Journal of Experimental Social Psychology 100 (2022) 104279

Contents lists available at ScienceDirect

Journal of Experimental Social Psychology

journal homepage: www.elsevier.com/locate/jesp

Belief system networks can be used to predict where to expect dynamic constraint[☆]

Felicity M. Turner-Zwinkels[a,*], Mark J. Brandt[a,b]

[a] Department of Social Psychology, Tilburg University, the Netherlands
[b] Department of Psychology, Michigan State University, USA

A R T I C L E I N F O

Editor: Kristin Laurin

Keywords:
Dynamic constraint
Belief systems
Psychological networks
Attitude change

A B S T R A C T

We test if a change in an attitude affects other related attitudes (i.e., dynamic constraint), a core prediction of belief systems theory. We use psychological network methods to represent the belief system and make preregistered predictions about which attitudes should change and to what extent. We collected data in two longitudinal experiments (N = 3004; N = 2999) and three pilot studies (combined N = 2788) from community samples of US Americans. We use data from T1 as pretest measures of attitudes and to estimate the structure of the sample's belief system from which to generate and preregister predictions. At T2 participants were randomly assigned to one of three conditions: a control condition (no manipulation), a terrorism attitude manipulation (Study 1), a crime attitude manipulation (Study 2) attitude manipulation, or a banking attitude manipulation (Studies 1 & 2). We successfully manipulated the targeted attitude and also observed changes in non-targeted attitudes in the belief system. Multilevel models provided evidence that changes in non-targeted attitudes were moderated by their distance from the targeted attitude within the belief system: Non-targeted attitudes closer to the experimentally targeted attitude typically changed more. Changes in non-targeted attitudes were generally related to (and mediated by) changes in the targeted attitude. We discuss the implications of our findings for belief systems theory and the value of network methods in studying attitude change.

In the 1950s and 60s social scientists (e.g., Converse, 1964; Shils, 1968) argued that well-structured political belief systems did not exist in the general public. Prominent among this was the seminal essay of Converse (1964), "The Nature of Belief Systems in the Mass Public". According to Converse, if people have well-structured political belief systems then they should show evidence of both static constraint (i.e., being able to predict one political attitude within the belief system from another; e.g., predicting economic attitudes from welfare attitudes) and dynamic constraint (i.e., a change in one political attitude within the belief system impacts other related attitudes; e.g., a change in taxation beliefs causes a change in social welfare beliefs). He determined that most people did not have static constraint. Recent research has further confirmed this: While some people do have static constraint (e.g., the more informed or higher educated; Stimson, 1975), most do not (Freeze & Montgomery, 2016; (Freeder, Lenz and Turney, 2018); Groenendyk, Kimbrough, & Pickup, 2020; Malka, Soto, Inzlicht, & Lelkes, 2014). Less attention has been paid to dynamic constraint. Despite featuring in Converse's theoretical writing, he did not test it empirically. Where dynamic constraint has been explored, there has been little evidence supporting its existence among political attitudes (Coppock & Green, 2021; Hopkins & Mummolo, 2017; Peffley & Hurwitz, 1992). A change in a political attitude within the belief system does not appear to readily impact other related attitudes, or propagate across a political belief system.

The key factor that distinguishes dynamic constraint research from other psychological research on attitude change is its focus on a (political) belief system. Dynamic constraint concerns the transmission of attitude change across a range of political attitudes that are interrelated within a belief system. As such, it may look at attitude change across attitudes with different strengths of association or with two or more degrees of separation (e.g., attitudes that are not directly connected, but neighbors of neighbors). In other words, if there is dynamic constraint, attitude change should be transmitted through neighboring attitudes to different attitudes about different topics (e.g., across different cultural attitudes about abortion, gender, and race) and across domains (e.g., from cultural attitudes to economic attitudes about business and

☆ This paper has been recommended for acceptance by Dr Kristin Laurin.

* Corresponding author.
E-mail address: f.m.turner@tilburguniversity.edu (F.M. Turner-Zwinkels).

https://doi.org/10.1016/j.jesp.2021.104279
Received 28 May 2021; Received in revised form 25 November 2021; Accepted 28 December 2021
Available online 18 January 2022

banking). For dynamic constraint to be established, it should be caused by a change in one focal attitude that is transmitted to political attitudes with different strengths of association, potentially from different domains. This is different than research that examines change in hierarchically nested attitudes (e.g., the belief that Biden is competent, friendly, charismatic are nested beliefs about U.S. President). Although dynamic constraint-like phenomena can be *implied* in social psychological research that manipulates one attitude to impact others (e.g., cognitive dissonance), this would not typically be considered as a full test of dynamic constraint because it usually focuses on a limited set of within-domain or hierarchically nested attitudes (e.g., dissonance research uses a counter-attitudinal task to change a target attitude and tests for changes in attitudes of the same topic and within the same domain; Cooper, 2019). So, although there is research that manipulates one attitude to impact others (e.g., Bian, Leslie, Murphy, & Cimpian, 2018), there is less research exploring dynamic constraint. Where it has been explored, there is much less evidence supporting it. This seems to suggest that changes in one political attitude do not typically ripple through the belief system to cause changes in other political attitudes (e.g., Coppock & Green, 2021). Thus, the key conclusion from the dynamic constraint literature is that dynamic constraint is rare or non-existent. As such, a core claim of belief systems research remains unverified.

We use psychological network methods to conduct a new test of dynamic constraint. We seek to understand whether dynamic constraint occurs in political belief systems, and if so, is attitude change moderated by the distance between political attitudes within this system? Network methods allow us to represent the multidimensional (Johnston & Ollerenshaw, 2020) structure of the belief system and conduct empirical tests using this information. This is important because prior research that has tested for dynamic constraint has selected a small number of theoretically interesting attitudes that scholars *assume* should be related. Consequently, it is unclear if a lack of dynamic constraint is because there is no dynamic constraint, or because these attitudes are not linked in the belief system in practice and therefore should not be expected to vary together in the first place. Furthermore, this work usually selects one focal attitude and a small selection of other attitudes with a varying strength of relation to this attitude. As such, dynamic constraint is limited to a bivariate space. To address these limitations, we use psychological network methods to estimate which attitudes are related and how strongly they are related within the context of a larger, political belief system. From this, we generated specific (preregistered) predictions about (a) which attitudes within the belief system should change as a result of a change in another attitude, and (b) how large this change should be relative to the distance between attitudes.

We conduct a comprehensive test of dynamic constraint with two longitudinal experiments (and three, large sample pilots) that manipulate targeted attitudes in a belief system. We then examine if these targeted manipulations cause subsequent changes in non–targeted attitudes. We used psychological network analyses to test if change in non-targeted attitudes is proportional to their distance from the targeted attitude in the belief system network. Together, this research takes a systems approach to testing dynamic constraint. In doing so, we aim to contribute to current knowledge on attitudes and their relationships to each other in belief systems.

1. Belief systems and dynamic constraint

A key characteristic of the study of belief systems in its different guises (e.g., ideologies, worldviews, etc.) is that they are built from multiple components (e.g., attitudes or values) that are related to each other in some way (e.g., Boutyline & Vaisey, 2017; Brandt & Sleegers, 2021; Converse, 1964; Gerring, 1997; Jost, 2006; Tedin, 1987). The relationships between the elements of a belief system are important because they allow people to go beyond having a set of isolated attitudes, to having a broader system of meaning which encapsulates their understandings of the world. These relations, otherwise referred to as *constraint* (Converse, 1964; Gerring, 1997), imply internal consistency between attitudes, so that attitudes that are similar in some way are more strongly connected to each other. Traditionally, this is expressed through two key forms of constraint.

First, static constraint is when a person holds attitudes that are qualitatively related (e.g., the attitude that social welfare should be increased is positively associated with the attitude that the government should be larger). In practice, a person is considered as having static constraint when they hold consistently liberal or consistently conservative attitudes. Although individuals are not always consistent in their attitudes (e.g., Converse, 1964), in general the more politically informed an individual is and the more frequently they think about their attitudes, the more consistent they become in terms of static constraint (Bartle, 2000; Converse, 1964; Federico & Schneider, 2007; Judd & Downing, 1990; Kalmoe, 2020; Keating & Bergan, 2017).

Second, dynamic constraint is the diffusion of attitude change throughout the belief system. As such, dynamic constraint concerns the *maintenance* of consistency, logic, and order within a changing belief system. The idea is that if an individual has a disposition towards continuity and consistency within their belief system, then a sufficient change in one attitude should have implications for the other attitudes it is connected to. If, for example, an individual were to change their attitude from anti- to pro-female empowerment, they may also feel the need to change their attitude from pro-life to pro-choice to reduce perceived inconsistency between these attitudes. Together, static and dynamic constraint act to maintain a sense of order and meaning within the belief system. They are not only the glue that holds the belief system together, but are considered a core premise of belief system research.

Although there is evidence for statistic constraint, evidence of dynamic constraint has been less forthcoming. This is because dynamic constraint is more difficult to study. Static constraint can be studied by calculating correlations between attitudes in different subpopulations (e.g., Kalmoe, 2020) or assessing the extent people hold consistently liberal or conservative attitudes (e.g., Federico & Schneider, 2007), with the idea being that higher correlations or higher consistency indicates more static constraint. Dynamic constraint has more stringent requirements than finding mere consistency.

In order to study dynamic constraint, at least three conditions should be met (cf. Coppock & Green, 2021). First, you need to cause a change in a targeted attitude in the belief system and this change needs to be reasonably substantial. If a change is too small, it will not create any notable inconsistency within the system and other attitudes will not readjust to accommodate it. Second, the changes in the targeted attitude should transmit across the belief system and extend beyond the non-targeted attitudes that are most closely related to the target attitudes. As such, attitude change should not be *limited* to attitudes sharing the same domain (e.g., to only cultural attitudes), though attitude change may well be strongest within the domain containing the targeted attitude. This is a key difference between dynamic constraint and spillover effects or hierarchically nested beliefs. For example, the focus of spillover effects is on within-domain effects (e.g., from one environmental behavior like recycling, to another like cycling a bike) and typically finds effects that are larger and more consistent (Nilsson, Bergquist, & Schultz, 2017; Whitmarsh & O'Neill, 2010) than in dynamic constraint research (Coppock & Green, 2021). Similarly, in the case of research on hierarchically nested beliefs, the targeted higher-order belief (e.g., values) causes changes in lower-order, associated beliefs (e.g., attitudes; Peffley & Hurwitz, 1985; Blankenship, Wegener, & Murray, 2012; Blankenship, Wegener, & Murray, 2015). Dynamic constraint, however, goes further and is evidenced by between-domain effects (i.e., between different types of attitudes at the same level of abstraction) which may be difficult to detect because between-domain effects are between attitudes that are not directly relevant to one another. Third, the changes in the belief system should be transmitted via the targeted attitude. If an event (e.g., an economic collapse, a pandemic, an experimental manipulation) directly changes a number of different beliefs, it would not show dynamic constraint, but rather that some events influence multiple attitudes (i.e., it would be a confounding variable).

Research exploring dynamic constraint-like effects in non-political attitudes suggests that it may happen. Research on dynamic constraint in non-political attitudes has been called indirect or lateral attitude change (Bohner, Elleringmann, Linne, Boege, & Glaser, 2020; Brannon, DeJong, & Gawronski, 2019; Glaser et al., 2015) or spatial inertia (McGuire, 1981; McGuire, 1990). These tests typically use stimuli further removed from meaningful political attitudes (e.g., novel space aliens, Bohner et al., 2020) or represent within-domain effects (McGuire, 1990). Furthermore, supportive results suggest that only non-targeted attitudes change if they are directly and strongly related to a targeted attitude, whereas more weakly/distally related non-targeted attitudes do not change (Brannon et al., 2019). This might suggest that dynamic constraint occurs among attitudes that are closely related, but is less likely to transmit more widely through a belief system to distally related attitudes.

When dynamic constraint is tested in the political domain, little evidence is found. Between-domain effects, in particular, appear to be rare or nonexistent (Coppock & Green, 2021). This work has used observational data (Peffley & Hurwitz, 1992) and experimental data (Brannon et al., 2019; Coppock & Green, 2021; Hopkins & Mummolo, 2017), with only little support. Experimental data is ideal for testing for dynamic constraint because it randomizes exposure to a persuasive appeal, allows targeting of only one attitude to change, and facilitates comparison with a control group. Coppock and Green (2021) present a comprehensive test of dynamic constraint across three experiments using a between-subjects designs. Despite successfully inducing differences in a targeted attitude between the experimental and control groups, they find no clear support for dynamic constraint in other attitudes. This might suggest that dynamic constraint does not exist among political attitudes, or is quite rare.

There are two key limitations of this research. First, research has tended to focus on a theoretically selected set of attitudes (Brannon et al., 2019; Coppock & Green, 2021). Measuring constraint based on researcher-defined assumptions about “what goes with what” involves imposing (elite) expectations on a sample. However, there may be various reasons why belief systems in the mass public deviate from these expectations. For example, non-elite belief systems may not be typically structured (Converse, 1964), groups may lack the motivation or have pragmatic reasons to deviate from normative belief systems (Groenendyk et al., 2020), or people may have competing interests which cause meaningful and intended deviations from a typical liberal-conservative ideological structure (Baldassarri & Goldberg, 2014). Such an approach may then focus on practically unrelated beliefs and could consequently underestimate dynamic constraint. Said another way, researchers may not have identified dynamic constraint in the political domain because the attitudes included in their studies are not connected in the structure of the belief system in the first place. Second, such experimental research (Brannon et al., 2019; Coppock & Green, 2021) also focuses on a small group of attitudes (usually approximately four), exploring the bivariate relation between one focal belief and some other beliefs that should be more closely and more distantly related to it. Our research will conduct a more detailed test of the impact of distance between attitudes on dynamic constraint, testing the relation among more attitudes with a wider variety of connection distances/strengths.

In our research, we build on Coppock and Green’s (2021) experiments in two ways. First, we use a longitudinal, experimental design to directly assess attitude change. Second, we model the sample’s belief system as a network of fifteen attitudes. This allows us to represent the structure of the belief system within our specific population under study, including estimates of which attitudes are connected, how strongly they are connected, and the proximity of their connections. From this, we can derive specific predictions about where we expect attitude change to occur and how strongly.

2. Belief systems as networks

To estimate the connections between attitudes in the belief system, we take a network approach to modelling belief systems. This approach conceptualizes belief systems as a network of interconnected attitudes or values (Boutyline & Vaisey, 2017; Brandt, 2020; Brandt, Sibley, & Osborne, 2019; Brandt & Sleegers, 2021; Fishman & Davis, 2019). These methods explicitly model (a) *multiple* attitudes and beliefs relevant to politics and (b) they model the *interrelationships* between them as partial correlations, in (c) a multidimensional system (Brandt et al., 2019). Prior research has demonstrated the utility of this network approach to modelling political belief systems (Brandt et al., 2019), individual attitudes (Dalege, Borsboom, van Harreveld, Waldorp, & van der Maas, 2017; see also connectionist models, Monroe & Read, 2008), and moral belief systems (Turner-Zwinkels, Johnson, Sibley, & Brandt, 2020), among others (e.g., stereotypes, Sayans-Jiménez, van Harreveld, Dalege, & Rojas Tejada, 2019; identity, Phua, Leong, & Hong, 2020; climate change beliefs, Verschoor, Albers, Poortinga, Böhm, & Steg, 2020; trust, Zhang, Liu, Brown, & Gil de Zúñiga, 2020). Relevant for our studies, networks have also been used to model how psychological constructs might change (Chambon et al., 2021).

Partial correlation-based network methods, like those we use here, are consistent with theoretical assumptions of belief systems (Brandt & Sleegers, 2021). Importantly for dynamic constraint, nodes that are positively connected to one another want to be like each other (e.g., both in a “liberal” state), whereas nodes that are negatively connected to each other want to be dissimilar (e.g., one in a “liberal” state and the other in a “conservative” state). When the connection is stronger, these tendencies are amplified (for an introduction see Dalege et al., 2016). These assumptions are consistent with findings that people tend to have consistent attitudes/belief systems (Festinger, 1957; Heider, 1946). As such, using a network approach to estimate the structure of belief systems may be an ideal way to anticipate how dynamic constraint plays out.

Notably, the partial correlation networks that we estimate here are most closely aligned with the idea of static constraint, but can be used to *hypothesize* how attitude change may ripple through the belief system. For example, if Attitude A in Fig. 1 changes, then other attitudes that it is directly connected to should also change (i.e., Attitudes B and C). Furthermore, the change in attitudes should be proportional to the strength of their connection with Attitude A (e.g., proportional to the partial correlation of the connection). Finally, this change (i.e., in Attitudes B & C) may be further transmitted to their neighboring attitudes (e.g., Attitudes, D & E). This means that the attitudes that are most strongly connected Attitude A change the most, and those further away

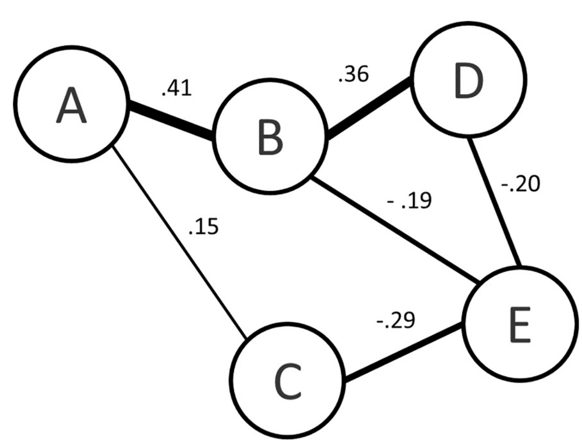

Fig. 1. A network of attitudes (A, B, C, D, E), where the links show the presence of a connection between attitudes, with thicker links showing stronger partial correlations.

or least strongly connected change less (or not at all).

There are at least three important advantages to estimating partial correlation networks prior to testing dynamic constraint hypotheses. First, the belief system network approach explicitly models the connections between attitudes as the strength of unique associations (i.e., partial correlations). Although some work has used bivariate correlations to measure constraint (Converse, 1964; Kalmoe, 2020; Boutyline & Vaisey, 2017), partial correlations are appropriate for estimating the transmission of attitude change required for dynamic constraint. This is because dynamic constraint concerns the transmission of attitude change from one attitude to another via unique causal associations between attitudes (i.e., not a confounding third variable; see condition #3 of dynamic constraint). Partial correlations fit this idea best. Second, networks allow us to include a variety of different attitudes because networks represent a multidimensional belief system. This allows us to estimate belief system structure and then test for dynamic constraint across a range of attitudes simultaneously. Third, partial correlation networks facilitate a more focused test of dynamic constraint by generating specific predictions about (a) which non-targeted attitudes should be impacted by the change in a targeted attitude, and (b) using the distance between attitudes in the belief system to predict precisely how strongly these non-targeted attitudes should be impacted by a change in the targeted attitude across a broader belief system. A similar network approach has also been recently applied in testing the impact of change in one node on its neighbors in a network of behavioral and psychological reactions to the COVID-19 pandemic, with somewhat promising results (Chambon et al., 2021). Together, this network approach should allow fresh insight into our knowledge about dynamic constraint in belief systems and the conditions under which it occurs.

3. Study overview

We present two longitudinal experiments testing for dynamic constraint in belief systems. In these studies, we report all measures, manipulations, and exclusions. Each experiment consists of two key stages. In the first wave (T1), we collected a premeasure of attitudes to (1) measure baseline attitudes against which we can assess if attitude change occurs, and (2) estimate partial correlations networks to estimate the structure of the belief system. We used the belief system structure to generate specific preregistered hypotheses about the extent that each non-targeted attitude will change (preregistration 1: https://osf.io/g8dsx; preregistration 2: https://osf.io/pqnwe). In the second wave (T2), we randomly assigned participants to experimental conditions. Experimental conditions included a persuasive appeal targeting one specific attitude and the control condition did not include an appeal. We then measured attitudes again. To test for dynamic constraint, we first tested if we successfully induced attitude change on the targeted variable (i.e., a baseline condition necessary for assess dynamic constraint), and then tested if other non-targeted attitudes also changed.

We test three key hypotheses. First, the opposition hypothesis (H1): The manipulation of an attitude only affects the specific attitude targeted. That is, there is no dynamic constraint (Coppock & Green, 2021). This is contrasted with our distance hypothesis (H2): the tendency for change in non-targeted attitudes (i.e., dynamic constraint) will decrease with distance from the targeted attitude in the belief system. That is, we expect that non-targeted attitudes that are more distant from the (experimentally) targeted attitude in the belief system structure will be impacted less than non-targeted attitudes that are closer to the targeted attitude.

The third hypothesis we test is the small world hypothesis (H3), which is the prediction that changes in non-targeted attitudes will be limited to the domain of the targeted attitude. For example, take the case where attitudes about terrorism and defense were in one domain (e.g., a security/violence domain) and attitudes about business and technology were in another domain (e.g., about the economy). H3 predicts that in the event where an individual's attitude to terrorism changed, then their attitude to defense should also change as this is part of the same attitude domain as terrorism. But H3 also predicts that attitudes to business and economy should not change as they are in a different domain. We used a bottom-up approach to identify domains. We conducted a community membership analysis, adopted from social network analysis, to identify domains of items in the belief system network. The idea is that attitudes identified as belong to the same community are in the same domain (i.e., roughly equivalent to a factor in factor analysis; Golino & Epskamp, 2017). The models testing this hypothesis are reported in supplementary materials because we found no evidence that attitude change was limited to the same domain as the targeted attitude. Instead, attitude change was similarly observed for attitudes both inside and outside of the targeted attitudes' domain. This is consistent with the idea that dynamic constraint happens across a broad section of the belief system. Finally, we also tested an explorative hypothesis (H4; not preregistered) that predicted that change in non-targeted attitudes will be transmitted via change in the targeted attitude (a key assumption of dynamic constraint). This hypothesis seeks to test if the data are consistent with the theory that changes in the non-targeted attitudes are caused by a change in the targeted attitude.

A key concept in our studies is the distance between attitudes. There are three key ways to operationalize distance between attitudes, which we will refer to as the *expected impact* of attitudes on each other. (1) *Direct impact* is the strength of direct connections (roughly approximate to a partial correlation). This is the simplest measure of impact that only takes into account the direct connections between non-targeted attitudes and the targeted attitude(i.e., if we manipulate Attitude A in Fig. 1, only Attitudes B and C should change proportional to the strength of connection with Attitude A). (2) *Indirect* impact is the strength of both direct and indirect connections. This calculates the shortest paths between a targeted attitude and the non-targeted attitudes. It therefore accounts for both attitudes that are directly connected and indirectly connected to the targeted attitude (i.e., if we manipulate Attitude A in Fig. 1, Attitudes B, C, D and E should all change to different extents proportional to the shortest distance to Attitude A). (3) *Ising* simulations of attitude change are the most complex measure of impact. This approach uses simulations of belief system dynamics (see also, Brandt & Sleegers, 2021; Dalege, Borsboom, van Harreveld, & van der Maas, 2017) to take into account the bidirectional influence between all attitudes, across the network as a whole, when calculating attitude change in non-targeted attitudes (explained in more detail below). As we have no a priori reason to prefer one measure of expected impact to another we employed three impact measures to test our distance hypothesis. However, results across the three impact measures were very similar. As such, we report only the Ising impact measure which fitted the data best (according to AIC/BIC values of models). The results of all impact measures are presented in supplementary materials.

4. Pilots 1 & 2

We conducted two, well-powered pilot studies to identify experimental manipulations that would reliably change participants' attitudes with a small-medium effect-size. We built on manipulations of attitudes used in prior work on dynamic constraint. Both pilots used the same materials and between participants design, with the exception that a different (banking) manipulation was tested in Pilot 2, following an unsuccessful manipulation result in Pilot 1. Results are summarized below, for a full report (i.e., including all measures, manipulations, and analyses) see supplementary materials.

Participants were recruited via Prolific (see sample details in Table 1) and randomly allocated to one of two framing conditions (Frame: terrorism v. banking). Depending on condition, they either read a short text highlighting the risk of terrorism in the U.S. or the need to change banking regulations. Pilot 1 sought to manipulate attitudes to increase support for (a) higher spending on terrorism (Hopkins & Mummolo,

Table 1
Sample information and means, standard deviations and independent samples *t*-tests comparing terrorism and banking attitudes for terrorism and banking conditions, Pilots 1 & 2.

Attitude	Condition	Stat	Pilot 1	Pilot 2
		N	790	792
		Payment	$1.20	$0.45
Terrorism attitude	Terrorism	M	3.75	3.76
		SD	1.49	1.44
	Banking	M	3.29	3.29
		SD	1.63	1.58
		95% CI:	0.24, 0.69	0.24, 0.68
		t-test	3.99	4.13
		df	730.94	720.57
		p-value	<0.001	<0.001
		Cohen's d	0.29	0.31
Banking Attitude	Terrorism	M	5.22	5.22
		SD	1.45	1.34
	Banking	M	5.05	5.91
		SD	1.51	1.91
		95% CI:	−0.05, 0.39	−0.87, −0.51
		t-test	1.59	7.42
		df	703	727.06
		p-value	0.14	<0.001
		Cohen's *d*	0.12	0.54

Note. Df differ between tests where they were adjusted to account for unequal between conditions in the Levene's test.

2017), and (b) reduced banking regulations (adapted from Coppock & Green, 2021). The banking manipulation was unsuccessful, so Pilot 2 tested a new banking manipulation adapted from Eadeh and Chang (2020) that aimed to increase support for increased regulations of banks. Following this, participants completed a series of items, the main dependent variables being *targeted attitudes* (i.e., targeted by the manipulation) towards spending on terrorism and banking regulations. These were each measured with a single item, on a seven point Likert-type scale (anchored at 1 = greatly decrease spending/regulations; 7 = greatly increase spending/regulations; 8 = don't know; 9 = haven't thought about it much; with 8 & 9 coded as missing). Thirteen other *non-targeted attitudes* were measured in Pilot 1 (e.g., crime, aid to the poor, benefits for the unemployed, healthcare, stimulating the economy) using the same item anchors. A reduced set of six *non-targeted attitudes* were measured in Pilot 2.

4.1. Manipulation check

We ran an independent samples *t*-test to test if the terrorism/banking manipulations successfully manipulated the targeted attitudes of terrorism spending and banking regulations, respectively (see Fig. 2/ Table 1). Results showed that the terrorism manipulation was successful in Pilots 1 & 2: As expected, participants in the terrorism condition endorsed higher levels of spending on the war on terrorism than participants in the banking condition. Although the banking manipulation did not successfully induce a difference between banking attitudes in the two conditions in Pilot 1, the new banking manipulation tested in Pilot 2 was successful. In this case, participants in the banking condition supported more regulations for banking in comparison to those in the terrorism condition. Thus, both manipulations in Pilot 2 successfully induced a change in the targeted attitude.

4.2. Discussion

Pilots 1 and 2 show that the terrorism and banking manipulations are effective: The manipulations induced differences on the targeted attitudes. We were also able to conduct a preliminary, and less ideal test of dynamic constraint (see supplemental materials for full details). Results for the terrorism (Pilots 1 and 2) and banking (Pilot 2) conditions generally aligned with the expectations of H2 – non-targeted attitudes closest to the terrorism/banking attitude showed the largest differences between groups. These tests, however, were based on belief system networks estimated after the experimental conditions, only tested between group differences, and did not include a control condition. In our main study, we use a longitudinal design for our key tests. This way we can estimate the belief system network prior to the manipulations, test the effect of the manipulation on within-subjects attitude change, and include a control group that will allow for a clearer interpretation of effects.

5. Study 1

Study 1 tested for dynamic constraint using a three-wave longitudinal design (with two weeks between each wave) to test if a change in a targeted attitude (terrorism/banking) would be associated with a change in non-targeted attitudes (e.g., defense/welfare). T1 premeasured attitudes. At T2 participants were randomly allocated to manipulations of terrorism or banking attitudes, or to a no manipulation control group. After this they completed the measure of attitudes again. At T3 participants completed an additional post-measure of attitudes. This design allowed us to make at least four key contributions. (1) To assess within-subjects change between T1 and T2. (2) To use a network analysis to estimate the structure of the belief system at T1 (i.e., before the experimental manipulation) and generate specific, preregistered predictions about where and how strongly dynamic constraint should

Table 2
Multilevel model predicting non-targeted T2 attitude, from non-targeted T1 attitude, condition (0 = control, 1 = terrorism/banking), and Ising impact, for the terrorism model (left) and banking model (right), Study 1.

	Terrorism model				Banking model			
	B	SE	B	SE	*B*	SE	B	SE
Step 1								
Attitude T1	0.63**	0.01	0.63**	0.01	0.63**	0.01	0.63**	0.01
Condition	0.09*	0.03	0.10*	0.04	0.05	0.03	−0.08**	0.03
Ising impact	−2.06**	0.46	−1.98**	0.46	1.90**	0.47	1.91**	0.47
Interaction: condition * impact	−0.47**	0.13	−0.53**	0.14	0.44**	0.12	0.45**	0.13
Constant	4.65**	0.05	4.63**	0.05	4.33**	0.06	4.38**	0.06
Step 2								
T1 targeted attitude			−0.00	0.02			0.06*	0.01
T2 targeted attitude			0.04*	0.02			0.20**	0.01
Log	−37,063.22		- 35,166.60		−36,855.30		−33,265.19	
AIC	74,142.44		70,353.20		73,726.59		66,550.38	
BIC	74,207.12		70,433.52		73,791.27		66,630.22	
Marginal R^2 model	0.43		0.44		0.45		0.48	
Marginal ΔR^2 (including impact)	0.06				0.06			
Marginal ΔR^2 (including interaction)	0.0001				0.0002			

Note. *p < .05; **p < .01.

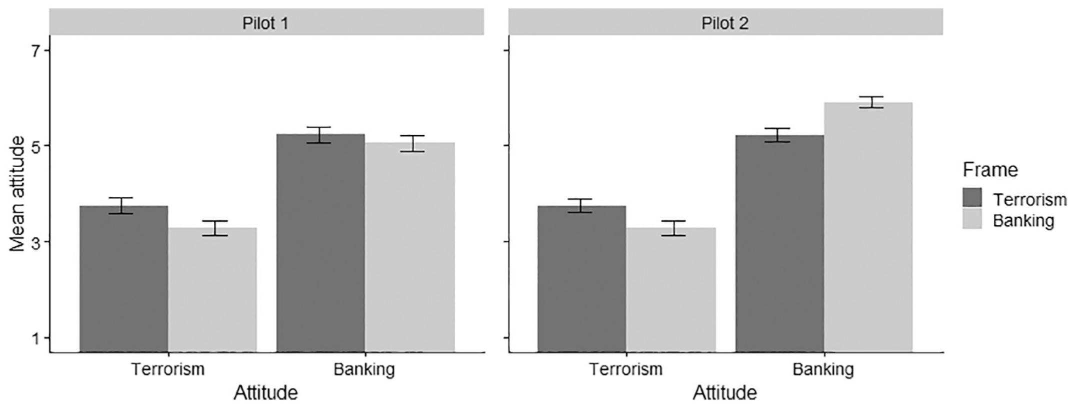

Fig. 2. Mean self-reported attitude to terrorism and banking according to condition, Pilots 1 & 2. Error bars are 95% confidence intervals of the means.

occur. (3) The addition of the control condition allowed us to test for attitude change against a stable baseline, controlling for natural changes in attitudes over the same time frame (e.g., via news media, etc.). (4) T3 allowed us to test for the consistency of possible attitude change overtime. We did not have predictions about over-time effects.

5.1. Method

5.1.1. Preregistrations

Methods and analyses were all preregistered at the OSF: https://osf.io/e3d8t. Specific predictions about which non-targeted attitudes that would show change and how strongly were preregistered at https://osf.io/g8dsx on the basis of the belief system network constructed from the T1 data. That is, the predictions about dynamic constraint were specific to the belief system estimated from this sample.

5.1.2. Design and sample

We ran a 3 (frame: terrorism v. banking v. control) × 3 (time point: T1 v. T2 v. T3) mixed within- and between-subjects design. Participants were recruited online via Prolific at the end of October 2019 to take part in T1. The two subsequent time points were measured at two week intervals (i.e., early and late November 2019), each being left open for one week for data collection. At T2, participants were randomly allocated to one of three frame conditions before completing the main survey. Participants were paid 75 cents for the completion of each wave (i.e., $2.25 in total if all surveys were completed). Rows in our data with duplicates of participants' unique identifiers were removed, leaving only their first data entry.

We recruited 3004 participants at T1 aged 18–81 (M = 34.44, SD = 12.27; 1429 = male, 1500 = female, 44 = non binary, 15 = transgender, 6 = other, 10 = prefer not to answer). This sample size was determined by an a priori power analysis in G*power (Faul, Erdfelder, Lang, & Buchner, 2007) using data from the pilot study (i.e., that the manipulation led to a minimum between subjects effect size on targeted attitude of Cohen's d = 0.12; see pre-registration). On the basis of this we determined that 1000 per condition (~3000 total) will give us 80% power for d = 0.125, 90% for 0.145, and 99% for 0.192. At T2, 2804 participants (93.34%) returned, and 2684 (89.34%) returned for T3. All participants were US residents and had not completed any previous version of this study/pilot. Participants were around the average level of education (M = 6.85, SD = 1.88, mode = 7; where 6 = trade/technical/vocational training, 7 = Associate degree), and household income (M = 3.76, SD = 1.72, mode = 4, $50,000 – $74,999; where 3 = $35,000 to $49,999, 4 = $50,000 to $74,999). The majority of participants were white/Caucasian (N = 2164, Asian = 287, Black/African American = 252, Hispanic/Latinx = 195, Prefer not to answer = 25, other = 81). Thus, our sample was fairly representative of the general US population on these dimensions tested.

Sensitivity analyses conducted in G*power at an alpha significance criterion of 0.05 and power of 80% indicate that this sample size should (a) be sufficient to detect an effect size of Cohen's d = 0.11 for independent samples ($n_{control} = 934$; $n_{terrorism} = 934/n_{banking} = 936$) *t*-tests of the difference in change in targeted attitude between the control and experimental conditions, and (b) a lower critical $R^2 = 0.0002$ in a linear random multiple regression with four predictors (H2).

5.1.3. Procedure

Participants volunteered to participate in a set of three studies titled "U.S. Policies - Attitudes and Opinions". They took part individually online. Participants completed the attitude measures at all three time points. At T1, participants also completed controls (ideology, income, political knowledge, importance of crime/banking,) and demographics. At T2, participants were randomly allocated to one of three conditions at the start of the survey (i.e., before completing the attitude measures). They read either the terrorism or banking frames (as in the pilot; followed by the persuasiveness and agreement items) or no additional text (i.e., the control). At T3, participants again completed the attitude measures. Full copies of all survey materials (including introduction, manipulations, survey items etc.) can be found at https://osf.io/vq7wg/.

5.1.4. Measures

Attitudes. Attitudes towards federal spending on 11 different policy areas were measured with single items on a seven point Likert type scale (anchored at 1 = greatly decrease spending; 7 = greatly increase spending; 8 = don't know; 9 = haven't thought about it much, where 8 & 9 were coded as missing). The topics were: defense, the war on terrorism, crime, aid to the poor, benefits for the unemployed, healthcare, stimulating the economy, improving the social and economic position of black people, controlling immigration, foreign affairs, and foreign aid. For example attitudes towards defense were measured with the item "Should federal spending on defense be increased, decreased, or kept the same?". Attitudes towards government regulation of four topics were measured on a seven point Likert type scale (from 1 = greatly decrease regulations; 7 = greatly increase regulations; 8 = don't know; 9 = haven't thought about it much; where 8 & 9 were coded as missing). This included banking, big-business, tech-companies and gun control. For example, big-business was measured as "Should regulations of big-business be increased, decreased, or kept the same?". We chose these topics to represent a range of issues in US American politics.[1]

Ideological strength was measured with the following item "When

[1] We analyzed the data to see if the number of people reporting attitudes as 8 or 9 differed systematically across conditions. We found no differences (in either Study 1 or 2) following correction for multiple comparisons. Thus, we conclude that our data does not have substantial differential response problems.

thinking about your political beliefs, do you see yourself as a liberal, conservative, moderate, or haven't you thought much about this?" (from 1 = strongly liberal; 7 = strongly conservative regulations; 8 = don't know; 9 = haven't thought about it much; where 8 & 9 were coded as missing).

Political knowledge was measured with a 5 questions from Carpini and Keeter (1993) of varying difficulty about the U.S. political system (e.g., Whose responsibility is it to determine if a law is constitutional or not…; with answer options: the President, the Congress, the Supreme Court). Participants' response to each question was coded as correct (=1) or incorrect (=0), and summed to a maximum score of 5 (minimum 0).

Importance of terrorism/banking were each measured with a single item: "How important or unimportant is the topic terrorism in the U.S.'/ 'Banking Regulation' to you"?", anchored at 1 = very unimportant, 7 = Very important).

Demographics. Gender, ethnicity, education, income and religion were measured. *Education* was measured using an 11-item categorization ranging from no schooling complete, nursey school to 8th grade, up till Doctorate degree. Total household *income* was measured using a six-item category: less than $20,000, $20,000–$34,999, $35,000–$49,999, $50,000–$74,999, $75,000-99,999, over $100,000.

Finally, *identity strength* and how much participants thought about terrorism/banking was also measured, but is not reported further as these items were not utilized in our analysis.

5.1.5. Analytical strategy

Network estimation. One network was constructed from the total sample at T1. The network was estimated using the qgraph (Epskamp, Cramer, Waldorp, Schmittmann, & Borsboom, 2012) and bootnet package (Epskamp, Borsboom, & Fried, 2018) in R version 3.5.1 (R Core Team, 2019). We estimated regularized polychoric, partial-correlation networks, applying the Extended Bayesian Information Criteria graphical LASSO (i.e., least absolute shrinkage; Foygel & Drton, 2010), following Epskamp and Fried's (2018) recommendations. Links between attitudes (i.e., nodes) potentially range from −1 to +1, and represent the (strength of the) relation between two attitudes, conditioning on all other attitudes present in the network. As such, links can be thought of as partial correlations; the association between nodes after controlling for the relations among all other nodes in the network. To simplify the final network models, increase replicability, and reduce type I error, the regularization LASSO method shrinks small links (i.e., which we are uncertain about) to zero, so that they are not included in the final network. All networks applied a tuning parameter of 0.5. Missing data was treated using full information maximum likelihood estimation.

Ising impact. This was calculated using Ising (1925) model simulations via the package "IsingFit" (van Borkulo & Epskamp, 2016) and igraph (Csardi & Nepusz, 2006). Ising models take a network as input and calculate the probabilities that nodes in this network will be "on" or "off" (i.e., whether attitudes will be endorsed or not; see also Dalege, Borsboom, van Harreveld, & van der Maas, 2017). This allows us to take the belief system network gathered in our sample at T1 and use simulations to model the extent we can expect non-targeted attitudes should change following a persuasion attempt on either the terrorism or banking attitude. Thus, we took the belief system network (consisting of all 15 attitudes) at T1 and ran 3 × 1000 simulations on this network. Each of these three sets of simulations sought to simulate the different conditions in one of our three (control/experimental) conditions:

First, we simulated the control condition with no persuasion attempt. Second, we simulated the terrorism condition with a persuasion attempt on the terrorism attitude. Third, we simulated the banking condition with a persuasion attempt on the banking attitude. In both the experimental conditions, the persuasion attempt was equally strong. As such, thresholds of the targeted attitude were set at 0.5; indicating that it has a disposition to be "on". Non-targeted attitudes all had their thresholds set at 0 (i.e., indicating they are likely to be "off"), which simulates the situation where the non-targeted attitudes were only impacted via the persuasion attempt on the targeted (terrorism/banking) attitudes. Whether non-targeted attitudes were ultimately impacted (i.e., if their end result was being in an "on" or "off" state, coded as 1 and 0 respectively), is therefore crucially determined by the network structure, including link strength between attitudes and the bidirectional, reciprocal influence of attitudes on each other (i.e., attitudes seek to align with their neighbors).

One final parameter that affects the simulation output is the 'temperature', which controls how much randomness acts on the attitudes. We set the temperature parameter as low (temperature = 1.2). This means there are low levels of randomness in our simulations, therefore increasing the extent that the network structure determines if non-targeted attitudes end up being "on" or "off". This allows us to test more directly the impact of network structure on dynamic constraint. Finally, in each of the three conditions, we summed separately the tendency for attitudes to finish the simulation in either an "on" (i.e., 1) or "off" (i.e., 0) state across the 1000 simulations ran, and divided this by 1000 to give the proportion of runs for which each attitude was "on" in each condition.

Finally, we calculated the effect size for the difference between proportions in two experimental conditions and the control condition (using the pwr package in R; Champely, 2020). This gave a final value quantifying the extent to which we expected attitude change in the experimental conditions would differ from the control. These effect sizes of the difference between conditions were used as our measure of Ising impact. In short, we used an Ising model to simulate attitude change within a belief system based on our manipulations. Then we used the effect sizes that were the outcome of that simulation as our measure of Ising impact. As such, Ising impact represents the extent that the simulations expected change in non-targeted attitudes in the experimental conditions to be larger or smaller than in the control condition, based on the structure of the belief system network.

Multilevel models. The models were estimated using lme4 (Bates, Maechler, Bolker, & Walker, 2015) and lmerTest (Kuznetsova, Brockhoff, & Christensen, 2017). Figures were made using ggplot2 (Wickham, 2016) and cowplot (Wilke, 2019).

5.2. Results

5.2.1. Manipulation

We first tested if the manipulations were successful at changing the targeted attitude. They were (see Fig. 3). From T1 to T2, participants in the terrorism condition (M = 0.31, SD = 1.21) supported a greater increase in spending on anti-terrorism than did participants in the control condition (M = 0.03, SD = 1.22; t(1720) =4.73, p <. 001, 95% CI difference: 0.16–0.39, d = 0.23). Looking at the effect of the banking manipulation, we see a larger attitude change, with participants in the banking condition showing a greater increase in support for banking regulations (M = 0.59, SD = 1.16) than those in the control condition (M = 0.05, SD = 1.31; t(1608.5) = 8.83, p <. 001, 95% CI difference: 0.42–0.66 d = 0.26). Both manipulations changed attitudes as expected, strengthening support for terrorism spending and banking regulations in the respective terrorism and banking conditions.

5.2.2. Dynamic constraint analysis

To test dynamic constraint, we generated the belief system network from T1 attitudes (i.e., before exposure to the experimental manipulation in T2) and ran a community analysis to find the different domains of attitudes (see Fig. 4). Four domains were retrieved (consistent with Pilot 1): (1) Violence (i.e., terrorism, defense, crime, immigration), (2) International (i.e., foreign affairs, foreign aid), (3) Inequality (i.e., Unemployment, healthcare, aid for poor, support for black people, guncontrol), and (4) Business (i.e., business, banking, technology).

We used this network as an estimate of belief system structure at T1 to make predictions about the magnitude of attitude change among fourteen non-targeted attitudes at T2 (following exposure to a

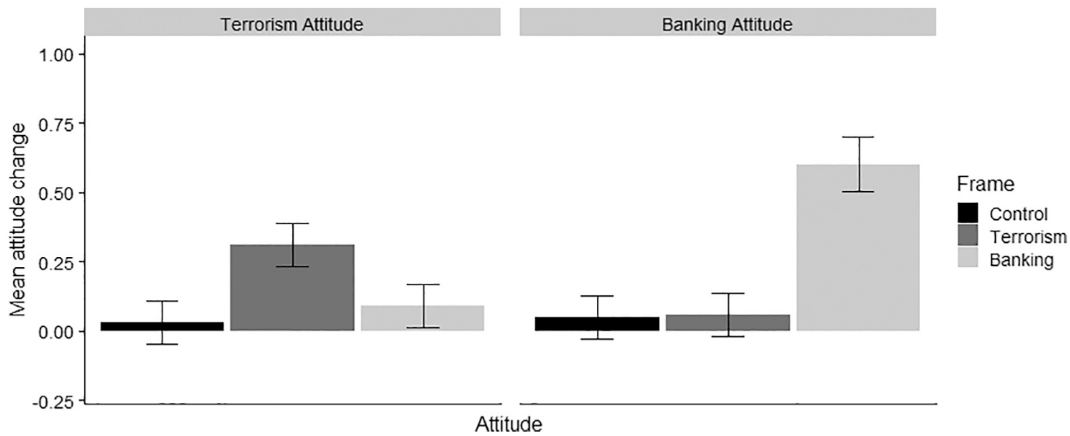

Fig. 3. Change in mean self-reported attitude to terrorism (left) and banking (right) between T1 and T2 according to condition. Error bars are 95% confidence intervals of the mean attitude change.

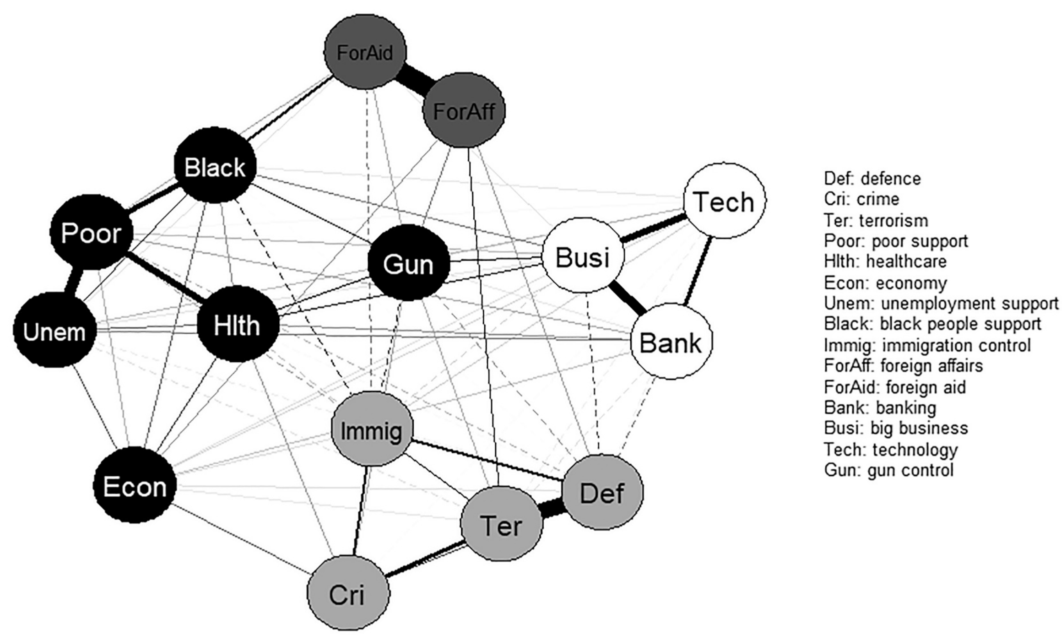

Fig. 4. Network visualization with nodes colored according to community membership (i.e., domain). Links between attitudes represent partial correlations. Thicker links represent stronger connections, with positive links shown as solid and negative links are dashed.

manipulation of one targeted attitude; i.e., terrorism/banking). Specifically, we used this network to generate measures of distance or impact between targeted and non-targeted attitudes. We expected that for participants in the experimental conditions, attitudes closest to the target attitude (i.e., with highest Ising impact) would show the largest changes between T1 and T2. We then preregistered these predictions (see https://osf.io/g8dsx).

5.2.2.1. Descriptive analysis. We ran a preliminary (non-preregistered) analysis to describe the amount of attitude change we observed in our sample for each attitude. Fig. 5 presents the difference in attitude change observed for each non-targeted attitude (between T1 & T2) between the experimental conditions and the control condition. This is plotted according to Ising impact from the manipulated node. Most notably, Fig. 5 presents evidence counter to the opposition hypothesis (H1): Non-targeted attitudes changed more in the experimental conditions than in the control condition, especially within the banking condition. This is consistent with dynamic constraint (i.e., that change in targeted attitudes impacts non-targeted attitudes). Specifically, 2 of 14 and 3 of 14 attitudes changed significantly more in the terrorism and banking conditions than in the control condition, respectively (controlling for the false discovery rate 0/14 and 3/14 differences remain, respectively). This suggests initial support for the presence of dynamic constraint, but does not control appropriately for within subjects effects (i.e., all non-targeted attitudes are nested within individuals).

5.2.2.2. Main analysis. To formally test if dynamic constraint is determined by the distance between these attitudes and the targeted attitudes in the belief system (terrorism or banking attitudes), we ran a series of preregistered multilevel models (see Table 2). Our models predicted non-targeted attitudes at T2 controlling for non-targeted attitudes at T1. This allows us to test for attitude change. We nested attitudes (level 1) within participants (level 2), assigning both random-intercepts. These multilevel models stack all fourteen non-targeted attitudes within individuals to control for statistical dependencies within individuals responses (Snijders & Bosker, 2011).

To test H2 and see if non-targeted attitudes with higher Ising impact from the targeted attitude show larger change than non-targeted

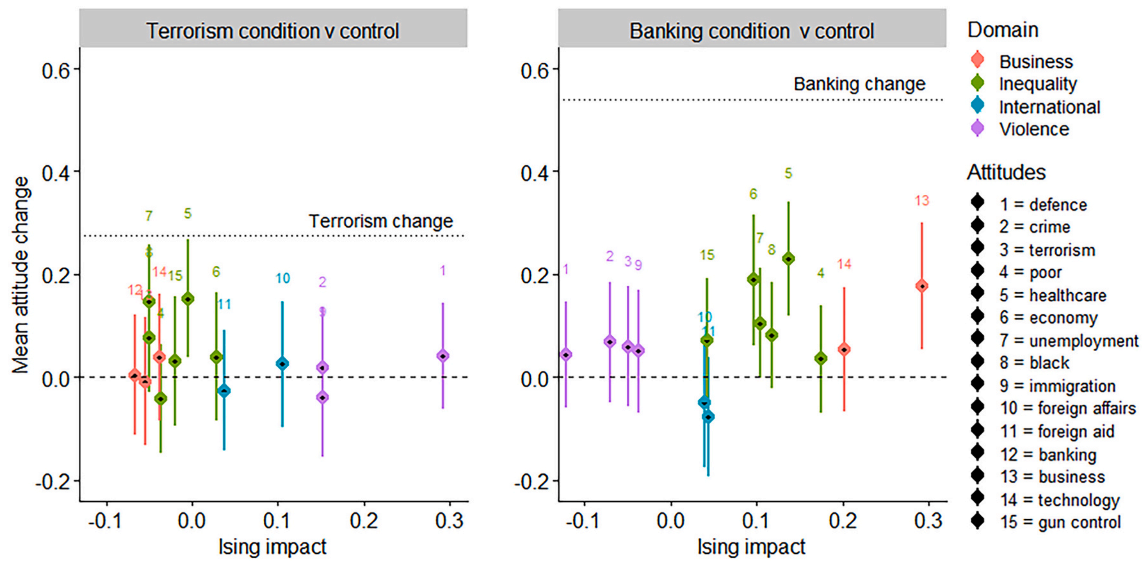

Fig. 5. Scatter plot of differences in attitude change between experimental and control conditions, plotted according to the expected Ising impact. Points represent differences in attitude change between T1 and 2 for the terrorism (left panel) and banking (right panel) conditions in comparison to the control condition. Higher numbers indicate that attitude change of participants in the experimental condition was more positive than for participants in the control condition. Attitudes are colored according to community membership (i.e., domain; A color version of the figure can be found in the online version of the manuscript). Error bars are 95% confidence intervals of the mean attitude change; Study 1.

attitudes that are further from the targeted attitude we used the following predictors in step 1 of our model: non-targeted attitude at T1, condition (0 = control condition, 1 = terrorism condition/banking condition), Ising impact, and a two-way interaction between condition and Ising impact. The interaction is the key test of the hypothesis because it tells us if the size of the attitude change of non-targeted attitudes at T2 is associated with its expected impact from the targeted attitude, within the experimental conditions. In step 2 of our model (which was not preregistered), we added the centered targeted attitude (i.e., terrorism/banking) at T1 and T2. This allows us to test if the change in non-targeted attitudes is caused by change in the targeted attitude (a mediation model testing this causal model is in the supplementary materials). This tests the explorative hypothesis (H4) that change in non-targeted attitudes will be transmitted via change in the targeted attitude.

We also checked if four preregistered controls – strength of ideology, income, importance of crime/banking, and political knowledge – moderated effects by including the interactions between these variables and the key predictors. In a deviation from our pre-registration, controls were mean centered to render model parameters meaningful and aid the interpretation of effects (for models see supplementary materials). Finally, we tested if effects were consistent at T3.

5.2.2.2.1. Terrorism models. Our terrorism models do not support the distance hypothesis (H2; see Table 2, left side), which predicts a *positive* two-way interaction between condition and impact. Instead, we find a significant *negative* interaction between condition and impact: the further the non-targeted attitude was from the terrorism attitude, the greater the increase in support for spending/regulations in participants in the terrorism condition compared to the control condition. We ran simple slopes to test if the terrorism condition differed from the control

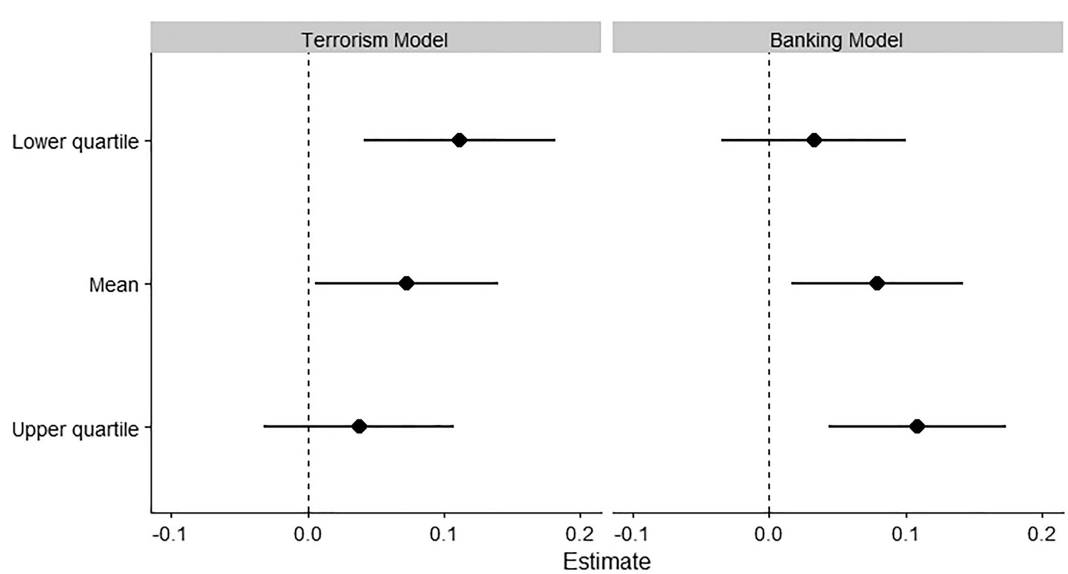

Fig. 6. Simple slopes and 95% confidence intervals for the terrorism model (left) and banking model (right) modelling the interaction between conditions (0 = control, 1 = terrorism/banking frame) and Ising impact at lower quartile, mean and upper quartiles of impact of Ising impact (right panel), Study 1.

at different levels of Ising impact (lower quartile, mean, upper quartile; see Fig. 6). Results unexpectedly showed that the terrorism condition differs positively and significantly from the control only at lower levels of impact. At mean levels of impact, the differences between the terrorism and control conditions are smaller but still significantly different from zero. At higher levels of impact the effect is non-significant. This shows that attitudes that are more distant/with lower Ising impact from the targeted terrorism attitude changed more in the terrorism condition than the control. This difference gets smaller for attitudes closer/with higher expected impact from the targeted terrorism attitude. This is the opposite pattern than we expected. This negative interaction remained, with the addition of the four controls – ideology, terrorism importance, income and political knowledge (although it becomes marginally significant at $t = 1.82$, $p < .07$). Thus, results present little support for our distance hypothesis: although impact was related to attitude change in non-target nodes, the effect of direction was opposite than expected. More distant non-targeted attitudes changed more strongly in the experimental condition than those that were closer to the targeted terrorism attitude.

In step 2 we added terrorism attitudes at T1 and T2 to test if change in the targeted terrorism attitude is related to change in non-targeted attitudes. Consistent with a pattern of dynamic constraint, we find that an increase in the terrorism attitude at T2 (i.e., controlling for T1 attitude) is significantly and positively related to change in the non-targeted attitudes. Additionally, we ran an explorative mediation model (see supplementary materials), which tested the predicted causal chain (i.e., that change in the non-targeted attitudes is caused by change in the targeted attitude). Although results were somewhat weak, we found the expected mediation that the terrorism condition positively predicted change in the targeted variable which positively predicted change in non-targeted attitudes. Consistent with H4, this mediation was also moderated, so that the effect was largest at higher Ising impact (i.e., the negative interaction between condition and Ising impact remains, but is partly offset by the positive interaction between T2 targeted attitude and Ising impact, suggesting moderated mediation).

5.2.2.2.2. Banking models. The banking models show support for the distance hypothesis (H2; see Table 2, right side). The expected positive interaction between impact and condition was observed. This showed that, within the banking condition, the closer the non-targeted attitudes were to the banking attitude (i.e., the higher Ising impact), the larger their change between T1 and T2, compared to the control. Simple slopes comparing the banking condition with the control at different levels of Ising impact showed the expected pattern (see Fig. 6). In comparison to the control, participants in the banking condition showed greater attitude change for non-targeted attitudes that had a higher expected impact from the targeted banking attitude. This difference became smaller and non-significant at lower levels of expected impact. Moreover, the predicted two-way interaction remained significant and positive across all models including controls. Thus, the banking manipulation supported predictions for the distance hypothesis, suggesting that non-targeted attitudes that were further away from the targeted banking attitude changed less than those that were closer to the targeted attitude.

In step 2, we added banking attitudes at T1 and T2 in order to test if change in the targeted banking attitude is related to change in non-targeted attitudes. Again, we find a significant positive effect indicating that change in the targeted banking attitude is related to change in the non-targeted attitudes. Furthermore, the mediation model analyses (in supplementary materials) align with the interpretation that changes in the non-targeted attitudes are mediated by changes in the targeted attitude, and that this effect is largest at higher levels of Ising impact (i.e., supporting H4). In other words, our data is consistent with the theory that changes in the non-targeted attitudes are caused by a change in the targeted attitude, and are greatest when the targeted attitude has a larger impact on (or is closer in distance to) the non-targeted attitude.

5.2.3. Consistency effects at T3

We included a final measurement two weeks later, at T3, to test if the effects of the manipulations were sustained over time. Attrition did not vary according to condition (χ^2 (2) = 0.04, p = .98). Results suggested that the banking manipulation had a small, but sustained impact on participants' banking attitudes at T3, with participants in the banking condition maintaining an increase in support of banking regulations in T3 compared to the control condition. Effects of the terrorism manipulation did not appear to have a sustained effect (i.e., there were no differences between the terrorism and control condition). Although we did find that some non-targeted attitudes differed significantly between the experimental and control conditions (range = 3–4 per condition), multilevel models showed that this relation was not strongly or consistently explained by Ising impact for the banking condition. For the terrorism condition, the unexpected negative interaction (between condition and impact) persisted. However, given that there was not a sustained effect of the manipulation on the targeted attitude it is hard to make sense of this result in the context of dynamic constraint. Thus, altogether, we did not find clear evidence supporting the presence of dynamic constraint at T3 (i.e., no clear support for H2). For a full report of analyses, see supplementary materials.

5.3. Discussion

We found mixed support for dynamic constraint predictions. Following the manipulation of targeted attitudes of terrorism and banking these attitudes increased between T1 and T2. Non-targeted attitudes also changed more strongly in comparison to a control group (i.e., controlling for changes taking part naturally in society via societal events), at least in the short-term. Importantly, consistent with H2, the distance between targeted and non-targeted attitudes in the belief system moderated change in the non-targeted attitude. For the banking condition, non-targeted attitudes that were further from the targeted attitude (i.e., those with lower expected impact) changed less than those that were more proximate to the targeted attitude (i.e., those with higher expected impact). This was consistent with our expectations. However, unexpectedly, the opposite relation as found within the terrorism condition. Although this presents evidence that distance between attitudes in a belief system is a determinant of attitude change, the direction of the change was not as expected. According to research on lateral attitude change (Brannon et al., 2019), this unexpected pattern observed in the terrorism condition may be due to displacement. Displacement occurs when an individual consciously rejects attitude change, so that there is no change in the non-targeted attitudes closest to the targeted attitude. Instead, the attitude change is displaced across the belief system, so that non-targeted attitudes further from the targeted attitude change instead (Glaser et al., 2015; Brannon et al., 2019; see discussion for further discussion). Notably however, the explorative mediation analysis suggests that this unexpected negative effect is offset somewhat by a positive change in non-targeted attitudes with higher Ising impact, proportional to the change in the targeted terrorism attitude. This suggests that our data shows both patterns of displacement (i.e., attitudes far away change more) and dynamic constraint (i.e., attitudes closer change more). Nevertheless, indirect effects observed in the exploratory mediation analyses for both terrorism and banking models support the interpretation that the attitude change observed in non-targeted variables in the experimental condition was transmitted via change in the targeted attitude. Although results from this mediation model should be taken with a grain of salt – given that such models rest on assumptions (e.g., sequential ignorability) which cannot be established here (Bullock, Green, & Ha, 2010) – this model is valuable in predicting the preconditions of dynamic constraint. Together with step 2 in our multilevel models, this mediation test supports the interpretation that non-targeted attitude change is caused by a change in the targeted attitude. As such, this is evidence supporting dynamic constraint.

The results of Study 1 were mixed and the manipulations we used both induced a sense of threat to change attitudes. To probe if our mixed results are due to chance and to extend beyond threatening manipulations, we sought to replicate and extend Study 1 by running a second longitudinal study. This study included a new manipulation that (a) did not induce a sense of threat in participants, and (b) pushed attitudes in the opposite direction (i.e., supporting reduced spending/regulations).

6. Pilot 3

We ran Pilot 3 to pretest three possible crime manipulations to use in Study 2. Results are summarized here, for a full report (including all measures, manipulations and exclusions) see supplementary materials. We recruited 1206 U.S residents via Prolific (male = 558, female =628; 14 = non-binary, 4 = transgender; 2 = missing/preferred not to say), who were paid $0.60. Participants were randomly allocated to one of four conditions (frame 1: crime threat v. frame 2: crime reducing v. frame 3: crime policy change v. control) and completed the same attitudes items as in Study 1. We selected the crime manipulation Frame 3 (crime policy; M = 3.09, SD = 1.62) that was associated with the greatest decrease in attitudes towards crime in comparison to the control (M = 4.36, SD = 1.39; 95% CI: 1.01–1.51, $t(544.02) = 9.92$, $p < .0001$, d = 0.84; see Fig. 7).

7. Study 2

7.1. Method

7.1.1. Preregistrations

We followed the Study 1 preregistration where relevant (e.g., banking manipulations and measures; https://osf.io/e3d8t). New materials (e.g., crime manipulation) and new predictions based on belief system networks from T1 of Study 2 were preregistered at https://osf.io/pqnwe.

7.1.2. Design and sample

We ran a 3 (frame: crime v. banking v. control) × 2 (time point: T1 v. T2) mixed within and between subjects design. All participants completed the survey online using Prolific. T1 (mid-February 2020) contained baseline attitude, control, and demographic measures. Participants who returned two weeks later to complete T2 were randomly allocated to one of three conditions and then completed the main survey. T2 remained open for participation for one week and took approximately 5 min to complete. Participants were paid 75 cents for the completion of each Wave (i.e., $1.50 in total if both surveys were completed). Rows with duplicate unique identifiers were removed, leaving only the first entry.

We recruited 2999 participants (1386 = male, 1539 = female, 47 = non-binary, 10 = transgender, 11 = prefer not to answer, 3 = other, 2 = Missing) for T1 (based on power and sensitivity analyses reported in Study 1). All participants were US residents and had not completed any previous version of this study/pilot. Participants were on average fairly highly educated (M = 6.82, SD = 1.90, mode = Bachelor degree), with an average household income (M = 3.81, SD = 1.76, mode = $50,000 – $74,999). The majority of participants were White/Caucasian (N = 2078, Asian = 299, Black/African American = 256, Hispanic/Latinx = 228, Prefer not to answer = 32, other = 105). Two thousand six-hundred and forty-two participants returned for T2 (an 88.09% retention rate).

Sensitivity analyses conducted in G*power at an alpha significance criterion of 0.05 and power of 80% indicate that this sample size should (a) be sufficient to detect an effect size of Cohen's $d = 0.12$ for independent samples ($n_{control} = 888$; $n_{terrorism} = 875/n_{banking} = 879$) t-tests of the difference in change in targeted attitude between the control and experimental conditions, and (b) a lower critical $R^2 = 0.0002$ in linear random multiple regression with four predictors (H2).

7.1.3. Procedure

The procedure was identical to that of Study 1 (T1 & T2), with two exceptions. First, at T2 participants completed the newly developed crime manipulation instead of the terrorism manipulation. Second, the 'terrorism importance'/ 'terrorism thinking' items measured in T1 were replaced with two items about crime (i.e., How important or unimportant is the topic 'crime in the U.S.' to you/I think about the topic 'crime in the U.S.' a lot). Both the banking condition, control condition, and all the other items remained the same as Study 1. This allows us to replicate the banking condition results from Study 1.

7.2. Results

7.2.1. Manipulation

In order to test if the manipulation was successful we calculated the raw difference between participant attitudes at T1 and T2 and ran an independent samples t-test comparing the two experimental conditions with the control, separately. Results confirmed that both manipulations were successful (see Fig. 8). First, participants in the crime condition supported a much greater reduction in spending on crime at T2 relative to T1 (M = −1.09, SD = 1.67), in comparison to participants in the control condition (M = −0.01, SD = 1.22; $t(1418) = 14.60$, $p < .001$, 95% CI of difference: 0.93–1.22, $d = 0.74$). Second, participants in the banking condition supported a greater increase in regulations on banking (M = 0.63, SD = 1.30) than the control (M = 0.06, SD = 1.23; t

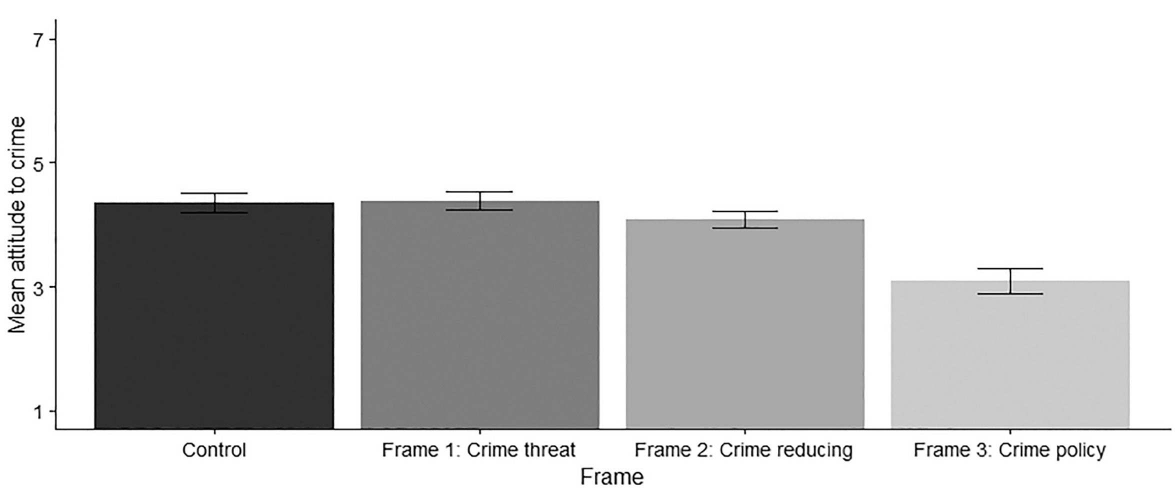

Fig. 7. Bar plot of support for spending increase/decrease on crime with 95% confidence interval around estimate according to three crime frames, Pilot 3. Error bars are 95% confidence intervals of the mean.

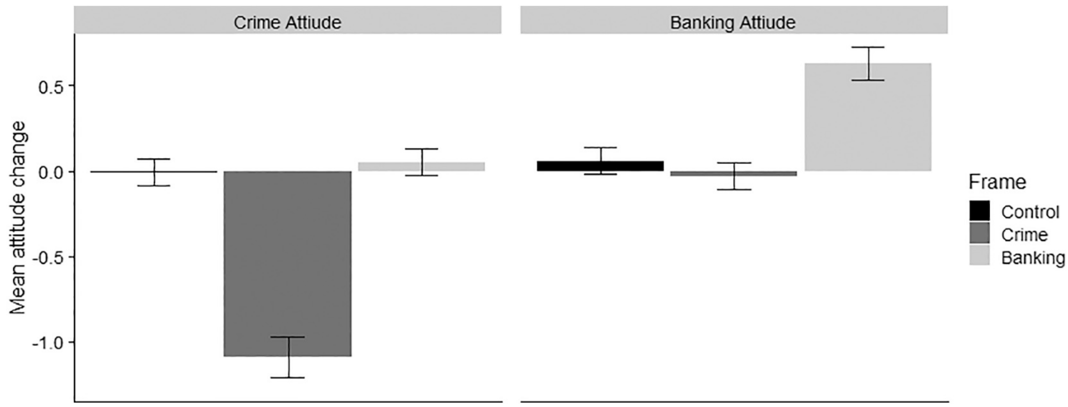

Fig. 8. Change in mean self-reported attitude to crime (left) and banking (right) between T1 and T2 according to condition, Study 2. Error bars are 95% confidence intervals of the mean difference.

(1508.8) = 8.84, p < .001, 95% CI difference: 0.44–0.70, *d* = 0.45). Thus, both manipulations changed participant attitudes in the expected direction with notable small to medium effect sizes.

7.2.2. Dynamic constraint analysis

7.2.2.1. Descriptive analysis. In order to test dynamic constraint, we first generated the belief system network from T1 attitudes (see Fig. 9). The same domains (i.e., violence, international, inequality and business) emerged as in Study 1. Predictions derived from this network (for direct, indirect and Ising impact) are presented in supplementary materials. Fig. 10 presents the difference in attitude change between experimental conditions and the control condition. Most notably, it presents evidence counter to the opposition hypothesis (H1). Results show that a small number of non-targeted attitudes also changed more strongly in experimental conditions than in the control: We observed 1/14 non-targeted attitudes that showed significantly greater change in the crime condition than the control condition, 2/14 non-targeted attitudes in the banking condition, and a number of marginal effects (controlling for the false discovery rate 1/14 and 0/14 differences remain, respectively).

7.2.2.2. Main analysis. In order to test if dynamic constraint among non-targeted attitudes is determined by their distance from the targeted attitudes (crime or banking) we ran the same pre-registered multilevel models as in Study 1, predicting non-targeted attitudes at T2.

7.2.2.2.1. Crime models. The crime models (see Table 3) offer some support for the distance hypothesis (H2). The predicted negative interaction between condition and impact emerged. This suggests that within the crime condition, non-targeted attitudes that were closest to/with highest Ising impact from crime show the greatest reduction. Simple slopes (see Fig. 11) demonstrated that attitude change in non-targeted attitudes followed the expected pattern: Slopes for the crime condition were significantly more negative than for the control conditions and this difference was largest and most negative at higher levels of Ising impact.

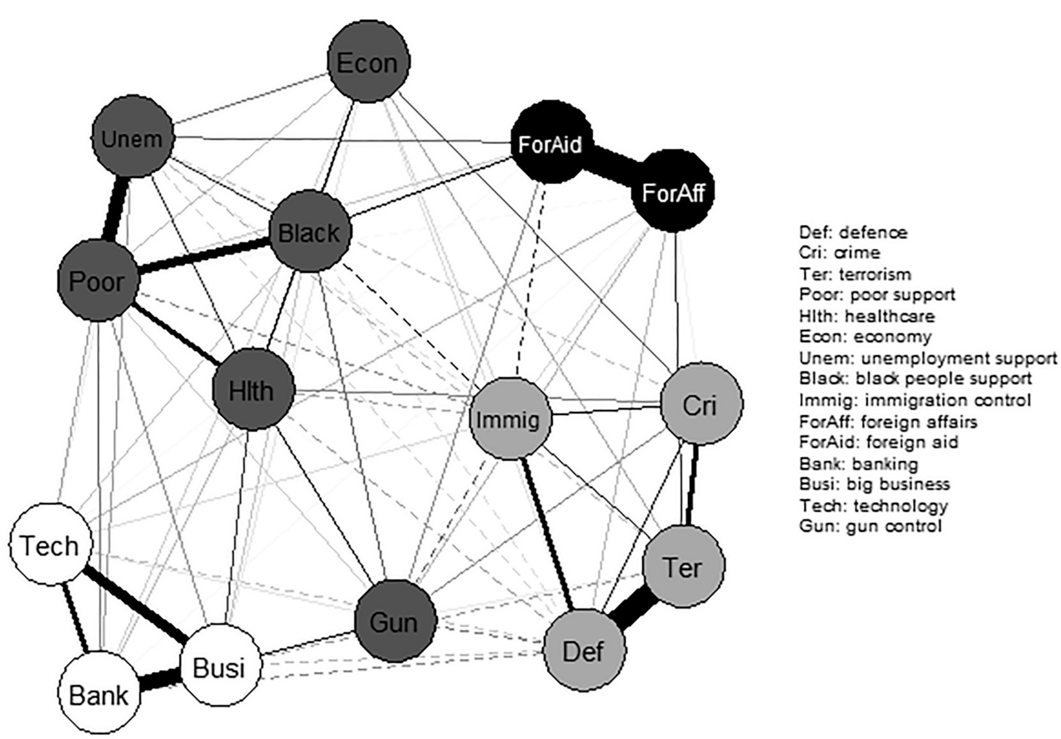

Fig. 9. Network visualization with nodes colored according to cluster membership. Thicker links between nodes represent stronger connections, with positive links shown as solid and negative links are dashed. Study 2.

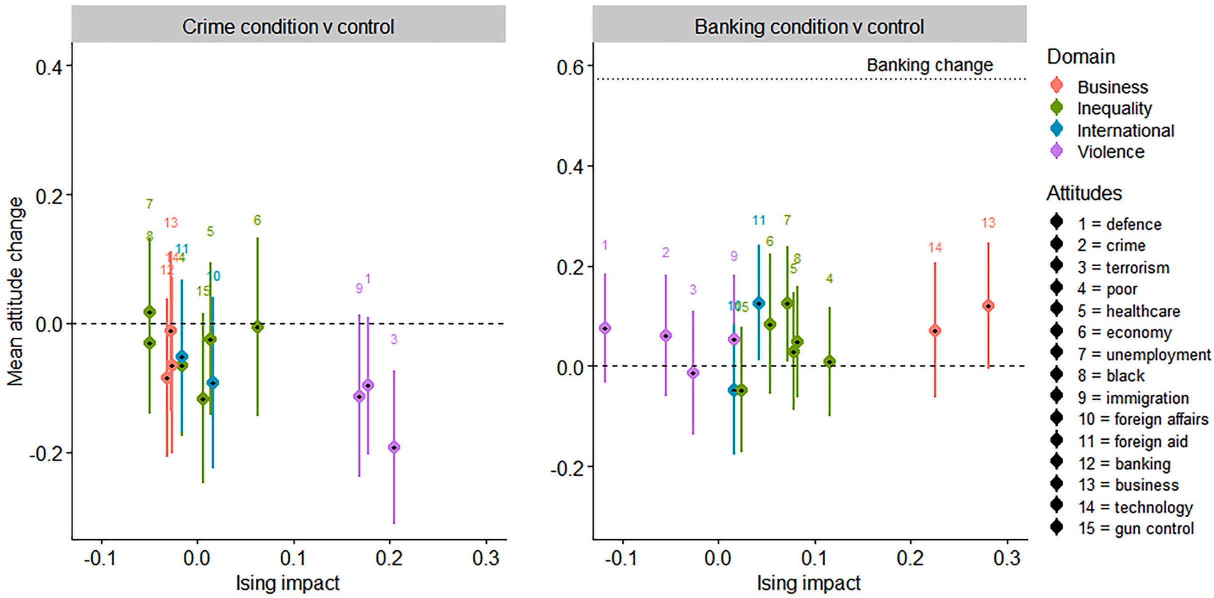

Fig. 10. Scatter plot of differences in attitude change between experimental and control conditions, plotted according to the expected Ising impact. Points represent differences in attitude change between T1 and T2 for the crime (left panel) and banking (right panel) conditions in comparison to the control condition. Higher numbers indicate that attitude change of participants in the experimental condition was more positive than for participants in the control condition. Attitudes are colored according to community membership (i.e., domain; A color version of the figure can be found in the online version of the manuscript). Error bars are 95% confidence intervals of the mean attitude change. The difference in crime attitude change between control and crime condition was −1.10, but was not added to the graph so as to allow approopriate scaling of the y-axis. Study 2.

Table 3
Multilevel crime models, predicting non-targeted T2 attitude, from non-targeted T1 attitude, condition (0 = control, 1 = crime and impact, Study 2.

	Crime model				Banking model			
	B	SE	B	SE	B	SE	B	SE
Step 1								
Attitude T1	0.62**	0.01	0.62**	0.01	0.61**	0.01	0.61**	0.01
Condition	−0.09**	0.03	−0.06	0.04	−0.00	0.03	−0.12**	0.03
Ising impact	−2.76**	0.83	−2.66**	0.83	2.26**	0.70	2.94**	0.83
Condition * Ising impact	−0.35*	0.17	−0.45**	0.18	0.34*	0.14	0.34*	0.15
Constant	4.67**	0.08	4.68**	0.08	4.39**	0.08	4.44**	0.08
Step 2								
T1 targeted attitude			0.02	0.01			0.08**	0.01
T2 targeted attitude			0.02	0.01			0.19**	0.01
Log	−35,024.33		−32,467.78		−34,670.36		−30,875.06	
AIC	70,064.65		64,955.57		69,356.73		61,770.11	
BIC	70,128.82		65,035.01		69,420.88		61,849.17	
Marginal R^2 final model	0.45		0.45		0.45		0.49	
Marginal ΔR^2 (including impact)	0.06				0.06			
Marginal ΔR^2 (including interaction)	0.00007				0.00007			

Note. ˆ $p < .10$, *$p < .05$, **$p < .01$.

However, the difference between low and high levels of Ising impact was smaller than expected. Notably, in models including controls of ideology, income, crimes importance and political knowledge the expected negative interaction between crime and impact remained significant (at $p < .02$). Thus, consistent with H2, participants in the experimental crime condition showed (slightly) greater support for the *reduction* in spending on the non-targeted attitudes that were closer to the targeted crime attitude, however the effect of Ising impact was smaller than expected.

In step 2 we added targeted crime attitudes at T1 and T2 to test the explorative expectation that dynamic constraint should be caused by changes in the targeted attitude (H4). Consistent with this, we found that T2 crime attitudes were positively related to targeted attitude change (i.e., controlling for T1 crime attitudes), but this relationship was smaller than expected and only marginally significant ($t = 1.54$, $p = .12$). In line with this, our mediation analysis (see supplementary materials) showed that the moderated mediation effects were very small. First, there is a negative, but small and non-significant ($p = .12$) indirect effect of condition on non-targeted attitude change at T2 via change in targeted crime attitudes. Second, the interaction between targeted attitude and Ising impact was significant and positive across all models (in line with explorative H4). This suggests that the bigger the change in the targeted crime attitudes the bigger the change in the non-targeted attitudes the closer they are to the crime attitude. However, overall, evidence of mediation was weak and inconsistent.

7.2.2.2.2. Banking models. The banking model (see Table 3) supports the distance hypothesis (H2): non-targeted attitude change was moderated by Ising impact. As expected, we observed a positive interaction effect, which indicates that the closer the attitude to the targeted banking attitude the greater the increase in spending/regulations support in the banking condition compared to the control. However, simple slopes (see Fig. 11) show that the difference between the banking and

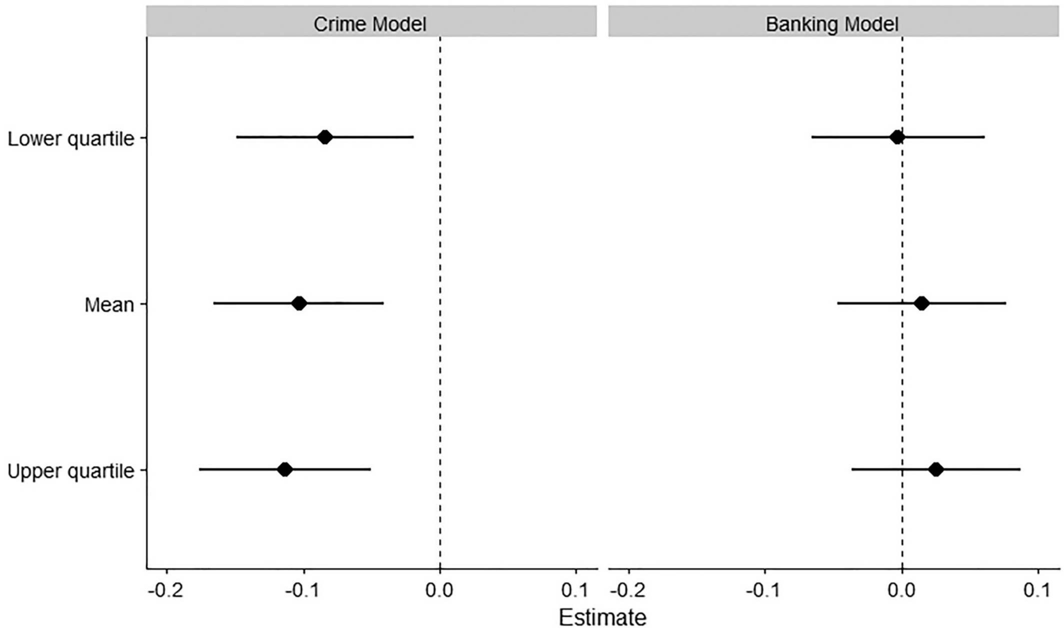

Fig. 11. Simple slopes and 95% confidence intervals for the crime model (left) and banking model (right) modelling the interactions showing the difference between conditions (0 = control, 1 = crime/banking frame) at lower quartile, mean and upper quartiles of Ising impact, Study 2.

control condition was non-significant. This difference did, however, follow the expected pattern: It was estimated as zero at lower quartiles, and slightly above zero at upper quartiles of impact (but only reached significance above the upper quartile). Moreover, this positive interaction was robust to the inclusion of the controls (i.e., ideology strength, income, banking importance, and political knowledge). In sum, results for the banking condition largely replicated. We found support for H2, showing that dynamic constraint is partly informed by distance between attitudes in a belief system, although effects were not as large as expected.

In step 2 we found that, as expected, T2 banking attitudes were positively and significantly related to non-targeted attitude change. The mediation analysis (see supplementary materials) showed the expected positive mediation effect of condition on increase in non-targeted attitudes via the change in targeted attitude. Specifically, we found positive and significant indirect effects of targeted attitude change on non-targeted attitudes. Moreover, in line with H4, this indirect effect is moderated by Ising impact so that it is stronger and more positive at higher levels of Ising impact. Together, these analyses suggest that changes in non-targeted attitudes are caused by changes in targeted attitudes, a condition of dynamic constraint.

7.3. Discussion

We found support for our expectation that distance between attitudes in a belief system matter for the transmission of attitude change from a targeted attitude to a non-targeted attitude (H2). Results were consistently in the expected direction, but differences between the control and experimental condition were weaker than expected. Step 2 in our multilevel models and the exploratory mediation show a pattern of results consistent with dynamic constraint for the banking condition: changes in non-targeted attitudes were associated with changes in targeted attitudes. Indeed, our explorative mediation revealed the expected indirect effects of condition on non-targeted attitude change via the targeted attitude, and there was evidence that this effect was moderated by impact, but again not as strongly as expected. However, evidence of mediation was more equivocal within the crime condition. Overall, this suggests that the structure of the belief system plays a role in determining the flow of attitude change between multiple attitudes.

8. General discussion

Across two longitudinal experiments and three pilot studies, we found mixed support for our hypotheses. First, our findings provide reasonable grounds for the rejection of H1, the opposition hypothesis, which claimed that a persuasion effort would only impact the attitude targeted. Across all 3 pilots and 2 longitudinal experiments, we found that not only was the attitude targeted by a persuasive appeal changed, but that non-targeted attitudes also changed across the belief system. This change was not always large, and only few significant effects survived chance correction. Nonetheless, there was change, which suggests the presence of spillover effects of attitude change on non-targeted attitudes. Second, in line with our distance hypothesis (H2), we found that this change was often moderated by the distance between non-targeted and targeted attitudes within the belief system. In other words, non-targeted attitudes with higher expected impact from the targeted attitude usually showed greater changes than non-targeted attitudes with lower expected impact (i.e., attitudes further away within the network). We found fairly robust support for this hypothesis using our banking manipulation, which generally replicated across Studies 1 and 2. Support using the terrorism manipulation in Study 1 was weaker. The terrorism manipulation supported the idea that distance between attitudes within a belief system mattered, but the pattern did not conform to expectations (i.e., non-targeted attitudes changed in a different direction than expected). The crime manipulation in Study 2 showed support for H2, although the differences between condition at high impact were not as large as expected. Together this provides support for the emergence of dynamic constraint using a network approach to belief systems.

8.1. Theoretical implications

First, this research provides support for the existence of dynamic constraint within belief systems in the mass public. Our results showed that not only did targeted attitudes change following an experimental manipulation, but also some non-targeted attitudes changed too. Of course, observed changes in non-targeted attitudes were often small and somewhat messy, but this is not unexpected given the complexity of dynamic constraint. We would not expect large effects (cf. Brandt & Sleegers, 2021). Moreover, the belief systems approach applied here used the aggregated belief system from the sample because it is not (yet) possible to assess each individuals' belief system, as would be preferred (more on this later). Importantly, despite these limitations, the effects that were observed occurred consistently over three pilots and two longitudinal studies, suggesting that these effects are not just random error.

We think that the changes in non-targeted attitudes and the moderation by impact should be interpreted as support for dynamic constraint because trends in our data arguably satisfy the three key preconditions of dynamic constraint laid out in our introduction. First, our manipulation of the targeted attitudes (banking, terrorism and crime) generated a medium to large effect size of attitude change. We think that this change in an individual's attitude should be enough to create dissonance within the belief system that other attitudes must adjust to. Second, our community analysis found no evidence for H3, instead suggesting that attitude change reached across domains. Specifically, the community analysis clusters attitudes that are densely connected, and is in many ways equivalent to a factor analysis (Christensen & Golino, 2020; Epskamp, Maris, Waldorp, & Borsboom, 2016; Golino & Epskamp, 2017). As such, attitudes that are grouped in the same community may be seen as existing within a shared domain (e.g., terrorism, defense, crime, and immigration attitudes may make a 'threat domain'), while those grouped in separate communities may be seen as existing in different domains. Thus, our finding that attitude change was *not* bound to non-targeted attitudes within the same community as the targeted attitude indicates that cross-domain change was observed. Finally, our explorative analyses showed that change in non-targeted attitudes was related to targeted attitude change. The explorative mediation analysis further showed that using the targeted attitude as a mediator (partially) mediated by change in the targeted attitude. Although mediation analyses have strong assumptions that our study design cannot fully fulfill (e.g., sequential ignorability; Bullock et al., 2010), it does provide evidence consistent with the idea that the attitude change in non-targeted attitudes was transmitted via the attitude targeted by the manipulation. Together, we believe that this provides fair grounds for the claim that the attitude change we observed on non-targeted attitudes can be conceptualized as dynamic constraint.

Dynamic constraint has important implications for the field. It supports a key pillar of belief systems theory, which suggests the existence of causal relations among belief elements. Specifically, our results present evidence that causal relationships between attitudes transmit influence among them, especially to those close together in the belief system. In this case, attitudes largely changed in order to be more similar to each other. So, when banking attitudes became more supportive of increased regulations, support for regulations in big-business also increased. It is notable that our findings were a product of a sample of the general U.S. population, with moderately educated participants. This is not a group that is traditionally considered to have a highly constrained belief system (Malka et al., 2014). This evidence runs somewhat counter to conclusions from Converse (1964) and more recent tests of dynamic constraint (Coppock & Green, 2021; although it is more consistent with the literature on lateral attitude change; Glaser et al., 2015). A likely reason why we find dynamic constraint where others did not is our application of the network method. Our network approach studies belief systems as a descriptive norm – something a group of people have (Cialdini, Reno, & Kallgren, 1990). In other contexts, belief systems have been discussed as an injunctive norm or a system of understanding people should have (Groenendyk et al., 2020), or as more of a latent variable (Coppock & Green, 2021), with expectations about where dynamic constraint should occur often being derived from an elite population (e.g., theoretically motivated). A benefit of our approach is that it takes serious the belief systems that exist in the mass public, in that it takes the associations between attitudes in the sample as leading in the construction of the belief system. In this way, our approach makes fewer assumptions about which attitudes should ideologically relate to one another and, perhaps as a result, is better able to find attitude change through dynamic constraint as a result. This may be one reason why our results are more similar to those obtained in the field of indirect or lateral attitude: This research tends to pre-test relatedness between items to select closer (i.e., more similar) or more distant (i.e., more different) items or attitudes (Bohner et al., 2020; Brannon et al., 2019). As such, scholars can interpret this research as evidence that dynamic constraint is supported when less stringent definitions of belief systems are applied (for a discussion of the value of less stringent definitions for belief systems/ideology, see; Jost, 2006).

Second, this research provides support for a network perspective on belief systems. Our results demonstrate that not only did dynamic constraint occur, but knowing the structure of a belief system (from T1) allowed us to better predict *where* and *how strongly* dynamic constraint occurred. Specifically, our results showed that the distance between non-targeted and targeted attitudes generally moderated change in the non-targeted attitudes (with attitudes closest to the targeted attitude/ with highest expected impact from the targeted attitude often changing the most). This network approach is a strength of this research, bringing together theoretical definitions of belief systems that have long existed in the literature (e.g., focusing on interrelations between belief system components) with a methodology that puts these components center stage. This allowed us to visualize the belief system, including multiple elements and interrelations among them, as they exist in the population, and generate specific (preregistered) predictions about where and how strongly dynamic constraint would occur and in what direction change would occur. Indeed, attitudes changed, becoming more 'similar' to their neighbors, and effects were spread across the belief system as a whole. Moreover, these predictions often provided a useful explanation of observed changes in attitudes nested within individuals. Consistently, across studies, Ising impact emerged as the best approximation of distance within a belief system network, according to the fit indices (AIC/ BIC), offering the most relevant predictions about this attitude change. In this way, results come together with other research (Brandt et al., 2019; Brandt & Sleegers, 2021) to show that a network approach to belief systems provides a useful method for measuring belief systems and assessing change in attitudes within this system (Chambon et al., 2021). Together this can help to better our understanding of belief systems and their consequences for individuals (e.g., persuasion processes).

8.2. Limitations

It is clear that our hypothesized network model did not provide a perfect explanation of attitude change or transmission within belief systems, and some results were somewhat messy. Most notably, the effect of the terrorism manipulation (Study 1) on attitude change was in the opposite direction than was expected (see displacement effects; Glaser et al., 2015). Also, across the board we found small effects of attitude change in non-targeted attitudes, along with a number of minor deviations from the predicted patterns of attitude change (e.g., some attitudes with lower impact had higher average change, as shown in Figs. 5 & 10). It is unclear precisely why these deviations occurred, but at least four explanations are possible. First, our network approach uses an aggregate model. This presents the average belief system in our target population. Although our networks appeared to offer a useful and replicable summary/average (e.g., path strengths and communities were similar; c.f. Brandt, 2020), and there are various benefits to this approach, one limitation is that an individual's own belief system is likely to deviate from this average (cf. ecological fallacy; see also Brandt & Morgan, 2021). As such, it is not surprising that our predictions were imperfect. We would expect more accurate predictions would be derived from individual belief systems. Although the methods exist to generate these, data required is intensive usually requiring approximately 50 time points or more (Bringmann, Ferrer, Hamaker, Borsboom, & Tuerlinckx, 2018), and so was not possible to test at this time. Until such ideal data exists, an alternative approach to this would be to apply recently developed correlational class analyses to test for existence of different structures of belief systems within the sample and group together individuals with similar belief systems (Hunzaker & Valentino, 2019; see also Baldassarri & Goldberg, 2014).

Second, tendencies for attitude change may not be equal across all attitudes – some (e.g., important) attitudes may be harder to change. Attitudes that are resistant to persuasion attempts may cause displacement to occur, where expected non-targeted attitudes do not change, but others attitudes do (as in the terrorism condition, Study 1; Glaser et al., 2015). Thus, an equal unit of dissonance caused by change in a neighboring attitude may have a greater impact on some attitudes rather than others, or result in a different pattern of change than anticipated. Third, our models assume bidirectional relations between attitudes, but it is imaginable that relations could be better conceptualized as one-way. This would mean that influence could flow more easily from terrorism to crime than the reverse. This may explain why some attitudes with higher expected impact showed only small or negligible attitude change. Fourth, despite our large sample, our power may still not have been adequate. Currently, we have little or no insight in the field about how large attitude change needs to be in order to cause dissonance within a belief system big enough to induce dynamic constraint. A single event causing attitude change may not be reliable enough, it could be that longer-term events and more stable change in a targeted attitude is required. This would be an interesting and important question for future research.

Lastly, this research was conducted in one specific U.S. context. Although our theory should apply across different contexts and countries, this would require replication. While our networks were shown to be consistent across different samples from the same target population (which according to demographics tested are fairly representative of the general US population), and networks in general are shown to be replicable (Brandt, 2020), variation is likely to occur between countries and individuals (Brandt et al., 2021).

9. Conclusion

This research presents evidence supporting the existence of a key pillar of belief system theory - dynamic constraint. We found that attitude change on a targeted attitude was transmitted to other non-targeted attitudes, across different domains. Attitude change of non-targeted attitudes not only generally followed the direction of attitude change in the targeted attitude (i.e., supporting increased or reduced spending), but the attitudes that were closest to the targeted attitude generally changed the most strongly. This shows that understanding the structure of the belief system can help to better understand and predict patterns of dynamic constraint. In other words, belief systems research may help further our understanding of how belief change happens and how to make effective persuasion appeals by facilitating a deeper understanding of the attitudes that are nested in a belief system and crucially the structures of relations linking these attitudes.

Acknowledgement

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 759320).

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.jesp.2021.104279.

References

Baldassarri, D., & Goldberg, A. (2014). Neither ideologues nor agnostics: Alternative voters' belief system in an age of partisan politics. *American Journal of Sociology, 120*(1), 45–95.

Bartle, J. (2000). Political awareness, opinion constraint and the stability of ideological positions. *Political Studies, 48*(3), 467–484.

Bates, D., Maechler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software, 67*(1), 1–48. https://doi.org/10.18637/jss.v067.i01

Bian, L., Leslie, S. J., Murphy, M. C., & Cimpian, A. (2018). Messages about brilliance undermine women's interest in educational and professional opportunities. *Journal of Experimental Social Psychology, 76*, 404–420.

Blankenship, K. L., Wegener, D. T., & Murray, R. A. (2012). Circumventing resistance: Using values to indirectly change attitudes. *Journal of Personality and Social Psychology, 103*(4), 606–621.

Blankenship, K. L., Wegener, D. T., & Murray, R. A. (2015). Values, inter-attitudinal structure, and attitude change: Value accessibility can increase a related attitude's resistance to change. *Personality and Social Psychology Bulletin, 41*(12), 1739–1750.

Bohner, G., Elleringmann, L., Linne, R., Boege, R., & Glaser, T. (2020). Lateral attitude change: Does acceptance versus rejection of focal change cause generalization versus displacement? *Research Report.*. https://doi.org/10.4119/unibi/2941633

van Borkulo, C., & Epskamp, S. (2016). IsingFit: Fitting ising models using the elasso method. R package version 0.3.1 https://CRAN.R-project.org/package=IsingFit.

Boutyline, A., & Vaisey, S. (2017). Belief network analysis: A relational approach to understanding the structure of attitudes. *American Journal of Sociology, 122*(5), 1371–1447.

Brandt, M. J. (2020). Estimating and examining the replicability of belief system networks. *Collabra: Psychology, 6*(1).

Brandt, M. J., & Morgan, G. S. (2021). Between-person methods provide limited insight about within-person belief systems. *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*(9), 1352–1364.

Brandt, M. J., & Sleegers, W. W. A. (2021). Evaluating belief system networks as a theory of political belief system dynamics. *Personality and Social Psychology Review, 25*, 159–185.

Brandt, M. J., Turner-Zwinkels, F. M., Karapirinler, B., van Leeuwen, F., Bender, M., van Osch, Y., & Adams, B. G. (2021). The association between threat and politics simultaneously depends on the type of threat, the political domain, and the country. *Personality and Social Psychology Bulletin, 47*(2), 324–343.

Brannon, S. M., DeJong, A., & Gawronski, B. (2019). Determinants of lateral attitude change: The roles of object relatedness, attitude certainty, and moral conviction. *Social Cognition, 37*(6), 624–658.

Bringmann, L. F., Ferrer, E., Hamaker, E. L., Borsboom, D., & Tuerlinckx, F. (2018). Modeling nonstationary emotion dynamics in dyads using a time-varying vector-autoregressive model. *Multivariate Behavioral Research, 53*(3), 293–314.

Bullock, J. G., Green, D. P., & Ha, S. E. (2010). Yes, but what's the mechanism?(don't expect an easy answer). *Journal of Personality and Social Psychology, 98*(4), 550–558.

Carpini, D. M., & Keeter, S. (1993). Measuring political knowledge: Putting first things first. *American Journal of Political Science, 37*(4), 1179.

Chambon, M., Dalege, J., Waldorp, L., van der Maas, H., Borsboom, D., & van Harreveld, F. (2021). *A complex systems perspective on compliance with behavioral measures during the COVID-19 pandemic in the Netherlands: How psychological networks can inform interventions*. Manuscript in preparation.

Champely, S. (2020). pwr: Basic functions for power analysis. R package version 1.3-0 https://CRAN.R-project.org/package=pwr.

Christensen, A. P., & Golino, H. (2020). *Statistical equivalency of factor and network loadings*.

Cialdini, R. B., Reno, R. R., & Kallgren, C. A. (1990). A focus theory of normative conduct: recycling the concept of norms to reduce littering in public places. *Journal of Personality and Social Psychology, 58*(6), 1015.

Cooper, J. (2019). Cognitive dissonance: Where we've been and where we're going. *International Review of Social Psychology, 32*(1).

Coppock, A., & Green, D. P. (2021). Do belief systems exhibit dynamic constraint? *Journal of Politics*. https://doi.org/10.1086/716294

Converse, P. E. (1964). In D. Apter (Ed.), *The nature of belief systems in mass publics*. New York: Free Press: Ideology and discontent.

Csardi, G., & Nepusz, T. (2006). The igraph software package for complex network research. *InterJournal, Complex Systems, 1695*(5), 1–9.

Dalege, J., Borsboom, D., van Harreveld, F., van den Berg, H., Conner, M., & van der Maas, H. L. (2016). Toward a formalized account of attitudes: The Causal Attitude Network (CAN) model. *Psychological Review, 123*(1), 2–22.

Dalege, J., Borsboom, D., van Harreveld, F., & van der Maas, H. L. (2017). Network analysis on attitudes: A brief tutorial. *Social Psychological and Personality Science, 8*(5), 528–537.

Dalege, J., Borsboom, D., van Harreveld, F., Waldorp, L. J., & van der Maas, H. L. (2017). Network structure explains the impact of attitudes on voting decisions. *Scientific Reports, 7*(1), 1–11.

Eadeh, F. R., & Chang, K. K. (2020). Can threat increase support for liberalism? New insights into the relationship between threat and political attitudes. *Social Psychological and Personality Science, 11*, 88–96.

Epskamp, S., Borsboom, D., & Fried, E. I. (2018). Estimating psychological networks and their accuracy: A tutorial paper. *Behavior Research Methods, 50*, 195–212.

Epskamp, S., Cramer, A., Waldorp, L. J., Schmittmann, V. D., & Borsboom, D. (2012). Qgraph: Network visualizations of relationships in psychometric data. *Journal of Statistical Software, 48*(4), 1–18. URL http://www.jstatsoft.org/v48/i04/.

Epskamp, S., & Fried, E. I. (2018). A tutorial on regularized partial correlation networks. *Psychological Methods, 23*(4), 617–634.
Epskamp, S., Maris, G. K., Waldorp, L. J., & Borsboom, D. (2016). Network psychometrics. In *Handbook of psychometrics*. New York: Wiley.
Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods, 39*, 175–191.
Federico, C. M., & Schneider, M. C. (2007). Political expertise and the use of ideology: Moderating effects of evaluative motivation. *Public Opinion Quarterly, 71*(2), 221–252.
Festinger, L. (1957). *A theory of cognitive dissonance* (Vol. 2). Stanford university press.
Fishman, N., & Davis, N. T. (2019). Change we can believe. In *Structural and content dynamics within belief networks*. Pre-print.
Foygel, R., & Drton, M. (2010). Extended Bayesian information criteria for Gaussian graphical models. In *Advances in neural information processing systems* (pp. 604–612).
Freeder, S., Lenz, G. S., & Turney, S. (2018). The importance of knowing 'what goes with what'. *The Journal of Politics, 81*(1), 274–290.
Freeze, M., & Montgomery, J. M. (2016). Static stability and evolving constraint: Preference stability and ideological structure in the mass public. *American Politics Research, 44*(3), 415–447.
Gerring, J. (1997). Ideology: A definitional analysis. *Political Research Quarterly, 50*(4), 957–994.
Glaser, T., Dickel, N., Liersch, B., Rees, J., Süssenbach, P., & Bohner, G. (2015). Lateral attitude change. *Personality and Social Psychology Review, 19*(3), 257–276.
Golino, H. F., & Epskamp, S. (2017). Exploratory graph analysis: A new approach for estimating the number of dimensions in psychological research. *PLoS One, 12*(6), Article e0174035.
Groenendyk, E. W., Kimbrough, E. O., & Pickup, M. (2020). *How norms shape the nature and origins of mass belief systems*. Available at SSRN 3541289.
Heider, F. (1946). Attitudes and cognitive organization. *The Journal of Psychology, 21*(1), 107–112.
Hopkins, D. J., & Mummolo, J. (2017). Assessing the breadth of framing effects. *Quarterly Journal of Political Science, 12*(1), 37–57.
Hunzaker, M. F., & Valentino, L. (2019). Mapping cultural schemas: From theory to method. *American Sociological Review, 84*(5), 950–981.
Johnston, C. D., & Ollerenshaw, T. (2020). How different are cultural and economic ideology? *Current Opinion in Behavioral Sciences, 34*, 94–101.
Jost, J. T. (2006). The end of the end of ideology. *American Psychologist, 61*(7), 651–670.
Judd, C. M., & Downing, J. W. (1990). Political expertise and the development of attitude consistency. *Social Cognition, 8*(1), 104–124.
Kalmoe, N. P. (2020). Uses and abuses of ideology in political psychology. *Political Psychology, 41*, 771–793.
Keating, D. M., & Bergan, D. E. (2017). Mapping Political Attitudes: The impact of concept mapping on ideological constraint. *Communication Studies, 68*(4), 439–454.
Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). LmerTest package: Tests in linear mixed effects models. *Journal of Statistical Software, 82*(13), 1–26. https://doi.org/10.18637/jss.v082.i13
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*(6), 1031–1051.
McGuire, W. J. (1981). The probabilogical model of cognitive structure and attitude change. In R. E. Petty, T. M. Ostrom, & T. C. Brock (Eds.), *Cognitive responses in Persuasion* (pp. 291–307). New Jersey: Erlbaum, Hillsdale.
McGuire, W. J. (1990). Dynamic operations of thought systems. *American Psychologist, 45*(4), 504–512.
Monroe, B. M., & Read, S. J. (2008). A general connectionist model of attitude structure and change: The ACS (Attitudes as Constraint Satisfaction) model. *Psychological Review, 115*(3), 733–759.
Nilsson, A., Bergquist, M., & Schultz, W. P. (2017). Spillover effects in environmental behaviors, across time and context: a review and research agenda. *Environmental Education Research, 23*(4), 573–589.
Peffley, M., & Hurwitz, J. (1992). International events and foreign policy beliefs: Public response to changing Soviet-US relations. *American Journal of Political Science*, 431–461.
Peffley, M. A., & Hurwitz, J. (1985). A hierarchical model of attitude constraint. *American Journal of Political Science*, 871–890.
Phua, D. Y., Leong, C. H., & Hong, Y. (2020). Heterogeneity in national identity construct: Example of Singapore using network analysis. *International Journal of Intercultural Relations, 78*, 20–32.
R Core Team. (2019). ***R: A language and environment for statistical computing***. Vienna, Austria: R Foundation for statistical computing. URL https://www.R-project.org/.
Sayans-Jiménez, P., van Harreveld, F., Dalege, J., & Rojas Tejada, A. J. (2019). Investigating stereotype structure with empirical network models. *European Journal of Social Psychology, 49*(3), 604–621.
Shils, E. A. (1968). The end of ideology? In C. Waxman (Ed.), *The end of ideology debate* (pp. 49–63). New York: Simon & Schuster (Original work published 1955).
Snijders, T. A., & Bosker, R. J. (2011). *Multilevel analysis: An introduction to basic and advanced multilevel modeling*. London: Sage.
Stimson, J. A. (1975). Belief systems: Constraint, complexity, and the 1972 election. *American Journal of Political Science*, 393–417.
Tedin, K. L. (1987). Political ideology and the vote. *Research in Micropolitics, 2*, 63–94.
Turner-Zwinkels, F. M., Johnson, B. B., Sibley, C. G., & Brandt, M. J. (2020). Conservatives' Moral Foundations are more Densely Connected than Liberals' Moral Foundations. *Personality and Social Psychology Bulletin, 47*(2), 167–184, 0146167220916070.
Verschoor, M., Albers, C., Poortinga, W., Böhm, G., & Steg, L. (2020). Exploring relationships between climate change beliefs and energy preferences: A network analysis of the European Social Survey. *Journal of Environmental Psychology, 70*. doi: 101435.
Whitmarsh, L., & O'Neill, S. (2010). Green identity, green living? The role of pro-environmental self-identity in determining consistency across diverse pro-environmental behaviours. *Journal of Environmental Psychology, 30*(3), 305–314.
Wickham, H. (2016). *Ggplot2: Elegant graphics for data analysis*. New York: NY: Springer-Verlag.
Wilke, C. O. (2019). Cowplot: streamlined plot theme and plot annotations for 'ggplot2'. R package version 1.0.0 https://CRAN.R-project.org/package=cowplot.
Zhang, R. J., Liu, J. H., Brown, G., & Gil de Zúñiga, H. (2020). A network analysis of Global Trust across 11 democratic countries. *International Journal of Public Opinion Research, 33*(1), 147–158.