

Effects of Deliberative Minipublics on Public Opinion: Experimental Evidence from a Survey on Social Security Reform

Sean Ingham¹ and Ines Levin²

¹Assistant Professor, Department of Political Science, University of Georgia

²Assistant Professor, Department of Political Science, University of California, Irvine

Abstract

This article examines the potential influence of deliberative minipublics on public opinion. Using data from a large-scale survey experiment with national coverage, we investigate whether learning the conclusions of a deliberative minipublic affects observers' support for changes to the Social Security program. Survey respondents in the primary treatment conditions were exposed to the findings of deliberative citizens' panels regarding proposed changes to Social Security. Respondents in control groups did not receive any information about the deliberative minipublic. Overall, our results suggest that deliberative minipublics have some ability to affect public opinion even if members of the public acquire only minimal information about them. In particular, they are able to influence the opinions of relatively uninformed citizens. The results also suggest, however, that the effects may be limited in their extent and magnitude—at least in the scenario, modeled by our experiment, in which citizens acquire only minimal information about deliberative minipublics.

Deliberative democrats argue that public deliberation should play a central role in democratic decision-making, but what does this require at the level of institutions? One model envisions “deliberative minipublics” in which small groups of ordinary citizens deliberate together about a policy problem before making a binding decision or conveying their conclusions to voters, legislators, public officials, or the public at large (Dahl, 1989; Fung, 2003; Goodin & Dryzek, 2006; MacKenzie & Warren, 2013; Merkle, 1996; Warren & Pearse, 2008). For example, the Citizens' Initiative Review (CIR) in Oregon convenes a

All correspondence concerning this article should be addressed to Sean Ingham, Department of Political Science, University of Georgia, 104 Baldwin Hall, Athens, GA 30602. Tel: 706-542-5925. E-mail: ingham@uga.edu

deliberative panel of ordinary citizens and publishes their conclusions about ballot measures in voters' pamphlets. The British Columbia Citizens' Assembly convened a random sample of citizens who spent 9 months deliberating on alternative electoral rules before submitting their recommended rule to a binding, province-wide referendum. James Fishkin's Deliberative Polls (DPs) bring together representative samples of citizens who participate in structured deliberation about a policy problem and whose postdeliberation judgments are measured in surveys and communicated to decision-makers or the public at large.

Deliberative minipublics often lack authority to decide policy directly; their potential to affect policy normally depends on their ability to influence public opinion or other decision-makers. An important question is therefore whether their messages influence the intended audiences. Can we expect Oregon voters to revise their opinions about a ballot measure on learning the more carefully considered and informed opinions of the CIR panelists? Can media coverage of DPs affect how the public thinks about an issue? Theorists have suggested that "citizens may come to support the substantive policy findings of a minipublic because that position is the product of reasoned discussion and open participation" (Fung, 2003, p. 352). Knowing what a group of my peers concluded about a policy after weighing expert testimony and deliberating on its merits may convey information to me about what I would conclude if I were better informed. The findings of a deliberative minipublic may serve as informational shortcuts ("cues") for citizens who are relatively uninformed about the policies under debate (Gastil, 2000; Gastil, 2014; MacKenzie & Warren, 2013).

Observers may wonder, however, whether the amateur participants in deliberative minipublics possess the sophistication and experience to reach informed, reasoned opinions about complex policy issues. They are typically volunteers from a randomly selected subset of the adult population. Why should outside observers revise their policy attitudes after learning the verdicts of random collections of ordinary citizens? In theory, participants in the minipublic acquire information and sophistication as a result of listening to expert testimony, reading briefing materials, and deliberating together; empirical studies of actual deliberative minipublics support the theory (Carson et al., 2013; Fishkin, 2009; Gastil & Richards, 2013; Gastil, 2000; Goodin & Dryzek, 2006; Warren & Pearse, 2008). Nonetheless, decision-makers and members of the general public may doubt the theory and react indifferently to the policy conclusions of deliberative minipublics.

Whether deliberative minipublics have the ability to persuade decision-makers and the public at large is critical to their assessment. The question is also understudied. A recent review of empirical literature on deliberation identifies it as a "topic worth considerable additional study" (Karpowitz &

Mendelberg, 2011, p. 268). Indeed, only a few studies have addressed it (Cutler et al., 2008; Gastil et al., 2013).

This article presents the results of a survey experiment we designed to learn whether and how Americans might revise their policy attitudes on learning the findings of a deliberative minipublic. We looked specifically at attitudes over several possible changes to Social Security and tested whether exposure to a deliberative minipublic's conclusions about the proposed changes affected these attitudes. Along with this primary question of interest, the study permits a comparison of the persuasive effects of messages from deliberative minipublics with messages from political elites. Previous literature has documented the effects of "cues" from political elites (Bullock, 2011; Lupia & McCubbins, 1998; Lupia, 1994; Zaller, 1992). If deliberative minipublics can persuade voters even in the face of contradictory messages from political elites, then they have the potential to serve as popular, democratic counterweights to elite political discourse.

Social Security is an appropriate case for studying the potential influence of deliberative minipublics. It is an area where citizens might be especially receptive to policy recommendations from deliberative minipublics, as most Americans are relatively uninformed about the state of the program and proposals for reform (Barabas, 2012). Advocates and politicians on different sides of the debate often present inconsistent and misleading views on the sustainability of the Social Security program (Cook et al., 2002; Jerit & Barabas, 2006). Ordinary citizens may view a deliberative panel of their peers as a more reliable signal of what they themselves would think, were they better informed. Indeed, a previous study found that participation in a deliberative forum where Social Security was discussed considerably affected participants' knowledge and opinions on the issue (Barabas, 2004). These are just the circumstances under which one would expect deliberative minipublics to be capable of influencing public opinion. Another reason for choosing this policy area is that real-world deliberative citizens' forums and Congressional commissions have tried to influence public opinion on the same proposals for reforming Social Security that the experimental treatments describe.

The experiment's results give partial support to the prediction that members of the American public are disposed to treat the findings of a deliberative minipublic as informative signals of their own counterfactual, better informed opinions. In line with the theory, learning about the minipublic's findings had stronger effects on the survey respondents who were less well-informed about the policy issue. The directions of these effects, however, were not always consistent with the theory. Toward the end of the article, we discuss possible explanations for the mixed results.

Deliberative Minipublics in Theory and Practice

A widely used implementation of the deliberative minipublic model is James Fishkin's DP (Fishkin, 1995; Fishkin, 2009; Merkle, 1996; Sturgis, Roberts, & Allum, 2005). A random, representative sample of a target population is polled on an issue and then invited to gather for a weekend and deliberate further on it. Participants are paid; they receive briefing materials, participate in moderated discussions, and have opportunities to interrogate experts.

In a few cases, DPs have been formally authorized to make a binding decision, but in most cases, they have served an advisory function. Their influence over decisions, such as it is, lies in persuading decision-makers or affecting the perceived public legitimacy of decision-makers' options (Fishkin, 2009, ch. 5). That the results of DPs have a causal effect on public opinion is the premise on which this second mechanism rests; yet, to our knowledge, the premise has never been subjected to a rigorous experimental test. Fishkin speculates that the results of an Australian DP influenced public opinion on a constitutional referendum, but acknowledges that "the effects of media coverage [of DPs] on the broader population are an element for further study when resources permit" (Fishkin, 2009, p. 150).

Another example of a deliberative minipublic was the 2004 British Columbia Citizens' Assembly, convened by the provincial parliament and authorized to submit a proposal for electoral reform to a binding province-wide referendum. [See Warren and Pearse (2008) for a collection of scholarly essays on the BC Citizens' Assembly.] Randomly selected residents were invited to participate. From the volunteers, 160 were selected to take part. The participants met over the course of a year to hear expert testimony and deliberate about the merits of different electoral systems. They became demonstrably better informed about the topic as a result (Blaise et al., 2008). Other citizens' assemblies have since been held in Ontario, the Netherlands, and Australia.

The Oregon CIR, begun in 2010, puts together a random sample of citizens to deliberate about ballot measures. They meet for several days, hear expert testimony and arguments from campaigns for and against the measure, and then draft a "Citizens' Statement" that indicates how many panelists support the measure and lists the reasons that those in favor and those opposed give for their respective positions. The statement is included in voters' pamphlets. Using an online survey experiment, Gastil et al. (2013) found that reading a CIR Statement influenced subjects' knowledge of the issue and their voting intentions; see also Gastil and Richards (2013) and Gastil (2014).

In many cases, the only mechanism by which deliberative minipublics can affect policy is their ability to influence public opinion. One explanation for why they would enjoy this ability is a *signaling* mechanism of persuasion. If

observers believe that the participants in the deliberative minipublic have information about a policy issue that they lack, and they believe that the participants' interests are sufficiently similar to their own, then they may be prepared to treat the conclusions of a deliberative minipublic as informative signals of what they themselves would conclude about a policy if they had access to the same information and could also deliberate about it with more care. That the deliberative minipublic supports a policy is evidence—admittedly inconclusive—that they would also support the policy if they were just as informed about the policy and had the same opportunity to deliberate about its merits as the members of the deliberative minipublic.

The mechanism we are describing would be captured by a “cheap talk” model of communication (Crawford & Sobel, 1982; Farrell & Rabin, 1996; Lupia & McCubbins, 1998). The “sender” (the deliberative minipublic) sends a message (its publicly announced conclusion about a policy) to the “receiver” (the public); the receiver then takes some decision (e.g., voting in a referendum) over which both parties have preferences. The sender has access to decision-relevant information that the receiver lacks. If their interests are sufficiently well aligned, the sender can credibly convey information about these facts to the receiver, and the receiver will update his/her beliefs accordingly.¹

If the observer treats the deliberative minipublic's findings as a reliable (albeit imperfect) signal of his/her counterfactual, better informed opinion, then he/she should update his/her beliefs on receiving this signal. Of course, no rational observer would regard the minipublic's findings as *conclusive* evidence. A rational observer may therefore strongly disagree with the minipublic's conclusion even after taking it into account and revising his/her beliefs. Alternatively, the effect of the revision may be that he/she is unsure about the conclusion and no longer confident that it is false.

The signal may have no effect on the opinion the observer expresses in a survey if the survey's measures of opinion are too coarse. For example, if p_0 is an observer's prior level of confidence that a particular policy should be undertaken, and p_1 is his/her (counterfactual) level of confidence conditional on learning that the deliberative minipublic approved of the policy, then the

¹In the case of the Oregon CIR, voters learn more than just whether a majority of panelists supported the measure. They also learn their “Key Findings”—factual claims about the measure—a “Citizen Statement in Support of the Measure,” listing reasons for approving the measure and indicating how many panelists took this position, and an analogous “Citizen Statement in Opposition to the Measure.” These reasons usually take the form of factual claims about the measure, but normative claims occasionally appear as well (e.g., “every voter should have equal rights in every election,” from the CIR Statement on Measure 90). For examples of past CIR Statements, visit www.healthydemocracy.org. While the messages from the Oregon CIR are richer than simple indications of whether a majority supported or opposed the measure, a signaling mechanism is still the most plausible explanation for why seeing the factual claims might influence voters. Whether reading about the panel's key findings and the factual claims offered in support or opposition influences a voter's opinion is likely going to depend on whether the voter trusts the panelists in the relevant sense.

theory predicts that $p_1 - p_0$ is positive. But if the survey asks the respondent whether he/she agrees, disagrees, or is unsure that the policy should be undertaken, then the signal may have no effect on his/her answer, even if $p_1 - p_0$ is positive: p_0 and p_1 may both be high, so that he/she would answer “Agree” whether or not he/she receives the signal, or they may both be middling, so that he/she would answer “Unsure” whether or not he/she receives the signal, and so on.

Nonetheless, if the theory is sound, then among a large population of such observers whose prior opinions vary, the conclusions of deliberative minipublics should have average effects that even coarse measures of opinion are able to detect. Some observers’ prior opinions will be close to the thresholds at which their survey responses would change. Among a population of such observers, those who have received the signal should be less likely to disagree with the conclusions endorsed by the deliberative minipublic than their counterparts who have not received the signal.

Political theorists’ explanations for how deliberative minipublics could influence public opinion are consistent with this explanation, although they are normally presented as accounts of “trust” or “opinion cues” rather than signals and inference. MacKenzie and Warren (2013) argue that properly designed deliberative minipublics can be “objects of public trust”; its members’ competence, achieved through deliberation, and the alignment of interests between its participants and the general public are central to their explanation. Warren and Gastil (2015) offer an explanation for why deliberative minipublics could serve as trusted sources of information that would allow citizens to make better judgments, and do so more effectively than professional politicians and other existing alternatives. Cutler et al. (2008) claim that the Citizens’ Assembly satisfied conditions for “warranted trust.”² Gastil (2000) describes deliberative minipublics as sources of “cues,” which might supplement or improve on the cues that citizens receive from other sources, such as politicians or interest groups. Fishkin et al. (2015) argue that a deliberative minipublic’s endorsement of a ballot initiative could be an “informational shortcut” for voters and could influence their support for the initiative because it would tell voters that “a proposal got on the ballot partly because a random sample of the people thought it was a good idea after they really thought about it in depth” (Fishkin et al., 2015, p. 1039). These claims can all be understood as variations on the theory presented above: observers will treat the conclusions of a deliberative minipublic as informative signals of what they themselves would think, if they were better informed and had deliberated about the issue. They are prepared to treat their conclusions as informative signals because they “trust” the participants in the minipublic: they believe that their interests are

²See, in particular, their conclusions about how the typical referendum voter reasoned about the Citizens’ Assembly on p. 187.

sufficiently well aligned with their own, and they believe that the participants, having deliberated about the issue, are better informed than they are.

The survey experiment was designed to test whether members of the American public would respond to the findings of a deliberative minipublic in the way this theory hypothesizes. In the next section, we describe the specific hypotheses that the theory supports in the context of the survey experiment.

Hypotheses

The theory outlined above predicts that observers should treat the policy conclusions of a deliberative minipublic as informative signals of what they would themselves think about the issue, if they had the opportunity to inform themselves and deliberate with others about it. The survey experiment allows us to test such predictions in the context of Americans' opinions about Social Security.

On the basis of the theory, one would predict that respondents who are told that a deliberative citizens' panel supports a given change to the Social Security program would be less likely to disagree that the proposed reform should be undertaken than respondents in a control condition who receive no such message (Hypothesis 1a). Relatedly, one would predict that they would be more likely to choose the proposed change over alternative proposed changes, when asked to report their top preference (Hypothesis 1b).

The next pair of hypotheses concerns how respondents would treat the findings of a deliberative minipublic when the findings conflict with messages from alternative sources. According to the theory, the value of deliberative minipublics derives in part from their recruitment of representative samples of ordinary citizens. They are supposed to be in this respect superior to other deliberative bodies such as expert panels or elected assemblies. The members of the latter deliberative bodies may be more knowledgeable about policy issues than the typical citizen who participates in a deliberative minipublic. But they may also be disproportionately white, male, and wealthy, and elected politicians and policy experts may have ties to special interests or face other incentives that bias their judgments. For this reason, the policy conclusions of a deliberative minipublic are thought to be more reliable indications of "enlightened public opinion" than the policy conclusions of expert panels or elected leaders (Fishkin, 2009; Gastil, 2000). Observers may therefore decide that they have reason to trust their policy recommendations over the conflicting policy recommendations of deliberative bodies composed of elected leaders and other political elites.³

³The connection between these hypotheses and the theory outlined above is admittedly loose: an observer may treat the message from the deliberative minipublic as an informative, reliable signal of his/her own counterfactual, better informed opinion, but believe that the signal from the panel of elites is more reliable still. For he/she may believe that the latter are more knowledgeable despite being less likely to have congruent preferences.

The experiment allows us to test whether this postulate holds true of the American public. Some respondents were randomly assigned to a treatment condition in which they read about the policy conclusion of a deliberative citizens' panel as well as the conflicting policy conclusion of a "bipartisan panel of Congressional leaders and business and labor representatives." On the basis of the postulate, one would predict that, by comparison with respondents in the control group, respondents in this treatment group would be less (more) likely to disagree that a proposed reform should be undertaken if they read that the deliberative citizens' panel supported (opposed) it even when presented with the conflicting policy conclusion of the bipartisan panel of Congressional leaders and business and labor representatives (Hypothesis 2a). One would also predict that they would be more (less) likely to prefer the policy change that the deliberative citizens' panel supported (opposed) relative to alternative policy changes (Hypothesis 2b).

The theory suggests a final hypothesis about how respondents' prior levels of familiarity with and knowledge about the policy issue are likely to condition the effects of learning about the deliberative minipublic's findings. Even if an observer regards a deliberative minipublic's conclusions as reliable signals of what he/she would conclude, were he/she better informed about a policy issue, the observer may revise his/her beliefs only slightly on learning its conclusion if the observer has a strong prior belief about the issue. Moreover, the observer's updated belief will likely translate into the same survey response: someone who is confident that a reform is a bad idea would likely still answer the survey question the same even after conflicting evidence causes him/her to revise downward his/her degree of confidence. Learning the policy recommendation of a deliberative minipublic should therefore have stronger effects among those observers who have, comparatively speaking, a weaker prior opinion about the policy issue.

The survey experiment allows us to test this prediction using indirect measures of whether a respondent has strong prior opinions about a policy issue. The survey includes self-reported measures of familiarity with the proposed changes to Social Security as well as a test of the respondent's knowledge about the payroll tax. We assume that respondents who claim to be unfamiliar with the proposed changes and who show themselves to be poorly informed about the payroll tax are less likely to have strong prior opinions about the proposed changes than their better informed and more familiar counterparts.⁴ Under this assumption, the theory predicts that the effects of the signal from the deliberative minipublic should be strongest among those

⁴Results from the survey are consistent with this assumption: respondents who report being familiar were more knowledgeable about the Social Security program, less likely to report being unsure about the merits of the proposed reforms, and less likely to report indifference when asked about their preferences over different reforms. See Table 1 A7 and A8 in the Supplementary Appendix.

who are relatively less familiar with and less informed about the proposed changes. Specifically, these uninformed respondents should be less likely than their more informed counterparts to disagree with a proposed change when they learn that the deliberative minipublic recommends it (Hypothesis 3).

Experimental Design

Our survey experiment was included in a national-representative Internet survey of individuals aged ≥ 18 years, conducted with funding from Time-Sharing Experiments for the Social Sciences. The survey was carried out by the GfK Group (formerly Knowledge Networks) between January 31, 2014 and February 13, 2014. Members of GfK's online panel were recruited using address-based probability sampling techniques applied on a sample frame—a list of addresses from the U.S. Postal Service's Delivery Sequence File—that covers 97% of U.S. households. Panel members were randomly selected to participate in the survey and received an invitation via email. Among the 9,326 members of GfK's KnowledgePanel who were invited to participate in the survey, 61.7% agreed to do so, yielding a final sample size of 5,758 respondents. In our study, we use poststratification weights provided by the polling firm to adjust for survey nonresponse and for selection deviations produced by the study's sample design (such as minority oversampling and noncoverage of unlisted households).

All respondents were given descriptions of the possible changes to Social Security as part of an initial pretreatment question. The three policies described to them were a proposal to raise the normal retirement age, a proposal to raise workers' payroll contributions to Social Security, and a proposal to raise the cap on payroll taxes. Respondents were then randomly assigned to one of eight treatment conditions, which differed along two dimensions: which policy they were asked about and what, if anything, was reported to them about the conclusions of a deliberative citizens' panel and a Congressional commission concerning the policy. We refer to these reports as "cues" and their sources as "citizens," in the case of the deliberative citizens' panel, and "elites," in the case of the Congressional commission. Table 1 summarizes the 4×2 factorial design. The eight treatment groups were broadly similar in terms of socio-demographic attributes and basic political attitudes; see Table A1 in the Supplementary Appendix.

Respondents in control conditions A0 and B0 received no information about the opinion of the panel of ordinary citizens, and were simply asked whether they "agree or disagree" that we should "raise the retirement age to 68" or "raise the cap on payroll taxes," respectively. Respondents in treatment condition A1 read the following paragraph before being asked whether they "agree or disagree that we should raise the retirement age to 68."

Table 1
Experimental design

Treatment condition	Raise retirement age	Raise cap on payroll taxes
No cues (control condition)	A0 (683, 11.9%)	B0 (773, 13.4%)
Citizens support	A1 (709, 12.3%)	B1 (715, 12.4%)
Citizens support, elites oppose	A2 (675, 11.7%)	B2 (703, 12.2%)
Elites support, citizens oppose	A3 (770, 13.4%)	B3 (731, 12.7%)

Note. The total sample size was 5,758. The survey experiment had a 4×2 factorial design. Respondents were randomly assigned to four different types of cues (see table rows) that made reference to one of two different Social Security reforms (see table columns). Each treatment condition represents a combination between a type of cue and a referenced reform. The table provides information about the label used to designate each treatment group as well as the number and percentage of respondents assigned to each treatment condition (between parentheses). Frequencies and percentages were computed using sampling weights.

A nonprofit organization recently recruited a random collection of ordinary citizens to participate in a deliberative forum on how to fund the Social Security program in the future. After hearing expert testimony, weighing evidence, and deliberating together, a majority supported raising the retirement age to 68.

Subjects in Condition A2 were shown, in addition to this description of the citizens' panel, the following:

But when a bipartisan panel of Congressional leaders and business and labor representatives recently heard expert testimony, weighed evidence and deliberated about the same issue, a majority opposed raising the retirement age to 68.

Condition A3 resembled this one but with the stances of the citizens' panel and Congressional panel reversed. Conditions B1–B3 were parallel, but with “raising the cap on payroll taxes” substituted for “raising the retirement age to 68.”

After answering the treatment question, subjects were asked which of the three changes to Social Security they would choose, if they could choose to implement just one, or if they would instead choose to leave the Social Security program as it is. The respondent again had the option of reporting indifference or being unsure.

Respondents in the treatment conditions received only a brief description of the deliberative forum. What is the justification for such a minimal experimental intervention? How should we assess its external validity? The treatment conditions should model what voters learn about deliberative minipublics and their conclusions when deliberative minipublics are incorporated into the political process. But different voters learn more or less about deliberative minipublics. According to evidence from a survey analyzed in Cutler et al. (2008), 40% of British Columbian voters claimed to know nothing about the Citizens' Assembly at the time of the referendum; in a survey of Oregon

voters, a week before the 2014 election, 46% reported being unaware of the CIR, while 35% and 20% reported being “somewhat aware” and “very aware,” respectively (Gastil et al., 2015, p. 35). Inevitably, then, the treatment conditions will contain *more* information about the deliberative minipublic than what many real-world voters would acquire, and *less* information than what many others would acquire. There is no uniquely best design, from the standpoint of external validity, given the heterogeneity of real-world voters and how much they know about deliberative minipublics. In choosing how to design the treatment, researchers face a trade-off. On the one hand, the more minimal the information contained in the treatment, the stronger the inference one can draw from *positive* results: if even minimal information about a deliberative minipublic influences respondents’ opinions, then, in the real-world, where many voters acquire even more information about the deliberative minipublic, one could more confidently infer effects on public opinion. On the other, the more minimal the information, the weaker the inference one can draw from *negative* results: even if minimal information fails to influence respondents’ opinions, real-world deliberative minipublics might influence public opinion if significant numbers of voters acquire more than minimal amounts of information. With this trade-off in mind, we designed the treatments to model the situation of voters who learn something, but not much, about the deliberative minipublic. We revisit this aspect of the design in the discussion below.

Results

Within each treatment condition, respondents were asked about one of two proposed changes to Social Security: raising the retirement age or raising the cap on payroll taxes. The first column in Table 2 presents the proportion of respondents assigned to the control condition who agreed, were unsure, and disagreed with the reform referenced in the treatment condition’s question. According to these numbers, >60% of control respondents disagreed with the proposal of raising the retirement age. Conversely, they were also highly likely to support raising the cap on payroll taxes: >50% agreed with the proposed change.⁵ The remaining columns in Table 2 present differences in agreement relative to the control group among respondents assigned to the three treatment conditions. Numbers in parentheses correspond to bootstrapped standard errors for the treatment effects.

⁵The baseline levels of support for raising the retirement age and for raising the tax on payroll taxes shown in Table 2 are consistent with those found in previous public opinion surveys, such as the July 2013 Social Security Survey conducted by Program for Public Consultation at the University of Maryland (Voice of the People 2013; see Table A2 in the Supplementary Appendix).

Table 2
Support for Referenced Reform

Agreement with reform	Baseline %	Treatment effect		
	No cues	Citizens support, no conflict	Citizens support, elites oppose	Citizens oppose, elites support
(A) Referenced reform: raise retirement age				
Disagree	61.1	-7.2 (2.7)	-4.5 (2.7)	-2.3 (2.5)
Unsure	12.3	3.0 (1.9)	6.9 (2.0)	3.9 (1.8)
Agree	26.7	4.2 (2.4)	-2.4 (2.3)	-1.6 (2.2)
(B) Referenced reform: raise cap on payroll taxes				
Disagree	20.3	1.4 (2.1)	4.8 (2.2)	8.2 (2.3)
Unsure	24.8	1.4 (2.3)	2.8 (2.3)	0.8 (2.3)
Agree	54.8	-2.8 (2.5)	-7.7 (2.6)	-9.0 (2.6)

Note. The table gives baseline percentages under the no-cues treatment (column 1) and point estimates of treatment effects (columns 2–4). Bootstrapped standard errors for the treatment effects are provided between parentheses. All analyses were adjusted using sampling weights.

When the reform in question was raising the retirement age, respondents who learned that the deliberative panel of ordinary citizens supported the reform were indeed less likely to disagree with the proposed reform than their counterparts in the control condition, as predicted. Specifically, their level of disagreement with the reform decreased by 7.2 percentage points. In the presence of conflicting messages from the deliberative citizens' panel and the panel of political elites, effects were smaller in magnitude, but still consistent with our hypotheses. When the citizens' panel supported and elites' panel opposed the proposed reform, respondents were significantly less likely to disagree with the proposed change, although here the effect was smaller—4.5 percentage points—than in the absence of conflicting messages. When the citizens' panel opposed and elites' panel supported the reform, respondents' level of disagreement did not change significantly relative to that observed in the control group. These results are broadly consistent with hypotheses H1a and H2a.

These hypotheses did not receive clear confirmation, however, when the reform in question was raising the cap on payroll taxes. In the absence of conflicting cues, being told that the citizens' panel supported raising the cap on payroll taxes caused no significant change in respondents' disagreement

with the reform. When given conflicting cues from the citizens' panel and the panel of political elites, respondents were more likely than their counterparts in the control condition to disagree with the proposed change—regardless of the position taken by the panel of ordinary citizens. Specifically, they were 8.2 percentage points more likely to disagree with raising the payroll cap when the citizens' panel was described as opposing it and the panel of elites was described as supporting it. This effect is consistent with the Hypothesis 2a. But, surprisingly, the other cue, in which the positions of the citizens' panel and panel of elites are reversed, *also* had this effect, albeit to a lesser extent: respondents were 4.8 percentage points more likely to disagree with raising the payroll cap when the citizens' panel was described as *supporting* it and the elites' panel was described as opposing it.

A puzzling pattern that was evidenced in all treatments is that the share of respondents that reported being *unsure* about their support for the proposed reform was consistently higher among groups exposed to the conclusions of the deliberative citizens' panel—although differences relative to the control group were often indistinguishable from zero. To further assess the statistical significance of this unexpected finding, we estimated a set of multinomial logit regressions, setting level of agreement (agreeing, disagreeing, and being unsure) as the dependent variable. The results of this analysis are presented in Table 3. These estimates indicate that learning about the position of the citizens' panel decreases the odds of agreeing and those of disagreeing with each proposal, relative to being unsure. To facilitate the visualization of results, Figure 1 provides a graphical depiction of the change in probability of each level of agreement caused by the introduction of deliberative cues, calculated on the basis of regression estimates. Results depicted in Figure 1 are entirely consistent with those reported in Table 2. These findings suggest that, while the increase in uncertainty is present for all treatments, it is only statistically significant for cases where respondents were exposed to conflicting cues, and only so when the referenced reform was raising the retirement age.

To gain a more complete understanding of the influence of deliberative cues, we used the regression coefficients reported in Table 3 to simulate marginal effects of all variables included in the model, and then compared the effect of changes in treatment assignment with those of marginal changes in a number of individual attributes that might contribute to explaining support for reforms to the Social Security program, such as age, income, and party identification (see Table A5 in the Supplementary Appendix). The results of this analysis indicate, among other things, that the decrease in disagreement caused by being told that the citizens' panel supported raising the retirement age is larger in magnitude than the decrease in disagreement associated with a change in party identification from Independent to Republican. When the citizens' panel was described as being at odds with the panel of

Table 3
Regression Analysis of Support for Referenced Reform

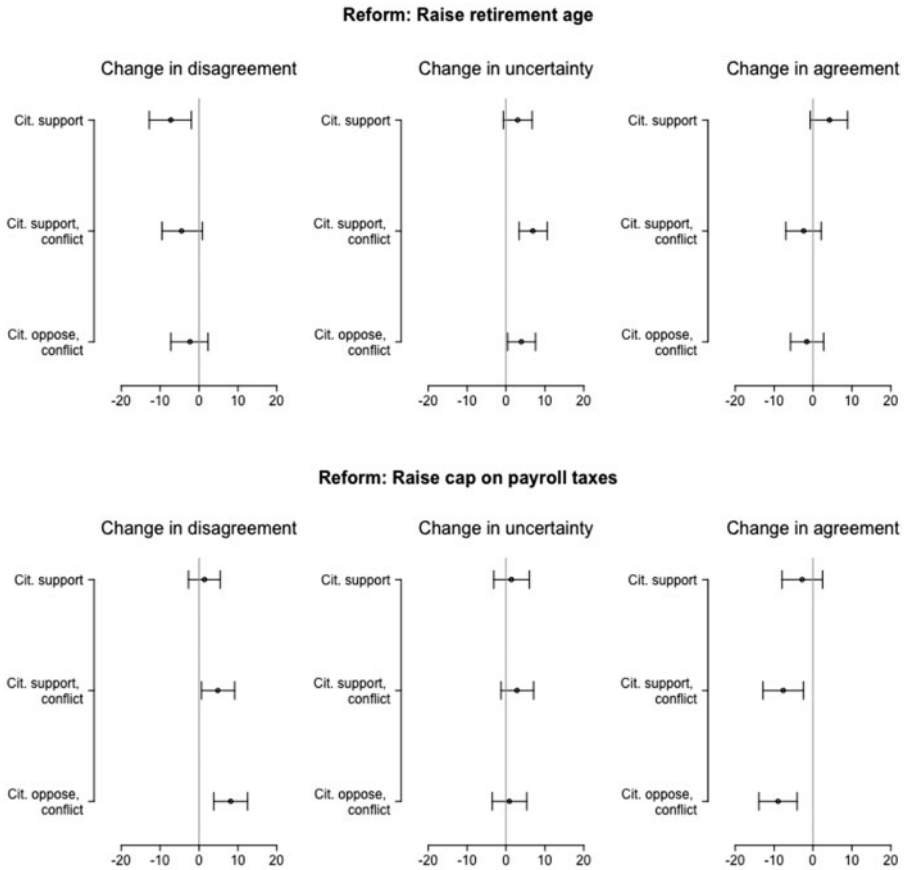
Treatment condition	Model A.1		Model A.2	
	Agree/ not sure	Disagree/ not sure	Agree/ not sure	Disagree/ not sure
(A) Referenced reform: raising retirement age				
Citizens support, no conflict	-0.111 (0.182)	-0.292 (0.168)	-0.075 (0.187)	-0.269 (0.169)
Citizens support, elites oppose	-0.514 (0.177)	-0.524 (0.16)	-0.551 (0.183)	-0.53 (0.162)
Citizens oppose, elites support	-0.438 (0.18)	-0.313 (0.162)	-0.491 (0.186)	-0.32 (0.164)
Includes control variables	No		Yes	
<i>N</i>			2,853	
	Model B.1		Model B.2	
	Agree/ not sure	Disagree/ not sure	Agree/ not sure	Disagree/ not sure
(B) Referenced reform: raising cap on payroll taxes				
Citizens support, no conflict	-0.25 (0.13)	-0.332 (0.16)	-0.252 (0.137)	-0.379 (0.164)
Citizens support, elites oppose	-0.411 (0.131)	-0.141 (0.155)	-0.381 (0.138)	-0.136 (0.159)
Citizens oppose, elites support	-0.324 (0.132)	0.019 (0.155)	-0.264 (0.139)	0.041 (0.159)
Includes control variables	No		Yes	
<i>N</i>			2,825	

Note. The tables provide point estimates and standard errors (between parentheses) of multinomial logit coefficients, with the baseline alternative set to “not sure.” The model on the right includes the following controls: income, age, education, gender, indicators of Democrat and Republican party identification, and ideology.

elites over raising the cap on payroll taxes the change in agreement was comparable in magnitude to that associated with a change in party identification from Independent to Democrat. The relationship between support for both reforms and socio-demographic attributes such as age and income was generally weak compared with the relationship between the same outcome variables and deliberative cues.

Thus, when the dependent variable was support for proposed changes to Social Security and the referenced reform was “raising the retirement age,” we found support for our hypotheses; but when the referenced reform was “raising the cap on payroll taxes,” the evidence was mixed. Next, we tested whether learning the position of the citizens’ panel affected which potential change the respondent identified as his/her top choice. Table 4 presents the proportion of respondents assigned to the control condition who designated each policy as their top choice when the reform referenced in the treatment condition was raising the retirement age (top panel) and when the referenced

Figure 1
Effect of treatment on support for referenced reform.



reform was raising the cap on payroll taxes (lower panel). Regardless of the referenced reform, control respondents identified raising the cap on payroll taxes as their preferred choice more frequently than any other possible reform: >42% selected this option. The second most popular option—preferred by >13% of control respondents—was raising the payroll tax rate, and the least popular—preferred by about 12%—was raising the retirement age. A large proportion of respondents (>27%) failed to identify any top choice. As before, the remaining columns in Table 4 present differences in proportions relative to the control group, and numbers in parentheses correspond to bootstrapped standard errors for the treatment effects.

In contrast to the earlier findings, where the dependent variable was respondents' level of agreement that the proposed reform should be implemented, we found that exposure to the views of the citizens' panel had no

Table 4
Preferences Over Reforms to the Social Security Program

	Baseline %	Treatment effect		
	No cues	Citizens support, no conflict	Citizens support, elites oppose	Citizens oppose, elites support
(A) Referenced reform: raise retirement age				
Raise cap on payroll taxes	44.1	-1.0 (2.8)	1.2 (2.8)	-4.4 (2.8)
Raise payroll tax rate	16.4	-1.3 (2.1)	-3.8 (2.0)	-1.7 (2.0)
Raise retirement age	12.3	0.2 (1.8)	-0.8 (1.8)	1.0 (1.8)
No top preference	27.2	2.1 (2.5)	3.4 (2.5)	5.0 (2.6)
(B) Referenced reform: raise cap on payroll taxes				
Raise cap on payroll taxes	42.1	3.2 (2.6)	0.9 (2.6)	-0.4 (2.6)
Raise payroll tax rate	13.9	0.2 (1.9)	-1.3 (1.8)	-0.3 (1.8)
Raise retirement age	12.2	-1.4 (1.7)	-0.3 (1.8)	-1.2 (1.7)
No top preference	31.9	-2.0 (2.4)	0.7 (2.5)	1.9 (2.5)

Note. The table gives baseline percentages under the no-cues treatment (column 1) and point estimates of treatment effects (columns 2-4). Bootstrapped standard errors for the treatment effects are provided between parentheses. All analyses were adjusted using sampling weights.

significant effect on which proposed change to Social Security the respondent selected as his/her most preferred option. These findings do not lend support to hypotheses H1b and H2b. As an additional test of these hypotheses, we estimated a new set of multinomial logit regressions, setting the top preference (raising the cap, raising the payroll tax rate, raising the retirement age, and absence of a top preference) as the dependent variable. The results of this analysis are presented in Table 5. These estimates confirm that the provision of deliberative cues did not significantly influence respondents' preferences over the reforms. Results depicted in Figure 2, which presents point estimates and confidence intervals estimated on the basis of regression estimates, indicate that treatment effects are indistinguishable from zero for all treatment conditions.

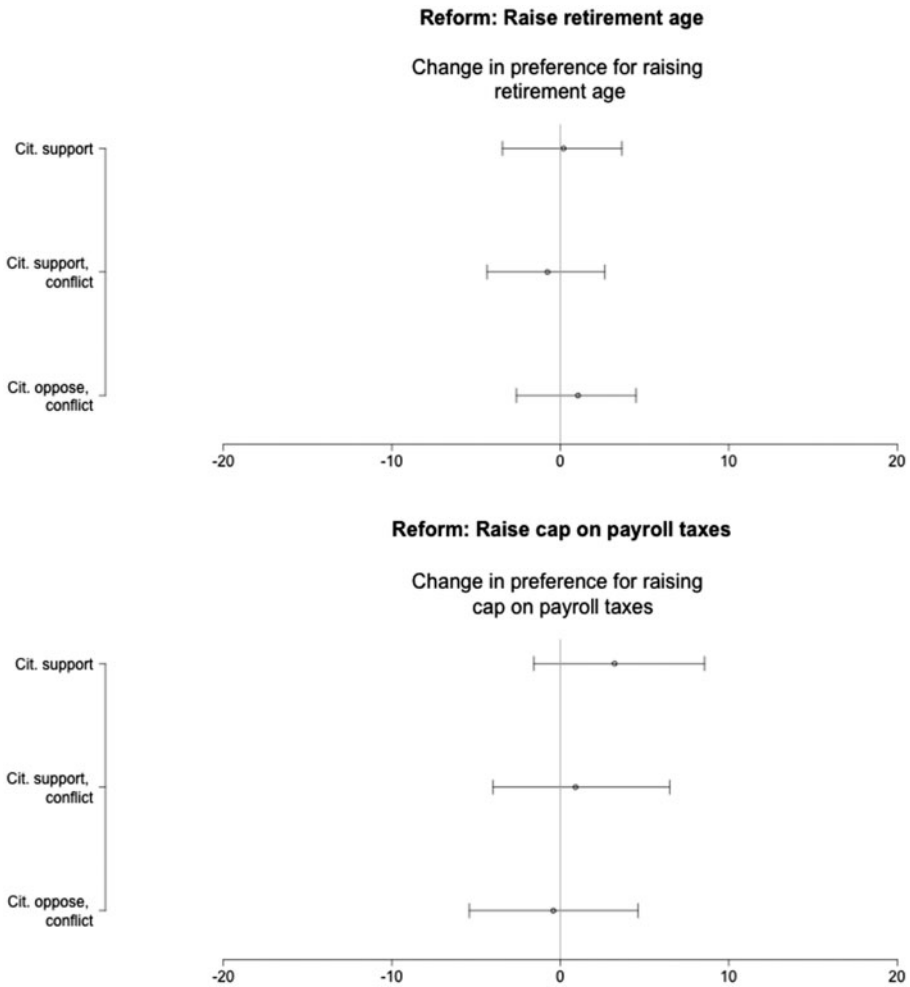
In sum, the main results of the survey experiment lent only partial support to our hypotheses. When the reform referenced in the treatment condition was

Table 5
Regression Analysis of Top Preferences

Independent variable	Model 1		Model 2		
	Raise cap/ no top pref.	Raise tax rate/ no top pref.	Raise age/ no top pref.	Raise tax rate/ no top pref.	Raise age/ no top pref.
Citizens support, no conflict	0.069 (0.134)	0.053 (0.175)	0.121 (0.188)	0.097 (0.178)	0.164 (0.193)
Citizens support, elites oppose	0.03 (0.131)	-0.139 (0.176)	-0.093 (0.191)	-0.108 (0.18)	-0.106 (0.196)
Citizens oppose, elites support	-0.086 (0.131)	-0.034 (0.171)	0.086 (0.183)	-0.052 (0.175)	0.033 (0.188)
Raise cap condition	-0.02 (0.13)	0.035 (0.169)	-0.011 (0.185)	0.054 (0.172)	-0.01 (0.19)
Raise cap condition*	-0.012 (0.186)	-0.174 (0.244)	-0.203 (0.265)	-0.195 (0.249)	-0.252 (0.273)
Citizens support, no conflict	-0.157 (0.183)	-0.301 (0.25)	0.024 (0.264)	-0.293 (0.254)	0.084 (0.271)
Raise cap condition*	-0.016 (0.184)	-0.113 (0.24)	-0.307 (0.264)	-0.05 (0.244)	-0.273 (0.271)
Citizens oppose, elites support					
Includes control variables	No	No	No	Yes	Yes
N					5,394

Note. The tables provide point estimates and standard errors (between parentheses) of multinomial logit coefficients, with the baseline alternative set to “no top preference.” The model on the right includes the following control variables: income, age, education, gender, indicators of Democrat and Republican party identification, and ideology.
*Represent multiplication signs used to indicate interaction terms in the regression mode

Figure 2
Effect of treatment on preferences over Social Security reforms



an increase in the retirement age, learning the position of the deliberative citizens' panel had the predicted effect on respondents' support for the reform, but when the referenced reform was an increase in the cap on payroll taxes, the same predicted effect failed to appear. Additionally, when the dependent variable was the respondent's choice of a most preferred reform, rather than support for the referenced reform, there was no evidence of the pattern of influence predicted by hypotheses H2a and H2b.

To test Hypothesis 3, we studied whether treatment effects varied as a function of individuals' knowledge of and self-reported familiarity with

proposals for reforming Social Security. We implemented this test using responses to two pretreatment questions. The first of these questions, measuring respondents' level of information about the Social Security tax rate, asked

Social Security, the federal program providing benefits to retired workers, is funded by a payroll tax. What is the Social Security tax rate that workers pay? In other words, what is the share of the worker's paycheck that goes toward Social Security?

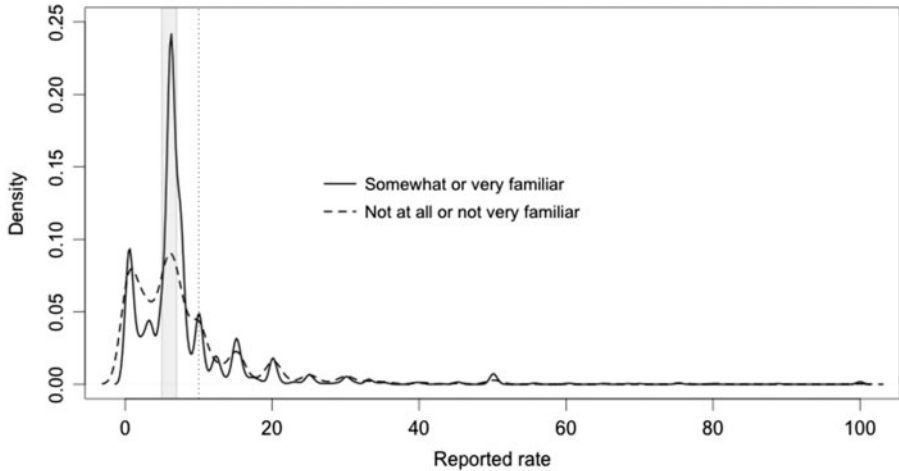
The correct answer to this question was 6.2%. Respondents were asked to answer this question by writing down a number ranging between 0.0 and 100.0 inside a number box with the following description "[NUMBER BOX]% of worker's paycheck going toward Social Security," with the ability to include decimal points. Only 8.5% of respondents got this question right. A larger percentage of respondents (30.2%), however, reported a number between 5 and 7, which is close to the actual payroll tax rate; an even larger proportion (55.2%) reported a number lower than 10.

The second question, posed immediately after the description of the three proposals for reforming Social Security, asked "How familiar are you with proposals for reforming Social Security along these lines?" The options for response were "Very familiar," "Somewhat familiar," "Not very familiar," and "Not at all familiar." Only 5.6% and 30.0% of respondents reported being very and somewhat familiar with the proposals, respectively. A majority reported being not very or not at all familiar (30.8% and 33.6%, respectively).

To examine whether self-reported familiarity correlates with being informed about elements of the Social Security program, we compared the accuracy of responses with the payroll tax-rate question among respondents describing themselves as more and less familiar with the proposals. Figure 3 depicts the distributions of the reported payroll tax rate for those describing themselves as at least somewhat familiar (smooth curve) and for those describing themselves as less than somewhat familiar (dashed curve). The gray area indicates tax rates $>5\%$ and $<7\%$ (i.e., close to the actual 6.2% rate) and the dotted vertical line indicates a 10% rate (3.8 percentage points above the actual rate). This figure suggests that the accuracy of responses to the tax-rate question varies considerably as a function of self-reported familiarity. Respondents who describe themselves as being familiar with the proposals for reform are considerably more likely to report relatively accurate rates (44% report a number between 5 and 7, and 66% a number <10 , compared with only 23% and 50%, respectively, among unfamiliar respondents).

We used responses to the payroll tax rate and familiarity questions to create a binary *information* indicator, taking the value of 1 if the person (i) reported an approximately accurate tax rate (specifically, a number between 5 and 7) and (ii) indicated being very or somewhat familiar with proposals for reforming Social Security, and the value of 0 otherwise. Based on this

Figure 3

Distribution of reported payroll tax rate

criterion, 15.5% of respondents, a total of 892 individuals, were classified as informed.

We then estimated a regression model of agreement with specific reforms (raising the retirement age and raising the cap on payroll taxes), including an interaction between treatment assignment and the binary information indicator. Subsequently, we conducted a series of simulations to test whether treatment effects vary within the two groups defined by the information indicator. Table 6 presents the results of this analysis.⁶ The first column in Table 6 presents the likelihood that a hypothetical control respondent agrees with, is unsure about, and disagrees with the referenced reform for each level of information and for each treatment condition. Differences in baseline percentages suggest that uninformed respondents are more likely to be unsure about their level of agreement, compared with more informed respondents, especially when asked about their support for raising the cap on payroll taxes. Columns 2–4 provide within-group treatment effects, measured as differences in levels of agreement, uncertainty, and disagreement, relative to the control group. Numbers given between parentheses correspond to standard errors associated with the treatment effects.

⁶Table A11 in the Supplementary Appendix presents the results of a similar analysis conducted using a more lenient information indicator, where respondents are coded as 1 if the person reported a tax rate <10% and being very or somewhat familiar with the proposals for reform.

Table 6

Model-Based Estimates of Support for Referenced Reform by Level of Information

Agreement with reform by level of information	Baseline %	Treatment effect		
	No cues	Citizens support, no conflict	Citizens support, elites oppose	Citizens oppose, elites support
(A) Referenced reform: raise retirement age				
<i>Proportion who disagree</i>				
Uninformed	65.7	-5.8 (2.8)	-5.6 (2.9)	-1.3 (2.7)
Informed	57.2	-0.9 (6.6)	3.0 (6.2)	6.1 (6.0)
<i>Proportion who are not sure</i>				
Uninformed	11.3	3.7 (2.0)	7.9 (2.1)	5.2 (2.0)
Informed	9.6	-6.7 (3.7)	-2.5 (3.8)	-2.8 (3.8)
<i>Proportion who agree</i>				
Uninformed	23.0	2.0 (2.5)	-2.3 (2.4)	-3.9 (2.3)
Informed	33.2	7.5 (6.2)	-0.5 (5.7)	-3.3 (5.5)
(B) Referenced reform: raise cap on payroll taxes				
<i>Proportion who disagree</i>				
Uninformed	18.2	-1.4 (2.1)	3.6 (2.3)	6.7 (2.4)
Informed	31.6	-12.6 (5.7)	-7.0 (6.1)	-12.2 (5.8)
<i>Proportion who are not sure</i>				
Uninformed	27.7	5.5 (2.8)	6.4 (2.8)	3.3 (2.8)
Informed	4.3	5.4 (4.3)	2.8 (4.3)	3.4 (4.2)
<i>Proportion who agree</i>				
Uninformed	54.2	-4.1 (3.0)	-10.0 (3.0)	-10.0 (3.0)
Informed	64.2	7.3 (6.4)	4.2 (6.8)	8.9 (6.5)

Note. The table summarizes the results of a multinomial logit regression of support for each referenced reform. Column 1 gives predicted probabilities of support under the no-cues treatment for a hypothetical individual with median characteristics. Columns 2-4 give point estimates of treatment effects. Standard errors for the treatment effects are provided between parentheses.

The opinion of the deliberative panel of ordinary citizens was generally more likely to influence uninformed respondents than informed respondents. When the referenced reform was raising the retirement age, the direction of the change among uninformed respondents depended on whether they were

exposed to conflicting cues. When told that the citizens' panel supported the policy and told nothing about the conflicting stance of the Congressional commission, uninformed citizens became considerably less likely to disagree with the reform (by 5.8 percentage points). If they received conflicting cues, uninformed respondents became less likely to disagree (by 5.6 and 1.3 percentage points when the panel supported and opposed the policy, respectively), but only significantly so when the citizens' panel supported the reform. Under no circumstance did we find statistically significant changes in disagreement among informed respondents, although we did find a statistically significant 6.7 percentage point decline in the likelihood of being unsure when the citizens' panel supported raising the retirement age and there were no conflicting cues.⁷ These findings are largely consistent with hypothesis H3.

We found different results when the reform in question was an increase to the cap on payroll taxes. In the absence of conflicting cues, learning the position of the citizens' panel did not significantly influence the extent of disagreement with the proposed reform among uninformed respondents. These respondents, however, were more likely to disagree with the proposed change when exposed to conflicting cues (by 3.6 and 6.7 percentage points when the citizens' panel supported and opposed the policy, respectively). Consistently with H3, the increase in disagreement with the referenced policy was stronger when citizens opposed and elites supported than in the opposite situation (where the change was smaller in magnitude and statistically insignificant). We also found that exposure to conflicting cues caused uninformed respondents to become more unsure about raising the cap on payroll taxes (although the change was only statistically significant when citizens supported and elites opposed), not unlike what we found when the reform in question was raising the retirement age.

Informed respondents, in turn, became considerably less likely to disagree with raising the cap on payroll taxes after learning about the position of the citizens' panel—by 12.6 percentage points in the absence of conflicting cues. The decline in disagreement translated into an increase in the proportion of informed respondents who either agreed or were unsure about their support for the reform. More surprisingly, the rate of disagreement also declined when informed respondents were exposed to conflicting cues (by 7.6 and 12.12 percentage points when the citizens' panel supported and opposed the policy, respectively, although only the latter change was statistically significant).⁸ With the exception of these changes, we found that information about

⁷The decline in uncertainty seen among informed respondents translated into a 7.5% increase in agreement, although this increase was not statistically significant. We observed a similar change when using the alternative (less strict) information indicator, although in that case the increase in agreement was significantly different from zero (see Table A11 in the Supplementary Appendix).

⁸When using the alternative (less strict) information indicator, we again found that conflicting cues lead to a reduction in disagreement among informed respondents, but the estimated effects are considerably smaller in size and lack statistical significance (see Table A11 in the Supplementary Appendix).

the views of the panel of ordinary citizens had no significant effects on opinions held by informed respondents. Thus, the results support hypothesis H₃, but subject to a qualification: the different types of cues mostly affected uninformed respondents, but in some cases the direction of the changes deviated from expectations.⁹

Discussion

The theory behind deliberative minipublics is that they reveal the counterfactual, enlightened public opinion that would materialize if the public were as well-informed and had the same opportunity for meaningful deliberation as the participants in the deliberative minipublic. If this theory is correct, then one reason to organize deliberative minipublics is that their policy conclusions provide observers with informative signals of what they themselves would conclude about a policy, if they were as well-informed and had the same opportunity for meaningful deliberation as the participants in the deliberative minipublic. If observers revise their policy opinions accordingly, then deliberative minipublics can influence the political process and bring its outcomes into closer alignment with enlightened public opinion. Indeed, this is the primary mechanism by which deliberative minipublics like the CIR aim to influence policy.

The experiment provided some evidence for the overarching hypothesis that the typical American would be prepared to understand deliberative minipublics in these terms. Respondents who were told that the deliberative citizens' panel supported an increase in the retirement age were indeed less likely than their counterparts in the control condition to disagree with the proposed change, as one would expect of observers who were prepared to treat the panel's policy conclusions as informative signals of their own counterfactual, enlightened opinions. The experimental results were not, however, uniformly supportive of the theory: no analogous effect was found when the treatment concerned an increase in the cap on payroll taxes rather than the retirement age. Moreover, there was no evidence that either treatment affected which possible change to Social Security respondents would most prefer to implement.

⁹We conducted additional analyses to determine whether the influence of deliberative cues was weaker among subsets of respondents with more at stake in the reforms to the Social Security program. These analyses produced mixed results. We found, for instance, that respondents on the verge of retirement were less likely to revise their level of disagreement with raising the retirement age in response to receiving deliberative cues, compared with younger respondents (see Table A13 in the Supplementary Appendix). However, individuals having more to lose were not always more resistant to revising their views. The reactions of wealthier respondents, for example, did not much differ from those observed among their less well-off counterparts when offered cues concerning citizens' support for raising the cap on payroll taxes (see Table A15 in the Supplementary Appendix).

One possible explanation for the mixed results is that the underlying theory is wrong, at least as applied to the American public: the typical American may not be ready to treat the policy conclusions of a representative, deliberative panel of ordinary citizens as evidence of what his/her own enlightened policy opinions would be. This posture could result from various causes. He/she may fail to reflect on the inferences to be drawn from the observation that a representative sample of citizens, after deliberating and informing themselves about a policy, came to a different conclusion from his/her own. He/she may react negatively to and resist transparent attempts to persuade or manipulate his/her opinions (Knowles & Linn, 2004; Wood, 2000). He/she may not trust the judgments of his/her fellow citizens.

One mark against these explanations is that they fail to explain why the description of the deliberative citizens' panel's conclusion about the retirement age *did* appear to have the predicted effect. An alternative explanation, which squares with this finding, rests on the reasonable assumption that how significantly an observer revises his prior beliefs on learning the policy conclusion of a deliberative minipublic depends on the policy in question. Respondents may have had weaker prior opinions about raising the retirement age than raising the cap on payroll taxes. A rational Bayesian, who treats the policy opinions of a deliberative minipublic as evidence of the policy opinions he/she should hold, may make smaller adjustments to his/her prior beliefs, the stronger his/her prior beliefs—that is, the closer the “subjective probabilities” attached to a proposition are to 0 or 1. This explanation could account for why messages concerning the retirement age affected respondents' propensity to support an increase to it, but the messages concerning the cap on payroll taxes had no analogous effect.

There might be a similar explanation for why the treatments had no significant effects on respondents' *preferences* over the reforms, as opposed to their *approval* of a given reform. On learning that a deliberative minipublic recommends an increase in the retirement age, an observer might come to agree with the recommendation, yet not revise his/her preferences over this proposed reform and alternative proposed reforms, even if he/she treats the policy opinions of a deliberative minipublic as evidence of his/her own counterfactual, enlightened opinions. For he/she may be relatively certain in his/her prior opinions about the *relative desirability* of the reforms, but uncertain in his/her prior opinion about the merits of raising the retirement age.

Finally, a plausible explanation for the mixed results is that the treatments may have been administered at too low a dosage, as it were. Respondents were told that a representative panel of ordinary citizens, after gathering evidence, listening to expert testimony, and deliberating, had come to support a policy. But it was up to the respondents to conjure up a mental picture of this event—a type of rare event which does not belong to the normal political

process and with which all or almost all of the respondents must have been unfamiliar. Moreover, respondents were given little information about this unusual event: they could only speculate about the quality of deliberation, the selection of the experts, just how representative the panel was, etc. The experiment's minimal intervention models the hypothetical scenario of a citizen who acquires only minimal information about a deliberative minipublic. Clearly, in any real-world situation, some citizens will acquire much more information than the respondents in our survey. But many citizens will acquire no more or even less information. Recall that in the case of the BC Citizens' Assembly, 40% of voters appeared to know nothing about the deliberative minipublic at the time of the referendum, and in the case of the Oregon CIR, 46% reported being "not aware" and 35% only "somewhat aware" of the CIR before the 2014 election (Cutler et al., 2008; Gastil et al., 2015). Determining the effects of acquiring only minimal information about a deliberative minipublic is therefore not without interest.

Nonetheless, future experimental work on the topic should expose respondents to a richer, more informative, and more credible account of what deliberative minipublics are, how their members are recruited, and how the deliberative process operates. The results of the present study justify administering "stronger dosages" of the treatment in future experiments, to determine whether the theory's predictions are valid for observers who learn more about deliberative minipublics than the survey respondents in our study did. Together with the results of the present study, evidence of effects when survey respondents receive more information about the deliberative minipublic would bolster the claims of activists and scholars who expect increased publicity and outreach to strengthen the electoral impact of real-world deliberative minipublics like the Oregon CIR (Gastil et al., 2015, p. 66).

Conclusion

The abstract idea of a deliberative democracy, in which collective decisions result not from the aggregation of citizens' unexamined preferences but instead from deliberation among free and equal citizens, holds wide appeal. But how can actual democracies be brought closer to this ideal? Deliberative minipublics strike many political theorists and activists as one possible solution. Small, representative groups of ordinary citizens can be brought together and given the resources to inform themselves and deliberate about a policy problem. They can be formally authorized to choose policies or to set policy agendas, as with the British Columbia Citizens' Assembly, or their conclusions can be conveyed to public officials or the public at large, as with DPs or the CIR. In the latter case, their capacity to affect public policy depends on their ability to influence observers' policy opinions.

The question of whether they have the ability to affect public opinion has received insufficient attention, given the interest deliberative minipublics hold for normative political theorists and activists. The study reported in this article is a step toward answering this question in the case of the American public. We conducted a nationally representative survey experiment where participants were randomly assigned to receive different “signals” of their own counterfactual, enlightened policy opinions: descriptions of the opinions of deliberative panels of ordinary citizens or elites regarding potential changes to Social Security. This policy issue was an ideal context in which to test our hypotheses about the readiness of American citizens to draw inferences from such messages: the issue is important to most citizens, but most of them also do not know enough to have fully developed, confident prior opinions about the merits of potential changes to Social Security.

The results of the survey experiment provide partial support for the hypotheses about how respondents would treat the messages. Learning of a deliberative citizens’ panel’s support for a change to Social Security made respondents less likely to disagree with the proposed change in the case of an increase to the retirement age, but not in the case of an increase to the cap on payroll taxes. By comparison with their better informed counterparts, the relatively uninformed respondents were often more likely to be influenced by the views of the deliberative citizens’ panel, as predicted. But the deliberative cues had hardly any effect on preferences over alternative ways to improve the sustainability of the Social Security program. Most survey respondents strongly preferred raising the cap on payroll taxes over raising the retirement age, and these stances were robust to exposure to the findings of the deliberative citizens’ panel.

On the one hand, the null results point to the limitations of deliberative minipublics’ influence over public opinion when members of the public acquire only minimal information about them. On the other hand, the positive results also suggest that at least for some issues, even observers who know little about deliberative minipublics—only that their members were randomly chosen citizens and deliberated together and heard testimony from experts—may treat their conclusions as informative signals. One explanation for why the experiment turned up negative results for one policy issue (raising the cap on payroll taxes) but positive results for another (raising the retirement age) is that respondents may have had, on average, stronger prior opinions about the former issue.

Another possibility is that some respondents may not have viewed the conclusions of the deliberative minipublic as an informative signal of their own counterfactual, enlightened opinions because they were told too little about the source of the signal. Future research should assess whether observers who learn more about deliberative minipublics than our experimental

subjects did would be more likely to treat their conclusions as informative signals of their own counterfactual, enlightened opinions.

References

- Barabas, J. (2004). How deliberation affects policy opinions. *American Political Science Review*, 98, 687–701. DOI: 10.1017/S0003055404041425
- Barabas, J. (2012). Social security knowledge. In D. J. Lamdin (Ed.), *Consumer Knowledge and Financial Decisions: Lifespan Perspectives* (pp. 217–237). New York, NY: Springer.
- Blaise, A., Carty, K. R., & Fournier, P. (2008). Do citizens' assemblies make reasoned choices? In M. Warren, & H. Pearse (Eds.), *Designing Deliberative Democracy* (pp. 127–144). Cambridge: Cambridge University Press.
- Bullock, J. G. (2011). Elite influence on public opinion in an informed electorate. *American Political Science Review*, 105, 496–515. DOI: 10.1017/S0003055411000165
- Carson, L., Gastil, J., Hartz-Karp, J., & Lubensky, R. (Eds.). (2013). *The Australian Citizens' Parliament and the Future of Deliberative Democracy*. University Park, PA: Pennsylvania State University Press.
- Cook, F. L., Barabas, J., & Page, B. I. (2002). Invoking public opinion: Policy elites and social security. *Public Opinion Quarterly*, 66, 235–264. DOI: 10.1086/340025
- Crawford, V. P., & Sobel, J. (1982). Strategic information transmission. *Econometrica*, 50, 1431–1451. DOI: 10.2307/1913390
- Cutler, F., Johnston, R., Carty, K. R., Blais, A., & Fournier, P. (2008). Deliberation, information, and trust: The British Columbia Citizens' Assembly as agenda setter. In M. Warren, & H. Pearse (Eds.), *Designing Deliberative Democracy* (pp. 166–191). Cambridge: Cambridge University Press.
- Dahl, R. (1989). *Democracy and Its Critics*. New Haven, CT: Yale University Press.
- Farrell, J., & Rabin, M. (1996). Cheap talk. *The Journal of Economic Perspectives*, 10, 103–118. DOI: 10.1257/jep.10.3.103
- Fishkin, J. S. (1995). *The Voice of the People: Public Opinion and Democracy*. New Haven, CT: Yale University Press.
- Fishkin, J. S. (2009). *When the People Speak: Deliberative Democracy and Public Consultation*. Oxford: Oxford University Press.
- Fishkin, J., Kousser, T., Luskin, R. C., & Siu, A. (2015). Deliberative agenda setting: Piloting reform of direct democracy in California. *Perspectives on Politics*, 13, 1030–1042. DOI: 10.1017/S1537592715002297
- Fung, A. (2003). Recipes for public spheres: Eight institutional design choices and their consequences. *Journal of Political Philosophy*, 11, 338–367. DOI: 10.1111/1467-9760.00181
- Gastil, J. (2000). *By Popular Demand: Revitalizing Representative Democracy Through Deliberative Elections*. Berkeley, CA: University of California Press.
- Gastil, J. (2014). Beyond endorsements and partisan cues: Giving voters viable alternatives to unreliable cognitive shortcuts. *The Good Society*, 23, 145–159. DOI: 10.5325/goodsociety.23.2.0145

- Gastil, J., & Richards, R. (2013). Making direct democracy deliberative through random assemblies. *Politics and Society*, 41, 253–281. DOI: 10.1177/0032329213483109
- Gastil, J., Knobloch, K. R., & Richards, R. (2015). Empowering voters through better information: Analysis of the citizens' initiative review, 2010–2014. Report prepared for the Democracy Fund. University Park: Pennsylvania State University. Retrieved from: <http://sites.psu.edu/citizensinitiativereview/wpcontent/uploads/sites/23162/2015/05/CIR-2010-2014-Full-Report.pdf> (last accessed: December 4, 2016)
- Gastil, J., Knobloch, K., Reedy, J., Henkels, M., & Cramer Walsh, K. (2013). Infusing micro- level deliberation into macro-level politics: Assessing the impact of the 2010 Citizens' Initiative Review on Oregon voters. Unpublished manuscript.
- Goodin, R. E., & Dryzek, J. S. (2006). Deliberative impacts: The macro-political uptake of mini-publics. *Politics and Society*, 34, 219–244. DOI: 10.1177/0032329206288152
- Jerit, J., & Barabas, J. (2006). Bankrupt rhetoric: How misleading information affects knowledge about Social Security. *Public Opinion Quarterly*, 70, 278–303. DOI: 10.1093/poq/nflo10
- Karpowitz, C. F., & Mendelberg, T. (2011). An experimental approach to citizen deliberation. In J. N. Druckman, D. P. Green, J. H. Kuklinski, & A. Lupia (Eds.), *Cambridge Handbook of Experimental Political Science*. Cambridge: Cambridge University Press.
- Knowles, E. S., & Linn, J. A. (2004). *Resistance and Persuasion*. New York, NY: Routledge.
- Lupia, A. (1994). Shortcuts versus encyclopedias: Information and voting behavior in California insurance reform elections. *American Political Science Review*, 88, 63–76. DOI: 10.2307/2944882
- Lupia, A., & McCubbins, M. D. (1998). *The Democratic Dilemma*. Cambridge: Cambridge University Press.
- MacKenzie, M., & Warren, M. E. (2013). Two trust-based uses of minipublics in democratic systems. In J. Parkinson, & J. Mansbridge (Eds.), *Deliberative Systems: Deliberative Democracy at the Large Scale* (pp. 95–124). Cambridge: Cambridge University Press.
- Merkle, D. M. (1996). Review: The national issues convention deliberative poll. *Public Opinion Quarterly*, 60, 588–619. DOI: 10.1086/297775
- Sturgis, P., Roberts, C., & Allum, N. (2005). A different take on the deliberative poll: Information, deliberation, and attitude constraint. *Public Opinion Quarterly*, 69, 30–65. DOI: 10.1093/poq/nfio05
- Warren, M. E., & Gastil, J. (2015). Can deliberative minipublics address the cognitive challenges of democratic citizenship? *The Journal of Politics*, 77, 562–574 DOI: 10.1086/680078
- Warren, M., & Pearse, H. (Eds.). (2008). *Designing Deliberative Democracy*. Cambridge: Cambridge University Press.
- Wood, W. (2000). Attitude change: Persuasion and social influence. *Annual Review of Psychology*, 51, 539–570. DOI: 10.1146/annurev.psych.51.1.539
- Zaller, J. (1992). *The Nature and Origins of Mass Opinion*. Cambridge: Cambridge University Press.