0.143	0.028	<0.001
Targeted-COVID-19 Quarantines	-0.090	0.026	0.001	-	-	-
Targeted-Gun Ownership	0.069	0.029	0.017	0.133	0.029	<0.001
Targeted-Government Run Healthcare	-0.108	0.026	<0.001	-0.154	0.029	<0.001
Targeted-Restrictions on Immigration	0.063	0.029	0.027	0.099	0.029	0.001
Targeted-Taxes on the Rich	-0.054	0.029	0.068	-0.139	0.029	<0.001

Figure 14. Study 4 unadjusted means, confidence intervals, and data points for the effect of cues on conceptual similarity between farm subsidies and additional attitudes.

Figure 15. Study 5 unadjusted means, confidence intervals, and data points for the effect of cues on conceptual similarity between food irradiation and additional attitudes.

Discussion

Studies 1 through 3 demonstrated how conceptual judgments of belief systems could be used to test and conceptually replicate prior work on belief systems, but by using an individual level estimate of belief system structure. Studies 4 and 5 aimed to test how that belief system structure might be influenced by outside forces. Both the Brandt-Sleegers and Goldberg-Stein models predict that people will respond to partisan frames by changing the structure of their belief systems. Consistent with this, I found that people changed how they connected their party affiliation with support for the targeted attitudes (farm subsidy or food irradiation policies) following both in-party and out-party frames. I also found that these changes in belief system structure mediated the effect of the frames on attitude change. In

additional exploratory analyses, I found that the partisan frames may change a number of different connections within the belief system. Together, these results suggest that outside information can change the structure of people's belief systems. These empirical results are consistent with the theoretical predictions of the Brandt-Sleegers and Goldberg-Stein models, which provides evidence for the validity of conceptual similarity judgments for measuring individual belief systems.

General Discussion

Belief systems are in the minds of individuals. This is an assumption shared by multiple theoretical approaches from multiple disciplines (Brandt & Sleegers, 2021; Converse, 1964; Goldberg & Stein, 2018; Kinder & Kalmoe, 2017; Martin, 2000). However, the methods for measuring belief systems have not been at the same level of analysis. Instead, dominant approaches for studying the structure of belief systems have relied on the cross-sectional associations between attitudes and identities (Boutyline & Vaisey, 2017; Brandt et al., 2019; Converse, 1964; Kalmoe, 2020), which can only tap into within-person processes, like a belief system, under stringent and unmet assumptions (Brandt & Morgan, in press; Hamaker, 2012; Molenaar & Campbell, 2009). Although methods have been developed to measure within-person systems for other psychological constructs (e.g., personality, Beck & Jackson, 2020), these same methods are unlikely to work effectively for political belief systems because of the stability of political attitudes (Brandt & In press; Kiley & Vaisey, 2020).

I suggested that and tested if conceptual similarity judgments provide a viable and valid measure of people's belief systems. Conceptual similarity judgments ask people to rate the extent different attitudes and identities are connected and use these ratings to construct person-level belief systems (cf. Stoller et al., 2020). In Studies 1-3, I found that these judgments were reliable and tracked other measures of similarity between political attitudes.

More importantly, there was evidence of the validity of these measures in assessing belief systems because people with more political engagement, political knowledge (only Study 3), attitude consistency, ideological extremity, and party extremity were associated with conceptual similarity ratings consistent with having more strongly interconnected belief systems. Studies 1-3 also showed that participants with similar conceptual similarity judgments also had similar patterns of responding to the political attitude items (i.e. preference congruence). Finally, Studies 4 and 5 showed that conceptual similarity judgments are affected by partisan frames and are associated with attitude change. Across these five studies, and across both cross-sectional and experimental validation tests, I was able to validate conceptual similarity judgments as measures of belief systems by demonstrating their convergent validity with the Brandt-Sleegers and Goldberg-Stein models of belief systems and their discriminant validity from political knowledge.

Benefits and Limitations

The conceptual similarity allows us to measure a person's belief system. The field has not been able to do that previously, despite the clear theoretical utility of such a measure. My approach builds on related techniques aimed at understanding how different people structure traits and trait perceptions (Stolier et al., 2020; Stolier et al., 2018; Xie et al., in press). This provides a solid foundation. The studies further show that the measures are valid and reliable. In addition, the conceptual similarity approach is resource efficient compared to alternative methods. It can be administered using basic survey methods (including pencil and paper), requiring no special tools or software. The primary cost is time. The method uses substantially fewer resources than, for example, intensive longitudinal studies that have been used to estimate within-person associations in other domains. For example, Study 1, the longest of the studies here, took an average of 24 minutes. In Studies 2 and 3, I showed that the study time could be reduced by randomly assigning participants to a subset of the stimuli.

These studies took about 15 minutes. These are not short studies, but are comparable in length to many lab studies and are substantially shorter than other methods that have been used to assess within-person systems (e.g., intensive longitudinal studies).

The conceptual similarity approach is also flexible. I chose attitudes and identities to use based on attitudes and identities that were salient in my target population during the time of study. Other researchers working in other contexts or with other interests can include alternative attitudes and identities. A key open question for people aiming to comprehensively map the entirety of the belief system is how to best identify a comprehensive set of attitudes and identities. This is an attitude that also applies to other methods of estimating belief system structure (e.g., Brandt et al., 2019; Converse, 1964; Kalmoe, 2020), with no consensus yet reached. One option is to draw attitudes from a comprehensive list of societally relevant attitudes, such as the attitudes that are the focus of public opinion polling (see Study 3). Such an approach can help scholars identify important issues.

This manuscript primarily focused on how particular links (or combinations of links) in the belief system were related to theoretically expected outcomes or predictors; however, it should be possible to use the conceptual similarity judgments to compute and study other indicators of belief system structure, such as centrality and density of the belief system (e.g., Ertan et al., 2022). Finally, the conceptual similarity approach advocated here aims to assess the structure of belief systems by assessing what attitudes and identities are connected to each other. Other approaches to the structure of belief systems are interested in the style of the belief system, such as distinguishing between what makes people more or less dogmatic (for example, Zmigrod, 2020). The results of the present study hint that strong links between ideologically consistent attitudes might be one indicator of more dogmatic thinking, but the studies were not designed to test these questions; they are open for future work.

A clear limitation is that the conceptual similarity approach relies on self-report. It thus suffers from the same shortcoming and ambiguities of any self-report measure in psychology. This points to a need for measures of conceptual similarity that do not rely on self-report. A possible point that may allay some concerns of self-report is that both versions of the measure (own vs other perspective) produces largely similar results.

A second limitation is that the conceptual similarity approach does not identify the causal directions between the attitudes and identities in the belief system. Although this assumed bi-directional association is consistent with the theoretical models I am building on (Brandt & Sleegers, 2021; Golberg & Stein, 2018), it is possible that some attitudes and identities are more causally potent than others. One possible solution to this limitation is to adopt perceived causality scaling, which has been used to map the relationships between symptoms of psychiatric disorders (Deserno et al., 2020; Frewen et al., 2012). This method asks about the perceived causality between items and has been used to, for example, explore how clinicians perceived symptoms to be related. This method may also be beneficial for the study of belief systems, however, it also relies on participants' insight into the connections between their attitudes.

It might be tempting to conclude that if resources are available, then intensive longitudinal studies should be preferred to the current approach. They do not require people to self-report associations and they can also handle different attitudes and identities. However, I think that this conclusion would be a mistake. Political attitudes and identities are quite stable across years (Brandt & In press; Kiley & Vaisey, 2020) and shorter time periods (see Study 3 of this manuscript and Brandt & In press). Intensive longitudinal studies often cover a much shorter time (e.g., multiple measurements over the course of a week). This means that any variation that *might* emerge in intensive longitudinal studies that cover shorter time periods is of unclear meaning and provenance. For this reason, I believe that conceptual

similarity judgments should be the preferred method for estimating within-person belief systems.

Which Version of the Measure?

I tested two versions of the measure. One asked participants about the conceptual similarity about their own attitudes, whereas the other asked participants about the conceptual similarity about the average American's attitudes. The Own Version of the measure is more face valid as it directly asks about people's own attitudes; however, by necessity it asks people to imagine that they support one attitude. This can be easy when someone does support the attitude, but may be more difficult when someone does not support the attitude. The Other Version of the measure does not require the imagination aspect that the own version does, but it also does not ask directly about the participant's own belief system. This builds on the intent of work in other subfields, such as stereotyping, which assumes that people may be more willing to divulge their opinions when asked about judgments about someone else.

Each version of the measure has trade-offs. Do the results help us choose? Both versions of the conceptual similarity measure had similar relationships with the observed correlations between political attitudes (Figures 2-4) and with the distance between attitudes within-persons (Figures S3, S7, & S8). Studies 1-3 also found that both versions of the measure were similarly internally consistent and had similar test-retest reliabilities (Figure 11). When it came to individual differences, across nearly every test the measures performed similarly (Figures 5-7). There was one exception. In Study 3, the versions of the measure differed in their associations with attitude positions (Figure 7, Right-most panel). Given that there were not clear theoretical predictions about this association, it is not clear if this difference favors one measure or another. There were significant positive association between belief system coherence (similarity with other's belief systems) and preference

congruence (similar patterns in attitude support) in Studies 1-3, although the effect is stronger for the Own Version than the Other Version in Studies 1 and 2. This might suggest support for the Own Version. In Studies 3 and 4 the experimental conditions affects in-party and out-party to a similar degree across versions of the measure (Tables S14 & S20), but there may be evidence that attitude change is related more to out-party similarity for the Other Version rather than the Own Version (Footnote 17). This might suggest support for the Other Version. Notably, the association between in-party similarity and attitude change is similar across versions.

Looking across all of these analyses, it appears that in nearly all of the use cases, the two measures are quite similar. In the few places where the measures do differ, the differences are relatively small and go in opposite directions; some suggest that the Own Version is ideal, whereas others suggest that the Other Version is ideal. At this stage, I think that researchers are justified in using either version of the measure. If they have the necessary resources, they might consider randomly assigning participants to one of the versions of the measure. As validation is an ongoing process, this would be useful for the continued testing of the validity of the conceptual similarity measures. Personally, when under resource constraints, I plan to use the Own Version of the measure because I value it's face validity and it is nearly always similar to the Other Version. Researchers with different priorities may make a different choice. Given the findings so far, I think that other choices can be justified.

Conclusions

Theoretical modeling of belief system dynamics highlight the need for a person-level measure of belief system structure (Brandt & Sleegers, 2021). The conceptual similarity judgment measure of belief systems opens up doors to direct empirical tests for how belief system structure affects attitude change and stability. This measure may also open up doors to identifying and testing otherwise unidentified ideological or personality differences in belief

systems. Ideological differences have typically focused on mean level differences in values or motivations (e.g., Jost, 2017); however, with a valid measure of belief system structure it will also be possible to identify ideological similarity and differences in the structure of the belief system as well. This helps illustrate how a measure of belief system structure can help open new research questions. By combining empirical observation with increasingly sophisticated theoretical models of belief system dynamics, we can put the study of belief system structure on firmer theoretical and empirical ground.

References

- Albarracin, D., & Vargas, P. (2010). Attitudes and persuasion: From biology to social responses to persuasive intent. In S. T. Fiske, D. T. Gilberg, & G. Lindzey (Eds.), *Handbook of Social Psychology* (pp.394-427). John Wiley & Sons.
- Baldassarri, D., & Goldberg, A. (2014). Neither ideologues nor agnostics: Alternative voters' belief system in an age of partisan politics. *American Journal of Sociology*, 120, 45–95.
- Bakker, B. N., Lelkes, Y., & Malka, A. (2020). Understanding partisan cue receptivity: Tests of predictions from the bounded rationality and expressive utility perspectives. *Journal of Politics*, 82, 1061–1077.
- Banaji, M. R., & Heiphetz, L. (2010). Attitudes. In S. T. Fiske, D. T. Gilberg, & G. Lindzey (Eds.), *Handbook of social psychology* (pp.353-393). John Wiley & Sons.
- Beck, E. D., & Jackson, J. J. (2020). Consistency and change in idiographic personality: A longitudinal ESM network study. *Journal of Personality and Social Psychology*, 118, 1080-1100.
- Boonen, J., Pedersen, E. F., & Hooghe, M. (2017). The effect of political sophistication and party identification on voter–party congruence. A comparative analysis of 30 countries. *Journal of Elections, Public Opinion and Parties*, 27, 311–329.
- Boutyline, A., & Vaisey, S. (2017). Belief network analysis: A relational approach to understanding the structure of attitudes. *American Journal of Sociology*, 122, 1371-1447.
- Brader, T., De Sio, L., Paparo, A., & Tucker, J. A. (2020). “Where you lead, I will follow”: Partisan cueing on high-salience issues in a turbulent multiparty system. *Political Psychology*, 41, 795–812.

- Brandt, M. J., Evans, A. M., & Crawford, J. T. (2015). The unthinking or confident extremist? Political extremists are more likely than moderates to reject experimenter-generated anchors. *Psychological Science*, 26, 189-202.
- Brandt, M. J., & Morgan, G. S. (in press). Measures of political ideology and beliefs are not appropriate for measuring ideologies in people's heads because they primarily measure stable traits. *Journal of Personality and Social Psychology*.
- Brandt, M. J., Sibley, C. G., & Osborne, D. (2019). What is central to political belief system networks?. *Personality and Social Psychology Bulletin*, 45, 1352-1364.
- Brandt, M., & Sleegers, W. (2021). Evaluating belief system networks as a theory of political belief system dynamics. *Personality and Social Psychology Review*, 25, 159-185.
- Broockman, D. E. (2016). Approaches to studying policy representation. *Legislative Studies Quarterly*, 41(1), 181-215.
- Bullock, J. G. (2011). Elite influence on public opinion in an informed electorate. *American Political Science Review*, 105, 496-515.
- Clifford, S., & Jerit, J. (2016). Cheating on political knowledge questions in online surveys: An assessment of the problem and solutions. *Public Opinion Quarterly*, 80, 858-887.
- Clifford, S. & Kennedy, R. (2021). The emotional landscape of American politics. Paper presented at the conference of the American Political Science Association.
- Cohen, G. L. (2003). Party over policy: The dominating impact of group influence on political beliefs. *Journal of Personality and Social Psychology*, 85, 808-822.
- Converse, P. E. (1964). The nature of belief systems in mass publics. In D. E. Apter (Ed.), *Ideology and discontent* (pp. 206-261). New York, NY: The Free Press.
- Deserno, M. K., Borsboom, D., Begeer, S., van Bork, R., Hinne, M., & Geurts, H. M. (2020). Highways to happiness for autistic adults? Perceived causal relations among clinicians. *Plos one*, 15(12), e0243298.

- Eagly, A. H., & Chaiken, S. (1993). *The psychology of attitudes*. Harcourt Brace Jovanovich College Publishers.
- Ellis, C., & Stimson, J. A. (2012). *Ideology in America*. Cambridge University Press.
- Epskamp, S., Waldorp, L. J., Möttus, R., & Borsboom, D. (2018). The Gaussian graphical model in cross-sectional and time-series data. *Multivariate Behavioral Research*, 53, 453-480.
- Ertan, G., Çarkoğlu, A., & Aytaç, S. E. (2022). Cognitive political networks: A structural approach to measure political polarization in multiparty systems. *Social Networks*, 68, 118-126.
- Fishman, N., & Davis, N. T. (in press). Change we can believe in: Structural and content dynamics within belief networks. *American Political Science Review*.
- Fiske, S. T., Cuddy, A. J., Glick, P., & Xu, J. (2002). A model of (often mixed) stereotype content: competence and warmth respectively follow from perceived status and competition. *Journal of Personality and Social Psychology*, 82, 878-902.
- Frewen, P. A., Allen, S. L., Lanius, R. A., & Neufeld, R. W. (2012). Perceived causal relations: novel methodology for assessing client attributions about causal associations between variables including symptoms and functional impairment. *Assessment*, 19, 480-493.
- Gidron, N. (2022). Many ways to be right: cross-pressured voters in Western Europe. *British Journal of Political Science*, 52, 146-161.
- Goldberg, A., & Stein, S. K. (2018). Beyond social contagion: Associative diffusion and the emergence of cultural variation. *American Sociological Review*, 83, 897-932.
- Gross, J. J., & John, O. P. (2003). Individual differences in two emotion regulation processes: Implications for affect, relationships, and well-being. *Journal of Personality and Social Psychology*, 85, 348-362.

- Gurney, K. (1997). An introduction to neural networks. CRC press.
- Hamaker, E. L. (2012). Why researchers should think “within-person”: A paradigmatic rationale. In M. R. Mehl, & T. S. Conner (Eds.), *Handbook of research methods for studying daily life* (pp. 43– 61). New York, NY: Guilford Press.
- Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79, 2554-2558.
- Imai, K., Keele, L., & Tingley, D. (2010). A general approach to causal mediation analysis. *Psychological Methods*, 15, 309-334.
- Imai, K., Keele, L., & Yamamoto, T. (2010). Identification, inference and sensitivity analysis for causal mediation effects. *Statistical Science*, 25, 51-71.
- Johnston, C. D., & Ollerenshaw, T. (2020). How different are cultural and economic ideology? *Current Opinion in Behavioral Sciences*, 34, 94-101.
- Jost, J. T. (2017). Ideological asymmetries and the essence of political psychology. *Political Psychology*, 38, 167-208.
- Kalmoe, N. P. (2020). Uses and abuses of ideology in political psychology. *Political Psychology*, 41, 771-793.
- Kam, C. D. (2005). Who toes the party line? Cues, values, and individual differences. *Political Behavior*, 27, 163–182.
- Kiley, K., & Vaisey, S. (2020). Measuring stability and change in personal culture using panel data. *American Sociological Review*, 85, 477-506.
- Kinder, D. R., & Kalmoe, N. P. (2017). Neither liberal nor conservative: Ideological innocence in the American public. University of Chicago Press.

- Kriegeskorte, N., Mur, M., & Bandettini, P. A. (2008). Representational similarity analysis-connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*, 2, 4.
- Leach, C. W., Van Zomeren, M., Zebel, S., Vliek, M. L., Pennekamp, S. F., Doosje, B., ... & Spears, R. (2008). Group-level self-definition and self-investment: a hierarchical (multicomponent) model of in-group identification. *Journal of Personality and Social Psychology*, 95, 144-165.
- Lupton, R. N., Myers, W. M., & Thornton, J. R. (2015). Political sophistication and the dimensionality of elite and mass attitudes, 1980–2004. *The Journal of Politics*, 77, 368–380.
- Malka, A., & Lelkes, Y. (2010). More than ideology: Conservative–liberal identity and receptivity to political cues. *Social Justice Research*, 23(2-3), 156-188.
- Malka, A., Soto, C. J., Inzlicht, M., & Lelkes, Y. (2014). Do needs for security and certainty predict cultural and economic conservatism? A cross-national analysis. *Journal of Personality and Social Psychology*, 106, 1031-1051.
- Mansueto, A. C., Wiers, R., van Weert, J., Schouten, B. C., & Epskamp, S. (2020, December 2). Investigating the feasibility of idiographic network models.
<https://doi.org/10.31234/osf.io/hgcz6>
- Martin, J. L. (2000). The relation of aggregate statistics on beliefs to culture and cognition. *Poetics*, 28, 5-20.
- Molenaar, P. C. (2004). A manifesto on psychology as idiographic science: Bringing the person back into scientific psychology, this time forever. *Measurement*, 2, 201-218.
- Molenaar, P. C., & Campbell, C. G. (2009). The new person-specific paradigm in psychology. *Current Directions in Psychological Science*, 18, 112-117.

- Morgan, G. S., & Wisneski, D. C. (2017). The structure of political ideology varies between and within people: Implications for theories about ideology's causes. *Social Cognition*, 35, 395-414.
- Price, E., Ottati, V., Wilson, C., & Kim, S. (2015). Open-minded cognition. *Personality and Social Psychology Bulletin*, 41(11), 1488-1504.
- Reicher, S. (2004). The context of social identity: Domination, resistance, and change. *Political Psychology*, 25, 921-945.
- Stolier, R. M., Hehman, E., & Freeman, J. B. (2020). Trait knowledge forms a common structure across social cognition. *Nature Human Behaviour*, 4, 361-371.
- Stolier, R. M., Hehman, E., Keller, M. D., Walker, M., & Freeman, J. B. (2018). The conceptual structure of face impressions. *Proceedings of the National Academy of Sciences*, 115, 9210-9215.
- Tingley, D., Yamamoto, T., Hirose, K., Keele, L., & Imai, K. (2014). mediation: R Package for causal mediation analysis. *Journal of Statistical Software*, 59, 1-38.
- Thornton, M. A., & Tamir, D. I. (2017). Mental models accurately predict emotion transitions. *Proceedings of the National Academy of Sciences*, 114, 5982-5987.
- Visser, P. S., & Mirabile, R. R. (2004). Attitudes in the social context: the impact of social network composition on individual-level attitude strength. *Journal of Personality and Social Psychology*, 87, 779-795.
- Xie, S. Y., Flake, J. K., Stolier, R. M., Freeman, J. B., & Hehman, E. (in press). Facial impressions are predicted by the structure of group stereotypes. *Psychological Science*.
- Zaller, J. R. (1992). *The nature and origins of mass opinion*. Cambridge University Press.
- Zmigrod, L. (2020). A Psychology of Ideology: Unpacking the Psychological Structure of Ideological Thinking. *PsyArXiv* <https://doi.org/10.31234/osf.io/ewy9t>