Do Belief Systems Exhibit Dynamic Constraint?

Alexander Coppock and Donald P. Green*

July 28, 2020

Abstract

As described in Converse (1964), belief systems are dynamically constrained if a change in one opinion causes a concomitant change in a related opinion. While an enormous literature is dedicated to the study of static constraint (the extent to which individuals hold political views that "go together"), dynamic constraint is rarely studied, especially using experimental research designs. We offer a new formalization of the theoretical argument that suggests an identification strategy for detecting dynamic constraint. We present evidence from survey experiments conducted with convenience samples of both the mass public and of political elites. Our results indicate that even among respondents whose belief systems are highly constrained in the static sense, a change in one attitude need not precipitate changes in related attitudes. These experimental results affirm and extend Converse's thesis about the limited extent of dynamically constrained ideological thinking in the mass public. The lack of dynamic constraint among our elite sample raises the question of how they come to hold political opinions that are constrained in a static sense. We present an experiment that suggests a potential explanation: elites may be more likely to be chided for expressing inconsistent positions.

*Alexander Coppock is Assistant Professor of Political Science, Yale University. Donald P. Green is J.W. Burgess Professor of Political Science, Columbia University. This research was reviewed and approved by the Institutional Review Board of Columbia University (IRB-AAAP1312 and IRB-AAAP9305).

A half-century after its publication, Philip Converse's (1964) essay "The Nature of Belief Systems in Mass Publics" remains at the forefront of public opinion scholarship. Its provocative thesis is that unlike "elites" (elected officials and party activists), members of the "mass public" (ordinary voters) typically hold political opinions that are a loosely structured mélange of liberal and conservative ideas. Members of Congress, for example, hold opinions that predictably reflect the worldview of either the left or right (Bafumi and Herron 2010). Conservatives in Congress resist social entitlement programs, support law-and-order policies, and seek to promote patriotism and traditional social values. Congressional liberals, by contrast, seek to expand the social safety net, promote the rights and well-being of disadvantaged groups, express concern about excessive police or military force, and challenge traditional values and the hierarchies that they engender. Among voters, these liberal or conservative ingredients are often jumbled together. On the whole, the public appears relatively moderate when their policy views are averaged across different issue domains, but as Ahler and Broockman (2018) point out, citizens who are centrist on average tend to support many policies that are not moderate.

Converse coined the term "constraint" to describe the interconnectedness of different political opinions. In a widely-quoted passage (p.2), he defines a "belief system" as "a configuration of ideas and attitudes in which the elements are bound together by some form of constraint or functional interdependence." The coherence of a belief system may be assessed at a given point in time or by tracing its elements over time. In reference to the former, Converse (p.3) uses the term "static" constraint "to mean the success we would have in predicting, given initial knowledge that an individual holds a specified attitude, that he holds certain further ideas and attitudes." Converse's empirical assessment of static constraint demonstrates the "fragmentation" and "narrowness" of mass belief systems (p.54) by showing the weak correlations between survey responses to questions about assorted policy topics (p.31).

In the decades since the publication of Converse's essay, static constraint has remained the focus of the vast empirical literature on belief systems.[1] Scholars have pointed out that static

[1] In addition to static constraint, another of Converse's methods for assessing ideological awareness is use of liberal and conservative parlance, which has also generated a large literature on the meaning and measurement of liberal-conservative self-identification (Jacoby 1991; Bauer et al. 2017)

constraint in the mass public may be underestimated when gauged by simple correlations (Achen 1975; Ansolabehere, Rodden and Snyder 2008), that constraint appears to operate more strongly now than in Converse's day (Freeze and Montgomery (2016) and Abramowitz and Saunders (2008), but see Jewitt and Goren (2016), who find that increases occurred only among the most politically interested and active), and that mass constraint may be higher outside the United States (Converse and Pierce 1986; Malka, Lelkes and Soto 2017). Nevertheless, it remains the case that in both absolute terms and in comparison to elites, ordinary voters hold opinions that are weakly correlated across different policy domains. In a recent analysis, Freeder, Lenz and Turney (2019) estimate that only approximately 20 to 40% of the U.S. mass public hold opinions that are constrained in the static sense.

Much less attention has been paid to what Converse calls "dynamic" constraint or "interdependence." This characteristic of belief systems concerns "the probability that a change in the perceived status (truth, desirability, and so forth) of one idea-element would psychologically require, from the point of view of the actor, some compensating change(s) in the status of idea-elements elsewhere in the configuration." (p.3) Converse contends that dynamic constraint operates weakly if at all among ordinary voters, who tend to acquire new policy opinions via "social" transmission (p.8) rather than through a principled reconciliation of new and old opinions. As a consequence, mass belief systems tend to span a narrow range of policy domains (p.5), too narrow to imbue election outcomes with meaning as ideological mandates (pp.58-64).

The data requirements for assessing *dynamic* constraint are more demanding because the measurement process requires an initial change in an "idea-element" whose ramifications can be traced within and across policy domains. In order to generate statistically meaningful conclusions, the initial change must be genuine (not a sampling fluke), sizable (so that its ramifications can be detected with reasonable precision), and brought about by domain-specific causes (so that concomitant changes in policy views across domains cannot be attributed to environmental forces that *directly* affect opinions in multiple domains). Converse (1964, 1970) did not offer direct evidence of this kind. In those rare instances where other scholars have assessed the degree of dynamic constraint in mass opinion, they have leveraged seismic political events, such as the abrupt change in U.S.-Soviet relations under Perestroika (Peffley and Hurwitz 1992), the 9/11 terrorist attacks (Kinder and Kam 2010), or the tax revolt of the late 1970s and early 1980s (Sears and Citrin 1982). Of these studies, only Peffley and Hurwitz (1992) both addressed the topic of dynamic constraint directly and used panel data to track individual attitudes over time.

Two important gaps persist in the literature on dynamic constraint. The first is a lack of formalization. Converse devotes just a few sentences to the topic of dynamic constraint and does not lay out an identification strategy that would allow researchers to detect it empirically. One contribution of the current paper is to present an array of psychological models that exhaust all of the ways that an exogenous shift in a given idea-element might affect other idea-elements. This exercise provides a mapping between theory and data, allowing us to identify patterns of empirical results that would be consistent with psychological processes that give rise to dynamic constraint.

The second is a dearth of experiments designed to introduce these exogenous shifts in a given idea-element. The before-after designs associated with terrorist attacks, Perestroika, or the Tax Revolt are suggestive but suffer the obvious limitation that events neither unfold randomly nor unfold in a way that targets just one specific idea-element. The advantage of deploying specifically-crafted interventions randomly and under controlled conditions is that they lead to much more confident conclusions about whether and to what degree exogenous changes in one idea-element set in motion changes in other attitudes. Several recent studies have contributed useful experimental evidence of this kind, and our paper builds on these contributions. In particular, Hopkins and Mummolo (2017) assesses the "breadth" of framing effects by randomly exposing subjects to short arguments made by U.S. senators, then measuring subjects spending preferences in both target and nontarget domains. For the most part, arguments against wasteful spending in one domain affected target spending preferences and but not spending preferences in other domains. Similarly, Trump and White (2018) find that information that changes survey respondents' beliefs about the degree of income inequality has little downstream effect on whether they perceive the economic system to be unfair. Other recent examples include Finseraas and Kotsadam (2017), who find that randomly induced exposure to immigrants reduces negative stereotypes about them but does not change policy preferences regarding immigration or the social welfare benefits to which they are entitled, and Peyton (2020), who shows that successful experimental manipulations of trust in government do not affect preferences for redistribution. One interesting exception to this pattern of narrow and ephemeral change is the finding by Broockman and Kalla (2016) that doorstep conversations between canvassers and voters on the subject of transgender rights affected not only policy opinions on that issue but also on the issue of gay rights. Although these studies do not explicitly address the topic of dynamic constraint, we believe that they are quite relevant and provide an important foundation for further investigation.

Our empirical contribution is to offer novel experimental results for both mass and elite samples. Consistent with Converse's argument about the narrowness of mass belief systems, our experiments show that change in one idea-element tends to precipitate little or no change in other cognate issue domains. Taken together, our experiments underscore Converse's contention that apparent static constraint may actually exaggerate the functional interdependence of a belief system. Even among elite respondents, who display a higher degree of consistency in their policy opinions across domains, opinion change occurs in a localized and fragmented fashion. These results suggest a hitherto unappreciated puzzle: If elites tend not to display dynamic constraint, how did they come to hold correlated attitudes that imply static constraint? We speculate, based on experimental results suggesting that respondents can be chided into expressing more consistent views, that the answer may lie in the environmental forces that monitor and reward elites' consistency.

Formalizing the Empirical Implications of Dynamic Constraint

Unlike constructs such as liberalism-conservatism or political knowledge, dynamic constraint is not a psychological trait that lends itself to measurement via standard psychometric tools. We cannot simply quiz people repeatedly about their opinions or beliefs in order to zero in on the latent trait of interest. Instead, dynamic constraint must be detected by tracing the ripples of some exogenous source of opinion change. For example, if a person becomes convinced that government spending is rife with corruption and waste, does she become more supportive of tax cuts designed to trim the fat in government (Sears and Citrin 1982)? The range of the belief system that guides a person's thinking is ultimately an empirical question. Converse's reading of survey data suggested that the average voter's belief system was quite localized and perhaps limited to topics such as race. Lane's (1962) depth interviews led him to conclude that few ordinary voters held belief systems structured along liberal-conservative lines but nevertheless argued that his respondents' idiosyncratic belief systems were broad-ranging. For our purposes, the range and structure of belief systems is suggested by the distance between the initial locus of change and responses to questions about less directly related topics.

Although evidence of such ripple effects would seem to suggest a dynamic process by which people maintain consistency among the various components of a far-reaching belief system, this interpretation hinges on additional assumptions. The purpose of this section is to formalize these assumptions so as to clarify the connection between the underlying theoretical process and observable patterns that could emerge from a randomized experiment. Doing so not only tells us what kinds of experiments might be informative; it also alerts us to the limits of what can be learned from such experiments.

We posit a set of four variables, only some of which are directly observed:

- I (unobserved). A belief system, or the set of functional interdependencies that may influence political attitudes. This belief system could conceivably be a tightly organized set of ideological principles or something much narrower, perhaps emanating from a set of group interests (e.g., a concern for the well-being of small business owners).
- Y_A (measured). Attitude A, as measured by the response to a survey question on issue A.
- Y_B (measured). Attitude B, as measured by the response to a survey question on issue B.
- Z_A (manipulated or set by nature). Information A, which is a causal factor that is directly relevant to issue A but not directly relevant to issue B. Z_A is the crucial factor that may set in motion the change in Y_A that may or may not be accompanied by a change in Y_B.

Notably absent from this list are the set of other unobserved factors that may influence the belief system, attitude A, or attitude B, such as background beliefs, demographics, and culture. The implications of the model do not hinge on whether these unobserved initial conditions result in the belief system or the intercorrelation of attitudes.

To this set of four variables, we apply a set of assumptions that are equivalent to a family of exclusion restrictions:

Assumption 1 (Exogeneity). We assume that Information A (Z_A) has no ancestors, i.e., it is randomly assigned or is otherwise exogenous. No paths lead from I, Y_A, or Y_B to Z_A.

Assumption 2 (Acyclicity) We assume a particular temporal ordering: A belief system (I) may affect attitudes (Y_A and Y_B) but attitudes do not affect the belief system. Y_A may affect Y_B but not vice-versa. This choice amounts to requiring that Z_A affect one attitude "first," and this ordering is reflected in the "A" subscript on the Z_A variable.

Assumption 3 (Functional Interdependence). A belief system (I) may or may not influence attitudes, but if it affects one attitude, it affects both. Suppose that, contrary to assumption 3, we allowed a belief system to affect one attitude but not the other. This belief system would then be too narrow to merit the name; instead it would be one of the many unobserved factors that independently set the level of one attitude or the other.

Figure 1: Model of Attitude Change with All Paths Consistent with Assumptions 1-3

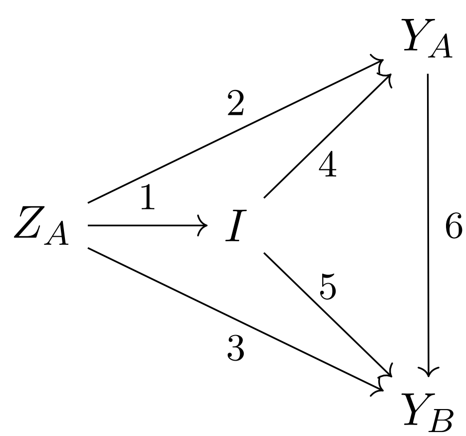

These assumptions can be used to zero in on the observable implications of dynamic constraint. We operationalize Converse's "functional interdependence" by defining dynamic constraint as an effect of Z_A that is mediated by I. For example, if ideology or group affinities are the *mechanism* by which Z_A affects attitudes, attitudes are dynamically constrained. If Z_A affects outcomes through pathways that do not include I (i.e., through direct paths only), attitudes are not dynamically constrained because there is no sense in which the idea-elements are bound together.[2]

Figure 1 represents a causal graph in which our four variables are interrelated by six direct paths. No other paths could be added to this figure without violating some or all of assumptions 1-3. The figure represents a model of attitude change in which dynamic constraint plays a role in shaping attitudes. Paths lead from Z_A to I and from I to Y_A and Y_B. Figure 1 is one of 32 possible configurations of paths that are consistent with assumptions 1-3. Each configuration of paths coincides with a distinct theory of how exogenous forces and belief systems interact to induce changes in opinions. The other 31 graphs have the same nodes but a different combination of paths connecting them. We arrive at 32 possible graphs because there 6 total paths, but by assumption 2, paths 4 and 5 are either both present or both absent: $2^5 = 32$. Some of these graphs include paths from Z_A to Y_A and Y_B that pass through I, some do not.[3]

In practice, two empirical strategies are used to draw inferences from putatively exogenous inputs such as information, argumentation, frames, and the like. The first is to track non-experimental changes in these inputs over time (e.g., due to a major political event) or across individuals (e.g., across segments of the public that are or are not aware of such an event). The problem with over time comparisons is that sudden events are often "compound treatments" in the sense that they bear on many policy domains simultaneously and could well influence both Y_A and Y_B directly. Comparisons across people with different levels of exposure to a given input pose problems of omitted variables bias, since we cannot tell whether different levels of exposure influence opinions or merely reflect unobserved attributes that are correlated with different opinions. The inferential challenges of omitted variables bias can be sidestepped by randomly assigning exposure to the exogenous intervention, which is the main reason we turn to randomized experimentation to study

[2] Although some combinations of pathways might be characterized as imposing "vertical" as opposed to "horizontal" constraint insofar as they induce change by invoking core values (Pollock, Lilie and Vittes 1993), we remain agnostic about this distinction. Converse does not use these terms, and either process could contribute to dynamic constraint.

[3] Appendix Table A.1 enumerates the 32 possible graphs according to the presence or absence of the paths shown in Figure 1. For example, model 1 in the first row of the table has paths from Z_A to I and Y_A but excludes all other direct relationships. Model 10 has all possible paths and is the model represented in Figure 1. In the appendix, we also provide a slight elaboration of our framework in which attitudes could affect the belief system, which then in turn influences other attitudes. The conclusions we draw from our data analysis do not depend on which version of the theory we employ, but we present both for completeness.

dynamic constraint.

Randomized experiments do not entirely resolve the empirical challenge of identifying the pathways through which an intervention affects outcomes. Causal mediation analysis – whether performed on observational or experimental data – relies on unverifiable assumptions and is prone to bias if they are not met (Green, Ha and Bullock 2010). One crucial assumption is the causal ordering. We must be correct about which variables are mediators and which variables are outcomes. If the variables are out of order, the estimates from the causal mediation analysis may be highly misleading. If we grant the causal model is correct, however, we still need to rely on an assumption of sequential ignorability (see Imai, Tingley and Yamamoto 2013, for a formal definition.), which requires both that the treatment be as-if-randomly assigned and that the level of mediators within each treatment be as-if randomly assigned. This latter assumption cannot be assured by design and is indeed implausible in our setting. I may emanate from unobserved factors that also affect attitudes Y_A and Y_B. For this reason, we adopt an alternative identification strategy for studying mediation. Rather than attempt to measure the indirect effects of I through causal mediation analysis, we enumerate the models that are consistent with the data and show that none of them implies a mediating role for I.

In the empirical sections below, we will report on a series of studies in which some exogenous event Z_A occurs and we measure subjects' attitudes Y_A and Y_B. As noted above, I is unmeasured. Our studies can turn out one of four ways: We could observe that Z_A affects Y_A, Z_A affects Y_B, both, or neither. Which theories of attitude change could have generated each pattern of evidence?

Our results come closest to the scenario in which Z_A affects Y_A only. For this reason, we explore the models consistent with that pattern of results in more detail in Figure 2. All three models are small variations on a theme: Z_A affects Y_A *directly*, i.e., without being mediated by a belief system. If it is the case that the effects of Z_A on Y_A are not mediated by I, we can conclude that subjects' belief systems are not dynamically constrained. If we do find a pattern of evidence that shows Z_A affects Y_A only, we will not be able to adjudicate among these three models.[4]

[4]An intriguing possibility is that different models may hold for different individuals. In model 2, the belief system still plays a role in constraining attitudes, but *not* in a dynamic fashion. In models 1 and 3, I plays no such role, and the only difference is whether or not Z influences this ineffectual belief system. Some individuals exhibit a high level of *static* constraint. Model 2 might do the best job of explaining their attitudes, while models 1 and 3 might

Figure 2: Models Consistent with Data Showing Effects of Information on Attitude A Only

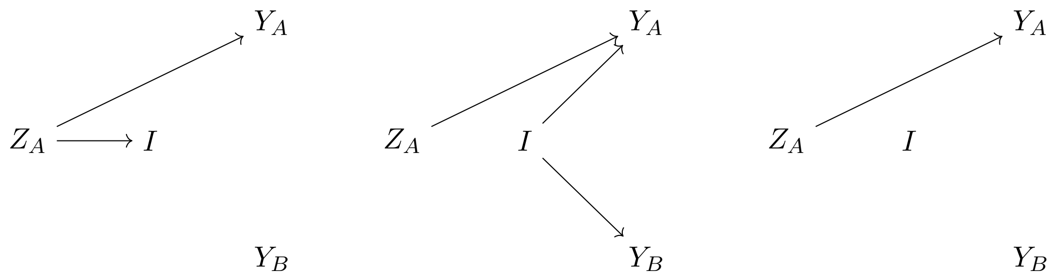

(a) Model 1 (b) Model 2 (c) Model 3

The experiments to follow will generate direct measures of Z_A, Y_A, and Y_B, but I is not directly observed. Nevertheless, this theoretical exercise shows how the data generated by these experiments would map on to the substantive quantities of interest. If we observe effects of our treatments on multiple attitudes, we cannot come to firm conclusions about the extent to which attitudes are dynamically constrained. If, on the other hand, we observe effects on target issues only, we may infer that the belief system does not mediate these effects and hence, attitudes are *not* dynamically constrained.

This identification approach departs markedly from the existing literature. Quite a bit of research on mass opinion is concerned with *scaling* subjects on one or possibly more ideological dimensions. Some approaches include asking subjects to place themselves on scales that range from extremely liberal to extremely conservative (see Kinder and Kalmoe 2017, Appendix A) Other approaches estimate subjects' latent ideological ideal points by positing an underlying factor structure, then fitting that model to subject responses to policy questions (e.g., Table 5 in Jewitt and Goren 2016). The I in our causal diagram is not posited with an eye toward direct measurement. We deliberately avoid imposing structure on I, which is unobserved in our studies. Second, in contrast to studies of static constraint, our estimand is not the correlation between policy attitudes. Instead, we seek to estimate the extent to which interventions that affect a target attitude also affect nontarget attitudes.

Following Converse, we look for dynamic constraint mindful of two potential sources of effect

best explain the attitudes of those exhibiting low levels of static constraint.

heterogeneity. The first is the distinction between elites and non-elites, which looms large in Converse's characterization of political belief systems. Because elites understand "what goes with what" and operate at a "level of conceptualization" that allows them to see the ideological links between policy issues, we expect elites to display more dynamic constraint than their non-elite counterparts. Study 1, therefore, examines not only the degree of dynamic constraint but also whether dynamic adjustments are detectably greater among elites. The second moderator is the salience of what Converse dubs the "logical" and "psychological" connections between issues. Issues involving race, in his assessment, lend themselves to constraint because it is easy for most people to recognize the winners and losers from different policies. In a similar vein, Hopkins and Mummolo (2017) predict that spillover effects across issues will depend on the conceptual distance between the issue content of the framing treatment and the nontarget issue. The studies below attempt in different ways to vary the strength of apparent connections by looking within and across issue domains. Study 1 considers a broad array of issues, whereas study 2 considers dynamic adjustment when a given consideration is applied in closely connected domains. Study 3 goes so far as to call respondents' attention to the connections between issue domains in an effort to facilitate dynamic adjustment.

Study 1: Four Tests of Dynamic Constraint with Mass and Elite Subject Pools

The empirical strategy of this paper requires strong and statistically robust "first-stage" effects in which interventions directly affect attitudes in a given domain. For this reason, we make use of existing studies that convincingly demonstrate direct effects. Our aim is not merely to replicate the original findings but rather to repurpose the data in order to examine the downstream effects of the experimental intervention on opinions in other domains.

We begin with a survey experimental investigation of the effects of newspaper opinion pieces.[5]

[5] The results of the MTurk and Elite versions of this experiment on target attitudes were first described in [citation withheld]. Here we focus on the effects of the manipulations on nontarget attitudes, which (as described in the previous section) provides the crucial test of dynamic constraint.

The experiment was conducted three times: first with a sample of 3,001 Mechanical Turk (MTurk) respondents, again with a sample of 2,181 "elite" respondents, and a final time with 2,524 Lucid respondents (Coppock and McClellan 2019) as a preregistered, out-of-sample test suggested during the review process.

The elite sample was constructed from mailing lists of Capitol Hill staffers, journalists, political professionals and financial industry workers. While we do not have direct measures of how much more ideologically sophisticated the elite sample is than the mass public samples, we do have information on their education and strength of partisanship. A full 96% of the elite sample holds a college degree, compared with 51% of the MTurk sample and 44% of the Lucid sample; 71% of the elites holds a graduate degree, compared with 13% for the MTurk sample and 15% for Lucid. Nearly half – 44% – of the elite sample describes themselves as "strong" partisans, but only a quarter of the MTurk sample chooses these most extreme options. Lucid is more similar to the elite sample than MTurk on this dimension at 46% strong partisans. The most distinctive attribute of the elite sample is its extremely high level of engagement with political news: 97% say they read the news at least daily, and 64% say they talk about politics with their family, friends, or coworkers at least once per day.[6]

All three sets of respondents were randomly assigned to read one of four op-eds or a control condition, before answering a series of policy attitude questions related to each of the four op-eds.[7] The treatment op-eds were all actual opinion pieces that were published in national outlets (the *New York Times*, the *Wall Street Journal*, *USA Today*, and *Newsweek*) and advocated libertarian policy positions on transportation infrastructure, the flat tax, veterans' healthcare, and the financial industry. Accordingly, the outcome variables are all scored so that higher values indicate higher agreement with the libertarian position.[8] Each treatment op-ed is associated with a "target" attitude that is measured with four outcome questions. For a visualization of the design of Study 1, see Figure 3. It describes how respondents in each sample were randomly assigned to the five

[6]See the supplemental materials for additional descriptive information about all three samples.

[7]In the MTurk and Lucid versions of the experiment, subjects could also be assigned to a fifth op-ed on climate change. We omit this treatment arm for clarity of presentation, but the substantive results we find for the other treatments are strongly paralleled in the climate change arm.

[8]See the appendix for the full text of the op-eds and all outcome questions.

experimental conditions before answering four attitudinal questions in each issue area, for a total of 16 outcomes. Within each treatment-control comparison, we can assess the effects of treatment on target and nontarget issues.

The op-eds themselves took up a diverse set of topics. The common thread across all four op-eds is criticism of wasteful government spending, bureaucratic inefficiency, and distortion of markets. The Amtrak op-ed concludes "This country doesn't need more infrastructure that it can't afford to maintain. Instead, it needs a more reliable system of transport funding, and that means one based on user fees and not tax subsidies." In his op-ed in support of the flat tax, Senator Rand Paul contends that "The tax code has grown so corrupt, complicated, intrusive and antigrowth that I've concluded the system isn't fixable." The veterans piece opens with a description of how poorly government delivers healthcare: "More than 57,000 veterans have been waiting at least three months for a doctor's appointment. Another 64,000 never even made it onto a waiting list. There are allegations that waits for care either caused or contributed to veterans' deaths." The op-ed defending Wall Street bankers closes with a call for further financial deregulation: "We could eliminate regulations that crowd out competition for the big banks." If indeed these op-eds were to affect attitudes on nontarget attitudes, we suspect that they would do so by eroding confidence in government's ability to deliver public goods and services efficiently and to create the conditions for the free market to deliver private goods and services. In this setting "functional interdependence" would mean applying criticisms of government wastefulness and market interference from one domain to another.

We first turn to the levels of *static* constraint, or the extent to which attitude levels are correlated. Consistent with a large empirical literature on the differing levels of static constraint among elites and ordinary voters, elites' attitudes are more strongly correlated than MTurk or Lucid respondents' attitudes. Table 1 shows the average Pearson correlations for all pairwise combinations of the 16 outcome questions (16 choose 2 = 120 pairs). We distinguish pairs of outcomes within the same outcome domain from pairs that span outcome domains. In the Elite sample, the average within domain correlation was 0.39, while the average across domains was 0.22. On MTurk, the average correlation between attitudes in the same domain was smaller at 0.29, and the average correlation

Figure 3: Study 1 Experimental Design

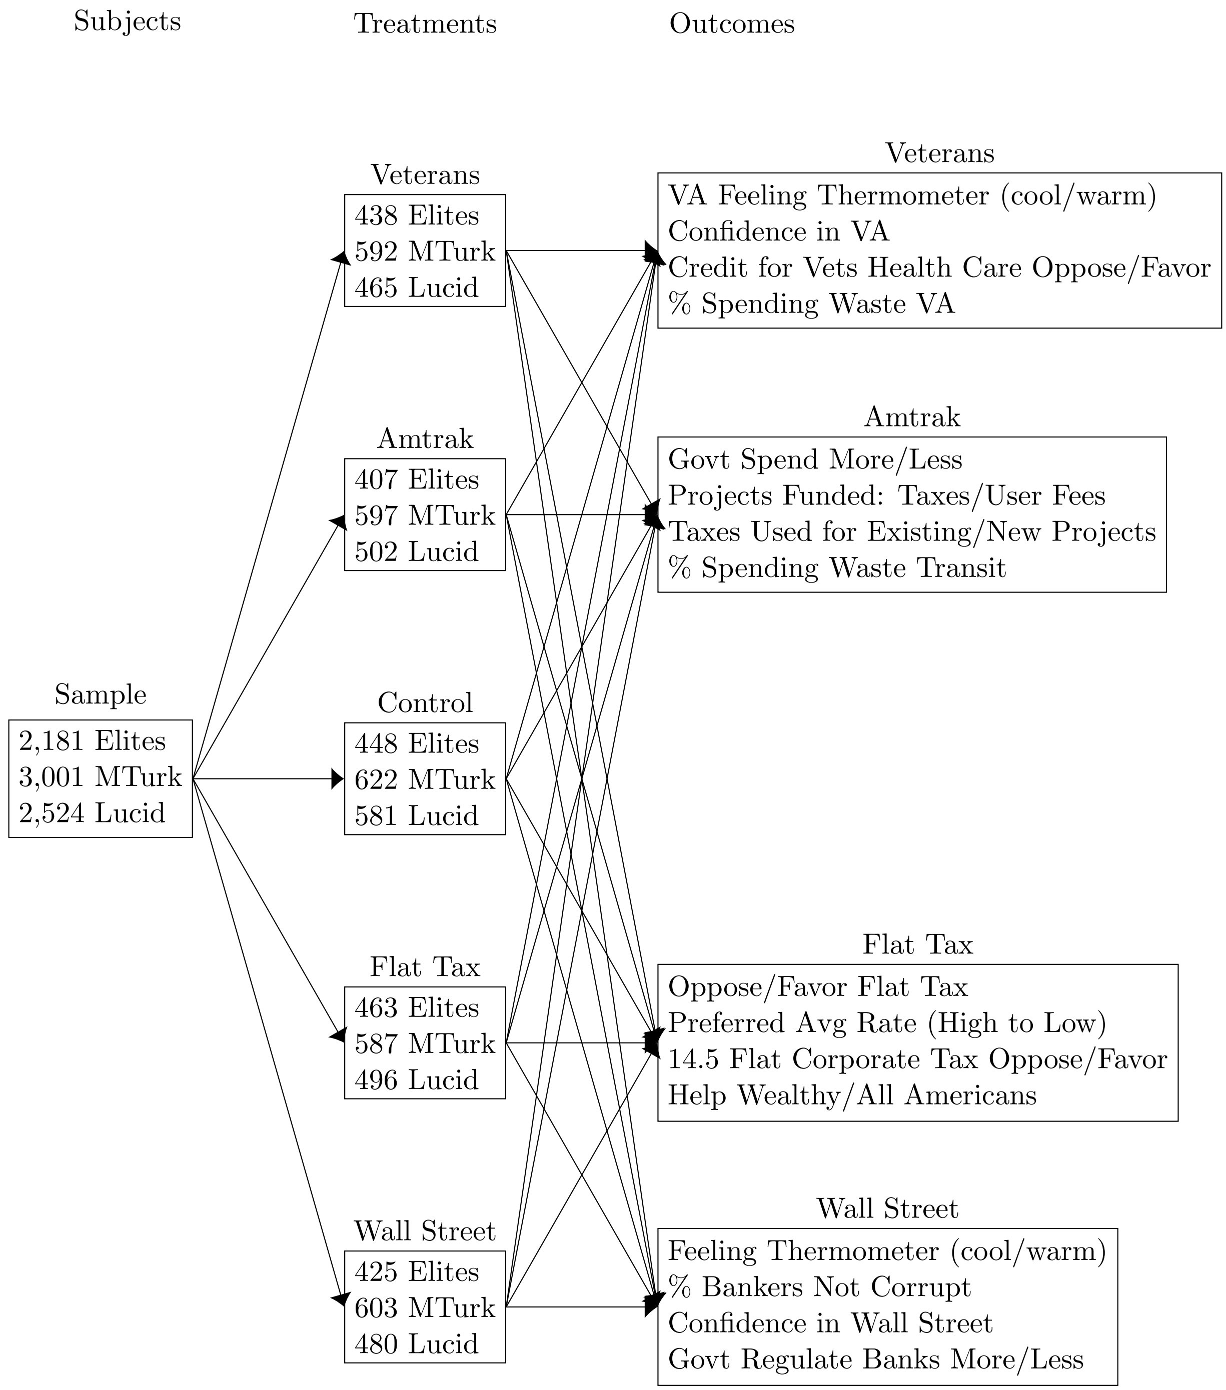

between attitudes in different domains was very small at 0.06. On Lucid, these low correlations are even lower, 0.12 and 0.00, respectively.[9]

Table 1: Levels of static constraint by study sample

	Across outcome family	Within outcome family
MTurk Sample	0.06	0.30
Elite Sample	0.22	0.39
Lucid Sample	0.00	0.12

This familiar pattern of static constraint (higher levels among elites but lower levels among the general public) is strongly contrasted by the uniformly low levels of dynamic constraint shown in Figure 4. The figure has 16 panels, each of which corresponds to the effects of an op-ed on a set of outcome measures. The rows of panels correspond to the four attitudes, each of which are measured with four questions. The columns are separate treatments. The points and bars within each panel represent treatment effect estimates and 95% confidence intervals; estimates obtained from each sample are distinguished by color and shape. The shaded panels on the diagonal indicate treatment effect estimates for each op-ed on *target* issues (the estimated effects of Z_A on Y_A, to use the terminology from above) while the nonshaded panels on the off-diagonals show the estimated effects of treatment on the *nontarget* issues (the effects of Z_A on a series of Y_B's.).

All of the reported effects are standardized with respect to the MTurk control group standard deviation, so an estimate of 0.25 indicates that the treatment group mean was 0.25 standard units higher than the control group mean. Proceeding diagonally down from the top left panel to the bottom right panel, we observe strong effects of information on target issues. The Wall Street treatment had the largest effects on its target issue of any of the treatments, averaging nearly 0.5 standard deviations in both samples. Similarly strong effects are present in the MTurk sample for the effects of the Flat Tax and Veterans treatments on their target issues, while the effects of these treatments are more muted in the elite sample. Among both elites and MTurk respondents, the effects of the Amtrak treatment on the Amtrak outcomes hover around one-third of a standard deviation. Treatment effects on the third question in the Amtrak set of outcomes (regarding how

[9] The reported correlations make no correction for measurement error in survey responses. Such corrections tend to make these correlations higher (Ansolabehere, Rodden and Snyder 2008; Freeze and Montgomery 2016), although the within-domain versus cross-domain pattern would remain.

subjects expect new taxes to be spent) were especially large, at 0.52 SDs among elites and 0.73 SDs on MTurk. Overall, out of the 48 estimated treatment effects on target issues, 39 are statistically significant at $p < 0.05$ or better. If we apply the Benjamini-Hochberg (1995) multiple comparisons correction to the associated p-values, 35 estimated effects remain significant. Follow-up interviews with these respondents shows that the this pattern of attitude change decayed somewhat but remained significant one month later, suggesting that the persuasive effects were more than mere priming.

The panels on the off-diagonals display the effects of op-eds on nontarget issues. For example, in the second row of the first column, we see how the flat tax treatment affected opinions concerning Wall Street; in the second column of the first row, we see how the Wall Street treatment affected attitudes about the flat tax. In nearly every case, the estimates are very close to zero. Out of 144 estimated effects on nontarget issues, only 9 of the effects are significant, and when we apply the Benjamini-Hochberg correction, none of these estimates is deemed significant. We examined follow-up interviews 10 days and one month later and found no evidence that attitudes had dynamically adjusted over time.[10]

In summary, Figure 4 provides clear evidence against dynamic constraint. The precision weighted average of all treatment effects on target issues is 0.42 on MTurk, 0.22 in the elite sample, and 0.245 on Lucid. By contrast, the precision weighted average of effects on nontarget issues is tiny in all three samples: 0.004 standard deviations on MTurk, 0.01 in the elite sample, and -0.02 on Lucid. This constitutes evidence that these treatments do indeed affect their target issues (Z_A affects Y_A) but do not affect nontarget issues (Z_A does not affect Y_B). This pattern of evidence is consistent with the three models described in Figure 2 only, none of which include a role for I as a mediator of exogenously-induced changes in opinions.

We can express the size of the of the average effect on nontarget issues as a proportion of the average effect on target issues. On MTurk, our best guess is that the average indirect effect is 0.8% of the average direct effect, with a very tight confidence interval[11] from -8.1% to 8.7%. On the

[10] In the appendix, we present figures analogous to Figure 4 for attitudes measured 10 and 30 days after treatments were delivered.

[11] Because the treatment effect estimates are nonindependent, we constructed a confidence interval around this ratio by bootstrapping at the respondent level.

elite sample, we estimate that indirect effects are 6.4% the size of the direct effects on average, with a wider confidence interval extending from -25.7% to 26.2%. On Lucid, the fraction is in fact negative, at -9.4%, with a confidence interval from -24% to +5%. If we pool the three studies together, our estimate is 1.4% (95% CI: -8%, 6%). In summary, this study was well-powered to detect even very small indirect effects on the order of 10 to 20 percent the size of direct effects, but we found no indication at all that these treatments induced any opinion change in nontarget domains.

Figure 4: Effects of Newspaper Op-eds on Target and Nontarget Issues

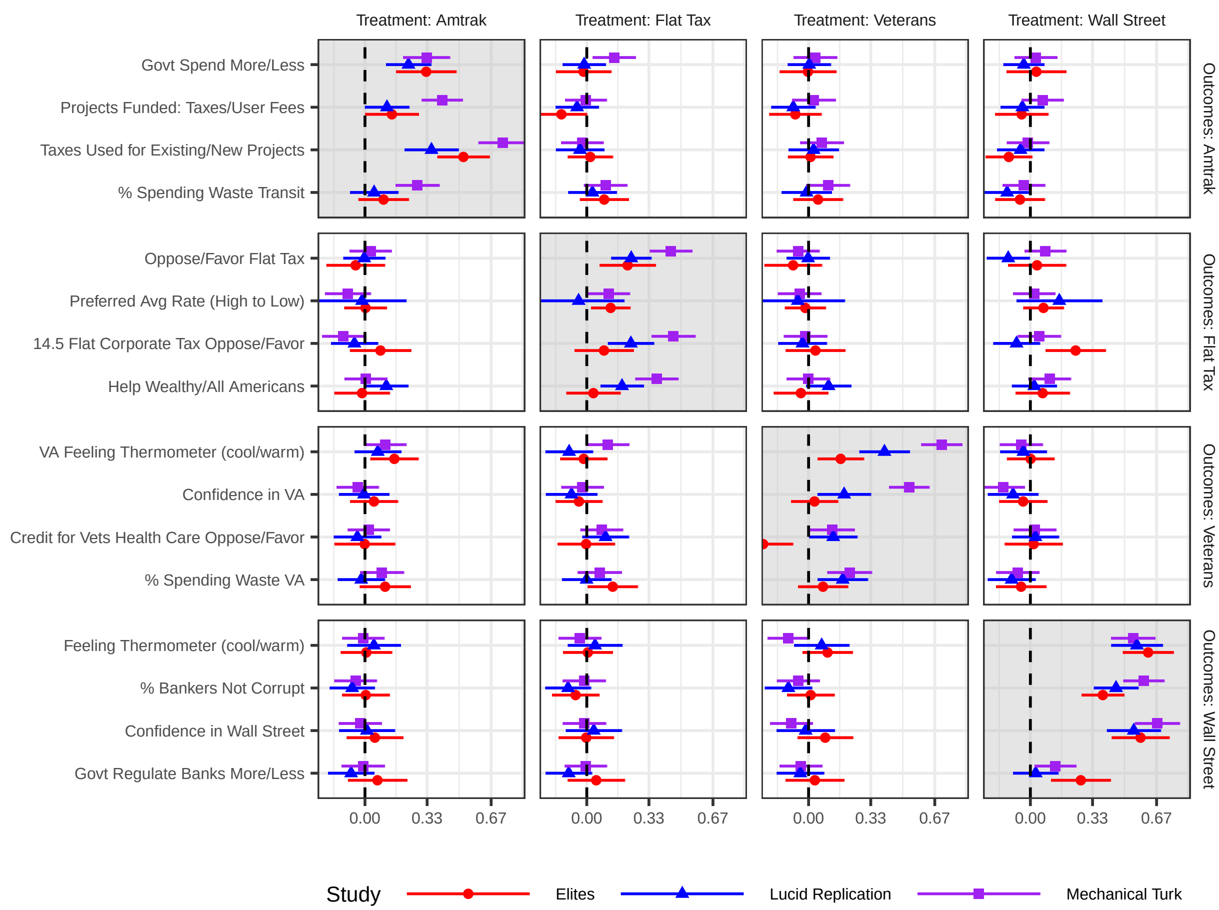

Study 2: Priming to Induce Constraint within a Policy Domain

Unlike Study 1, which considers a range of issues that are connected by the abstract principle of government involvement, Study 2 addresses the "easy" issue (Carmines and Stimson 1980) of employer discrimination against gay men and lesbians, which is neither technically complex nor primarily about means versus ends. The 1,087 subjects for this experiment were also recruited via Lucid. As shown in the Appendix, this Lucid sample contains a higher fraction of women and is somewhat better educated than the national population.

Subjects were shown one version of the following question at random, with the experimental variation in wording indicated in brackets: "Do you agree or disagree with the following statement: Employers [who have deeply held religious convictions] should have the right to refuse to hire gay men and lesbians." [1: Strongly disagree, 7: Strongly agree] We expected the religious convictions manipulation to increase our subjects' view that discrimination against gay men and lesbians is acceptable, which is what we found. We wanted to know if raising the religious objection in subjects' minds would change their responses to subsequent questions on closely related topics that all draw on the same pool of considerations. One outcome measure raises the same objection in a different domain (e.g., religious objections among officials who oversee marriage licenses), while the other two outcome measures broach the subject of same-sex marriage more generally or the acceptability of prayer in public school settings. On their face, these topics seem closely related insofar as they all revolve around the primacy of traditional religious values. Empirically, the level of static constraint among them is rather high. Responses to these nontarget outcomes are positively correlated among each other and with the target outcome of the manipulation (correlations range from 0.3 to 0.7), so it is reasonable to imagine that subjects would apply the religious conviction consideration when answering the three further questions.

> *Local Officials*: Sometimes local officials who issue marriage licenses have religious objections to same-sex marriage. Should these local officials be required to issue marriage licenses to same-sex couples, or should local officials be excused from issuing licenses that violate their religious convictions? [Binary: excused = 1]

Same-sex Marriage Amendment: Would you favor or oppose an amendment to the United States Constitution that would allow marriage only between a man and a woman, and outlaw marriages between people of the same sex? [1: Strongly oppose, 7: Strongly favor]

Prayer: Do you agree or disagree with the following statement: "Football coaches at public high schools should be allowed to lead their players in Christian prayer during games." [1: Strongly disagree, 7: Strongly agree]

Prayer Amendment: Would you favor or oppose an amendment to the United States Constitution that would require a daily prayer to be recited in public school classrooms? [1: Strongly oppose, 7: Strongly favor]

Table 2: Employment Discrimination Experimental Results

	Initial Outcome	Subsequent Outcomes			
	Employers	Local Officials	SSM Amend	Prayer	Prayer Amend
Religious Exemption	0.444*	−0.003	0.217	−0.075	0.087
	(0.133)	(0.029)	(0.143)	(0.115)	(0.131)
Constant (Control Mean)	3.016	0.381	3.679	5.115	4.164
	(0.090)	(0.021)	(0.099)	(0.081)	(0.093)
N (Lucid Original)	1,087	1,086	1,087	1,087	1,085
		Replication Study			
Religious Exemption	0.410*	−0.038*	0.049	−0.014	0.019
	(0.080)	(0.018)	(0.085)	(0.073)	(0.079)
Constant (Control Mean)	3.061	0.383	3.634	4.788	4.041
	(0.056)	(0.013)	(0.060)	(0.051)	(0.056)
N (Lucid Replication)	2,954	2,950	2,953	2,950	2,953

*p < .05. Robust standard errors are in parentheses.

Table 2 shows that the religious exception treatment increased subjects' agreement that employers should have the right to discriminate by 0.44 scale points in the original study and 0.41 points in the replication. These estimates are highly statistically significant in both samples and represent large first-stage effects. However, the four additional outcome measures did not change in a manner consistent with the initial response, as shown in columns 2 through 5. In both samples, two of the estimated effects are positive, and two are negative, but all are close to zero and statistically insignificant. When primed to consider a value – affirmation of others' religious convictions – people become substantially more accepting of job discrimination against gay men and lesbians, but this predisposition evidently has no bearing on the way they adjudicate the tension between religious convictions and processing marriage licenses for same-sex couples, nor does it change the weight they accord religious conviction on other church-state issues. What Singer (1966, p.69) called the "drive for consistency" seems to operate very weakly in the domain of political attitudes, even when respondents are asked to address closely connected issues.

Study 3: Replication of Trump and White (2018)

Experimental investigation of the specific trios of treatment, target attitude, and nontarget attitude considered in studies 1 and 2 reveal scant evidence of dynamic constraint. The extent to which this absence of dynamically constrained ideologies travels to other (Z_A, Y_A, Y_B) trios is an open question that will have to await further study. In the meantime, we can speculate on some mostly likely cases for dynamic constraint. All else equal, we expect that treatments with larger treatment effects on target attitudes will be more likely to generate dynamic constraint, for the main reason that large effects on Y_A suggest large effects on the mediators between Z_A and Y_A, including possibly those mediators (the belief system) that are also linked to Y_B. We also expect more dynamic constraint the more tightly linked Y_A and Y_B are. In the limit, when Y_A and Y_B are not distinct attitudes but are instead two separate measures of the same attitude, we would predict indeed that if Z_A moved Y_A that it would move Y_B as well (the within-domain consistency of treatment effects in Study 1 follows this pattern). In general, we expect that two distinct attitudes will exhibit more dynamic constraint as they share more causal mediation pathways from treatment to outcome in common.

In the spirit of searching for a most likely case of dynamic constraint, we turn now to a replication and reanalysis of Trump and White (2018), a survey experimental investigation of whether changing beliefs about income inequality also changes system justification attitudes. Their study maps neatly into our theoretical setup. The experimental treatment (Z_A) manipulates the vertical scale of time-series graphs of income inequality in the US. In the control condition, the figure shown to subjects has a vertical scale that extends from 0 to 100, so the changes in income inequality appear small. In the treatment condition, the axis zooms in to the range of the data such that income inequality appears to have risen dramatically in recent years. Belief in income inequality (Y_A) is the attitude targeted by this intervention and economic system justification (Y_B) is the attitude that ought to change if ideologies about economic fairness are dynamically constrained. The economic system justification scale is built from subjects agreement with statements like "The existence of widespread economic differences does not mean that they are inevitable" and "There are many reasons to think that the economic system is unfair." This experiment provides an appropriate test for the presence of dynamic constraint because of the tight theoretical link between inequality and fairness: when beliefs about income inequality go up, expressions of system justification should go down.[12]

Despite these reasonable theoretical expectations, the results of Trump and White (2018) are consistent with low levels dynamic constraint. They find that the graphical manipulation of inequality data has clear and unmistakable effects on beliefs about income inequality, but small, nonsignificant effects on a measure of economic system justification (Jost and Thompson 2000). Our replication study conducted on Lucid found precisely the same pattern.[13] The results from both studies are shown in Table 3. The average effects of the treatment on belief that income inequality increased are very large (approximately 25 percentage points in both samples). By contrast, the effects on economic system justification are small and close to zero: -0.18 standard units in the original study and -0.07 units in the replication.

By no means do we conclude from Study 3 that dynamic constraint never occurs; it plainly occurs everytime a person applies a political principle learned in one domain to a second. Nevertheless, the results of Study 3 suggest to us that the share of people with dynamically constrained attitudes is relatively small. The treatment changed many people's beliefs about income inequality (large effects of Z_A on Y_A) but these effects did not have concomitant effects on the views that respondents' expressed shortly thereafter about the fairness of the economic system (Y_B), even though the link between Y_A and Y_B is in theory quite clear.

[12] See the supplementary materials for the full text of both the treatments and outcomes.

[13] In our replication, we included a pure control condition in which subjects saw no inequality information at all. We find that the "low inequality" control used in the original study decreases belief in income inequality, thereby maximizing the contrast between control and treatment. Even with this very effective manipulation, we find no discernible movement on system justification.

Table 3: Inequality Experimental Results

	Belief in Inequality (binary) (1)	Economic System Justification (standard units) (2)
Inequality Treatment	0.261*	−0.181
	(0.030)	(0.103)
Constant (Control Mean)	0.394	−0.000
	(0.022)	(0.077)
N (Original)	1,020	336
	Replication Study	
Inequality Treatment	0.229*	−0.071
	(0.037)	(0.066)
Pure control	0.122*	−0.063
	(0.039)	(0.068)
Constant (Control Mean)	0.446	−0.254
	(0.029)	(0.050)
N (Lucid Replication)	1,041	1,026

*p < .05. Robust standard errors are in parentheses.
A random third of original study responents was asked the Economic System Justification outcome question.

Discussion

Since the publication of Converse's famous essay, the empirical and theoretical literature on ideological constraint has focused almost exclusively on static constraint, often measured by the correlation among respondents' political attitudes across issues. With the exception of Peffley and Hurwitz (1992), which focused on foreign policy views, Converse's notion of dynamic constraint has largely escaped academic scrutiny.

One reason why dynamic constraint may have been understudied to date is that the research design challenges are somewhat daunting. First, we need some exogenous shock that changes a target attitude. Second, we have to be able to credibly assume that the shock does not itself directly affect nontarget attitudes, but instead affects nontarget attitudes only through an intermediate variable, the belief system. Finally, the object of study is a causal mediation process (always difficult to study, even experimentally (Imai, Tingley and Yamamoto 2013)), but here we have the added challenge that the mediator is unobserved. We addressed the first two challenges by using a randomized design in which we were in control of both the allocation and the content of the treatments. We grappled with the third challenge by enumerating the full set of theoretical models that met our foundational assumptions, mapping those into observable empirical patterns, then using our experimental estimates to rule out some models in favor of others.

What did we find? In Study 1, the intervention induced significant opinion change on target issues in 39 of 48 opportunities, providing a quite powerful first-stage platform for the exploration of dynamic effects. But in only 9 of 144 opportunities did it induce significant changes on nontarget issues, and the average spillover effect is just 1.4% of the first-stage direct effect. Study 1 covered an assortment of substantive issues, and our results are consistent across all of them. Inducing attitude change in one domain does not bring about change in related domains. The only theoretical models that are consistent with this empirical pattern rule out the belief system as a mediator. On this basis, we conclude that respondents' attitudes were not dynamically constrained.

Study 2 examined the spillover effects of raising considerations in the context of the questions that are posed to respondents. Reminding subjects that some employers may have religious objections clearly increases subjects' support for discrimination, but this effect does not alter what people say shortly thereafter in response to closely related issues like school prayer or same-sex marriage. In other words, those randomly induced to condone discrimination against gay men and lesbians are no more likely than their control group counterparts to condone discrimination on the closely related topic of same-sex marriage. Study 2, like Study 1, offers little evidence of dynamic constraint.

A possible criticism of both studies 1 and 2 is that heterogeneity in treatment effects could be correlated with heterogeneity in dynamic constraint. It could be that subjects who "comply" with the treatment by updating their target attitude are the subjects who are least dynamically constrained – after all, they were the easiest to move. Subjects who are unresponsive to treatment – the noncompliers – may be those who are dynamically constrained, but we fail to detect them because of their steadfastness on target attitudes. Study 3 provides a partial response to this point. We generate very large first stage effects, which is to say we generate a large share of compliers, yet we still find little evidence that beliefs about inequality themselves shift economic system justification. We concede, however, that further research is needed to overcome this limitation. One approach is to repeatedly induce first-stage opinion change among the same set of respondents in an effort to nudge every subject into compliance at least once. Another would be to intervene in ways that call respondents' attention to the connections between issues, perhaps by showing an opinion leader modeling this kind of thinking (e.g., "gay rights are civil rights").

Another reason that dynamic constraint has received less scholarly attention than it deserves is that observers may have assumed that because levels of *static* constraint have been low, the scope for politically meaningful *dynamic* constraint is relatively limited. This intuition seems flawed – there is no necessary connection between levels of dynamic and static constraint. To see this, consider a set of subjects among whom Y_A and Y_B are very weakly correlated, which is to say static constraint is low. Now imagine a treatment Z_A that successfully increases target attitude Y_A for all subjects and, via dynamic adjustment, also increases nontarget attitude Y_B for all subjects by the same amount. The two post-treatment attitudes Y_A and Y_B remain just as uncorrelated as before but have moved together via a dynamically constrained process. While we do not find this pattern in our experiments, it was of course a possibility ex ante.

The lack of a necessary connection between static and dynamic constraint is underscored by a comparison of the elite and mass public samples in Study 1. These two sets of respondents displayed the same pattern that Converse documented decades ago: low static constraint in the mass public, but comparatively higher static constraint among the elites. Although we were able to generate strong treatment effects on target attitudes for both mass and elite, our treatments had essentially zero effect on nontarget attitudes for both groups. Although elites have a clearer sense of how idea-elements in an ideological belief system fit together, when one idea-element changes, elites do not spontaneously adjust these elements so as to make them consistent.

This disjuncture raises a theoretical puzzle. Why do we observe relatively high levels of static constraint among elites if they did not arrive at them dynamically? That is, how do elites construct their belief systems, if not by applying lessons from one domain to others?

The surveys for Study 2 (both the original and replication) included a small additional experiment that suggests one possibility. We gave subjects an opportunity to contradict themselves when expressing their opinion about spending preferences. Among people who did contradict themselves, we randomized whether we pointed out their inconsistency or not. We then measured their subsequent spending preferences to see if these subjects become more consistent with their "general" spending view. The general spending question asked: "Generally speaking, would you say you favor smaller government with fewer services, or larger government with more services?" The next question asked: "Thinking about the federal budget, do you want to see the President and Congress increase spending on education, decrease spending, or keep it about the same?" As is well documented (Sears and Citrin 1982; Ahler and Broockman 2018), inconsistency between preferences for smaller government and increased social spending is common. Among the MTurk sample, 30% of the respondents were inconsistent; this figure was 24% on Lucid.

We conducted our experiment among the inconsistent subset only. Half of these subjects were randomly assigned to receive a query: "We noticed that you said that generally speaking, you preferred a [smaller/larger] government with [fewer/more] services but that you also want to [increase/decrease] federal spending on education. Just to check, we'd like to ask you about education spending again." The experimental outcome is not their response to the education spending question, but subsequent spending questions, asked directly after: "Still thinking about the federal budget, do you want to see the President and Congress increase spending on each of the following, decrease spending, or keep it about the same?" [Five spending categories, Social Security, Medicare, Health Care for Low-income Families, National Defense, Foreign Aid]. We average all five responses so that -1 on the resulting scale indicates that a respondent preferred decreasing spending in all five categories and +1 indicates a preference for increases in all categories.

The results are shown in Table 4. When induced to become more consistent, people who prefer smaller government prefer less spending (-0.172 scale points on MTurk, -0.153 scale points on Lucid). We see very similar results among inconsistent subjects who prefer larger government, though owing to their comparatively small numbers, our effect estimates are less precise in this subgroup. Overall, we interpret these results to indicate a clear shift toward greater consistency among those who receive an encouragement, suggesting that respondents are capable of maintaining consistency when that becomes a salient concern.

When party elites and members of Congress express ideologically inconsistent views, they immediately encounter exactly the sorts of criticism that induce consistency. We therefore hypothesize

Table 4: Correction Experiment

	Average Federal Spending Preferences	
	Prefer Small Gov.	Prefer Large Gov.
	(1)	(2)
Correction Treatment	-0.172^{*}	0.174
	(0.047)	(0.224)
Constant (Control Mean)	0.316	0.043
	(0.030)	(0.174)
N (MTurk)	299	26
	Replication Study	
Correction Treatment	-0.153^{*}	0.077
	(0.030)	(0.074)
Constant (Control Mean)	0.232	0.064
	(0.021)	(0.049)
N (Lucid replication)	609	111

$^{*}p < .05$
Robust standard errors are in parentheses.

that it may be the strength and frequency of encouragement that distinguishes elites from masses, as the former are continually admonished by party leaders, interest groups, and commentators to maintain their fealty to a body of ideological ideas. For those on the front lines of politics, the encouragement that we administered experimentally is meted out daily and often with great vehemence.

With the advent of large-scale experimental studies of public opinion, contemporary scholarship is now able to fill in the gaps in Converse's empirical case against dynamic constraint. Our studies only consider a few of the issue domains across which dynamic constraint could manifest itself, and more work is needed to assess whether our results hold for other opinion domains and other types of interventions that could precipitate opinion change. That said, judging from the evidence at hand, it appears that support for Converse's position is perhaps even stronger than he suspected. Changes in opinion do not ramify across domains, at least not without strong encouragement, which for most people rarely occurs.

References

Abramowitz, Alan I. and Kyle L. Saunders. 2008. "Is Polarization a Myth?" *Journal of Politics* 70(2):542–55.

Achen, Christopher. 1975. "Mass Political Attitudes and the Survey Response." *American Political Science Review* 69:1218–1231.

Ahler, Douglas J. and David E. Broockman. 2018. "The Delegate Paradox: Why Polarized Politicians can Represent Citizens Best." *The Journal of Politics* 80(4):1117–1133.

Ansolabehere, Stephen, Jonathan Rodden and James M. Snyder. 2008. "The Strength of Issues: Using Multiple Measures to Gauge Preference Stability, Ideological Constraint, and Issue Voting." *American Political Science Review* 102(02):215–232.

Bafumi, Joseph and Michael C. Herron. 2010. "Leapfrog Representation and Extremism: A Study of American Voters and their Members in Congress." *American Political Science Review* 104(03):519–542.

Bauer, Paul C., Pablo Barberá, Kathrin Ackermann and Aaron Venetz. 2017. "Is the Left-Right Scale a Valid Measure of Ideology?" *Political Behavior* 39(3):553–583.

Benjamini, Yoav and Yosef Hochberg. 1995. "Controlling the False Discovery Rate: a Practical and Powerful Approach to Multiple Testing." *Journal of the Royal Statistical Society. Series B (Methodological)* pp. 289–300.

Broockman, David and Joshua Kalla. 2016. "Durably Reducing Transphobia: A Field Experiment on Door-to-Door Canvassing." *Science* 352(6282):220–224.

Carmines, Edward G. and James A. Stimson. 1980. "The Two Faces of Issue Voting." *American Political Science Review* 74(1):78–91.

Converse, Philip E. 1964. The Nature of Belief Systems in Mass Publics. In *Ideology and Discontent*, ed. David E. Apter. Ann Arbor: University of Michigan Press.

Converse, Philip E. 1970. Attitudes and Non-attitudes: Continuation of a Dialogue. In *The Quantitative Analysis of Social Problems*, ed. Edward R. Tufte. Vol. 168 p. 189.

Converse, Philip E. and Roy Pierce. 1986. *Political Representation in France.* Cambridge, MA: Harvard University Press.

Coppock, Alexander and Oliver A. McClellan. 2019. "Validating the Demographic, Political, Psychological, and Experimental Results Obtained from a New Source of Online Survey Respondents." *Research & Politics* 6(1):1–14.

Finseraas, Henning and Andreas Kotsadam. 2017. "Does Personal Contact with Ethnic Minorities Affect Anti-immigrant Sentiments? Evidence from a Field Experiment." *European Journal of Political Research* 56(3):703–722.

Freeder, Sean, Gabriel S. Lenz and Shad Turney. 2019. "The Importance of Knowing "What Goes with What": Reinterpreting the Evidence on Policy Attitude Stability." *The Journal of Politics* 81(1):274–290.

Freeze, Melanie and Jacob M. Montgomery. 2016. "Static Stability and Evolving Constraint Preference Stability and Ideological Structure in the Mass Public." *American Politics Research* 44(3):415–447.

Green, Donald P., Shang E. Ha and John G. Bullock. 2010. "Enough already about "black box" experiments: Studying mediation is more difficult than most scholars suppose." *The Annals of the American Academy of Political and Social Science* 628(1):200–208.

Hopkins, Daniel J. and Jonathan Mummolo. 2017. "Assessing the Breadth of Framing Effects." *Quarterly Journal of Political Science* 12(1):37–57.

Imai, Kosuke, Dustin Tingley and Teppei Yamamoto. 2013. "Experimental designs for identifying causal mechanisms." *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 176(1):5–51.

Jacoby, William G. 1991. "Ideological Identification and Issue Attitudes." *American Journal of Political Science* 35(1):178–205.

Jewitt, Caitlin E. and Paul Goren. 2016. "Ideological structure and consistency in the age of polarization." *American Politics Research* 44(1):81–105.

Jost, John T. and Erik P. Thompson. 2000. "Group-Based Dominance and Opposition to Equality as Independent Predictors of Self-Esteem, Ethnocentrism, and Social Policy Attitudes among African Americans and European Americans." *Journal of Experimental Social Psychology* 36(3):209 – 232.

Kinder, Donald R. and Cindy D. Kam. 2010. *Us Against Them: Ethnocentric Foundations of American Opinion.* Chicago, IL: University of Chicago Press.

Kinder, Donald R and Nathan P Kalmoe. 2017. *Neither liberal nor conservative: Ideological innocence in the American public.* University of Chicago Press.

Lane, Robert Edwards. 1962. *Political ideology: Why the American common man believes what he does.* Free Press of Glencoe.

Malka, Ariel, Yphtach Lelkes and Christopher J. Soto. 2017. "Are Cultural and Economic Conservatism Positively Correlated? A Large-Scale Cross-National Test." *British Journal of Political Science* .

Peffley, Mark and Jon Hurwitz. 1992. "International Events and Foreign Policy Beliefs: Public Response to Changing Soviet-US Relations." *American Journal of Political Science* pp. 431–461.

Peyton, Kyle. 2020. "Does Trust in Government Increase Support for Redistribution? Evidence from Randomized Survey Experiments." *American Political Science Review* 114(2):596–602.

Pollock, Philip H, Stuart A Lilie and M Elliot Vittes. 1993. "Hard issues, core values and vertical constraint: The case of nuclear power." *British Journal of Political Science* 23(1):29–50.

Sears, David O. and Jack Citrin. 1982. *Tax Revolt: Something for Nothing in California.* Cambridge, MA: Harvard University Press.

Singer, Jerome. 1966. Motivation for Consistency. In *Cognitive consistency: Motivational antecedents and behavioral consequents*, ed. Shel Feldman. New York, NY: Academic Press pp. 48–75.

Trump, Kris-Stella and Ariel White. 2018. "Does Inequality Beget Inequality? Experimental Tests of the Prediction that Inequality Increases System Justification Motivation." *Journal of Experimental Political Science* 5(3):206–216.