



Means, Motive, & Opportunity in Becoming Informed About Politics: A Deliberative Field Experiment

Kevin M. Esterling (UC Riverside)
Michael A. Neblo (Ohio State)
David M.J. Lazer (Harvard)

PNG Working paper No. PNG07-006

Research on political knowledge typically measures citizens' ability to recall political information on the spot. Most citizens appear appalling ignorant using such criteria. We argue that, for Madisonian and deliberative theories of democracy, this "pop quiz" conceptualization is much less normatively relevant than citizens' capacity to *become* informed. However, we know relatively little about the temporal dynamics of citizens' motivation and capacity for learning. In the summer of 2006 we conducted a field experiment in which we recruited current members of the U.S. Congress to discuss immigration policy with random samples of their constituents. We develop an innovative statistical method to identify average treatment effects from field experiments, and find that constituents demonstrate a strong capacity to become informed in response to this opportunity. Participants learned directly from the sessions, but were also motivated to process the study's briefing materials more diligently than non-deliberating subjects, and to increase their attention to politics outside the context of the experiment. Perhaps most encouragingly, this capacity seems to be spread widely throughout the population, in that it is unrelated to demographics and prior political knowledge.

This project is funded by a grant from the Digital Government Program of the NSF (award number IIS-0429452). This paper is prepared for presentation at the Annual Meetings of the American Political Science Association, Chicago, Ill., Aug. 30 to Sept. 2, 2007. We presented an earlier version at the Program on Network Governance, Kennedy School of Government, Harvard University, May 2007, and thank the participants for very useful comments.

Keywords: deliberation, experiment, e-government

Since at least Converse (1964) and Stokes and Miller (1962), survey research has painted a rather grim picture of the average citizen's knowledge of politics. From these findings, many scholars draw strong normative conclusions regarding the health of contemporary democracy and its capacity to reach egalitarian ideals (e.g., Luskin 1990: 333). Researchers in this tradition advance what, on first glance, appears to be a persuasive claim: citizens must possess substantial factual knowledge about government, policies, and politicians in order to judge political actors and to be able to induce effective accountability (e.g., Delli Carpini and Keeter 1996: 65). In survey after survey, however, the typical American scores poorly across a variety of political knowledge measures. Moreover, the United States is hardly alone in this regard (Almond and Verba 1963). It follows that the promises of contemporary democracy are realized for the privileged few at the top of society's knowledge pyramid, while for most citizens democracy amounts only to "tragedy or farce" (Delli Carpini and Keeter 1996: 60).

We argue that such pessimistic normative conclusions derive from a hasty application of normative theory to survey methods. In contrast to this pessimistic vision, both Madisonian and deliberative democratic theory place less emphasis on a static picture of citizens carrying around factual knowledge with them in order to enforce accountability uniformly. Instead, citizens begin with inchoate opinions about the relative priority of policy problems. They expect their representatives to gather information and formulate concrete solutions, exercising appropriate judgment in an iterated division of political labor (e.g., Habermas 1996:357, Hamilton, Madison, & Jay 1961, Bianco 1994, Neblo 2005). For accountability to be effective, citizens need to possess the *capacity* to monitor their representative on the potentially rare occasions when they choose to gear up on an issue (Esterling 2004). It is this (often latent) capacity that causes representatives to exercise judgment on their constituents' behalf (Pitkin 1967: 222-3). On this view, standard general knowledge measures are not well tailored to the normative questions regarding representation and accountability (Prior and Lupia, forthcoming). As a result, the typical citizen's capacity to meet democratic ideals remains an open question.

In the summer and fall of 2006, we conducted deliberative field experiments wherein current members of the U.S. Congress discussed immigration policy with their constituents via a web-based platform. In the experiment, we randomly assigned constituents to experimental conditions as well as to control groups, and test whether deliberating¹ with a current member of Congress affects, among other things, constituents' factual knowledge about the topic of immigration. In all, twelve current members participated in twenty deliberative sessions, along with two sessions wherein a neutral expert on immigration substituted for the Member as a kind of control condition.² To our knowledge, we are the first in political science to involve sitting members of Congress in field experiments.

To measure changes in issue knowledge, we administered baseline and follow-up surveys to deliberators as well as to constituents randomly assigned to two control groups. A well known inferential problem arises with field experiments when subjects do not perfectly comply with their randomized treatment. Treatment and control groups are not directly comparable to the extent that those who comply with the treatment differ systematically from those who do not comply. To address this problem, we propose the *generalized endogenous treatment* model (or the GET model) which accommodates both self selection into the treatment and non-random attrition by failing to respond to the follow-up survey. The GET model identifies the average treatment effect in a multisite experiment by allowing the dependent variable to correlate with a subject's decision both to take up the treatment and to follow through on the post-treatment measurement, where the correlation can occur at both the individual and the site (i.e., congressional district) level.

We find that constituents demonstrate a strong capacity to become informed in response to the opportunity to discuss an issue with their Representative. Such gains stem from three separate mechanisms: participants learned directly from the sessions, but were also motivated to process the study's briefing materials more diligently than non-deliberating subjects, and to increase their attention to politics outside the context of the experiment. Perhaps most encouragingly, this capacity seems to be spread widely throughout the population, in that it is unrelated to demographics and prior political

knowledge. On Madisonian and deliberative grounds, then, our results suggest that contemporary democracy has somewhat more solid underpinnings than previously thought.

Different Theories, Different Standards

One's theory of representation maps a set of empirical findings regarding political knowledge to a normative assessment of contemporary democracy. Much, if not all, of the political knowledge literature assumes, either implicitly or explicitly, an individualist and outcome based theory of representation. On this setup, citizens are charged with retrospectively monitoring the performance of parties and elected officials for whether political outcomes were consistent with the individual's preferences (see, e.g., Delli Carpini and Keeter 1996: 55, Luskin 1990: 333, Converse 1964: 240-1, and Stokes and Miller 1962: 532). Process considerations and their Representatives' rationales play little or no role. On the standard view, citizens cannot exercise accountability effectively unless they are well informed regarding governmental structures, policies, and political actors (Delli Carpini and Keeter 1993, 1996: 10, 65). Thus, to measure citizens' capacities as participants in a democracy, the survey researcher need only administer a set of political knowledge questions in a "pop quiz" format, and then assess the distribution of responses.

Of course, as with any topic, these findings and their implications for democracy have been subject to debate. Some controversies center on the content validity of items (Delli Carpini and Keeter 1993; Graber 1994), the reliability of the measures (Luskin 1987; Mondak 2001), and the extent of cross sectional variation in knowledge across demographic classes (Luskin 1990). Another controversy centers on the kinds and quantity of factual information citizens must hold in long term memory in order to make retrospective evaluations effectively. One side holds that cognitive shortcuts such as cueing, heuristics and online information processing are typically adequate for citizens to make reasonable retrospective judgments (e.g., Lodge et al 1989, Lupia 2006). The other side argues that such shortcuts are far less effective than advocates claim (Delli Carpini and Keeter 1996: 269).

All parties, though, seem to agree, at least implicitly, that retrospective evaluations of individual preference-policy conformity should inform the methods for measuring and assessing citizen knowledge

(Thompson and Bell 2006). The debates in this literature do not revolve around alternative conceptions of democracy, but instead whether or not Americans measure up to the normative expectations implicit in individualist-outcome conception.

However, such a conception is not the only, or even most interesting, conception from a normative point of view. In Madisonian theory, for example, citizens delegate the responsibility of gathering information and exercising judgments to their elected representatives (Hamilton, Madison, & Jay 1961: 82, Pitkin 1967: 207-255). A representative must blend diverse perspectives from within a district, between local and national interests, and often among competing experts (Pitkin 1967: 215). Citizens assess their representatives' performance primarily on the quality of their representative's rationales, not directly on whether the outcome matched her preferences (Esterling 2004: 35-8). Though deliberative theory might seem to require large amounts of "pop quiz" style knowledge, institutionally sophisticated versions do not. For example, Habermas, in developing the institutional implications for his theory of deliberative democracy, makes a similar distinction, conceding the need for representatives in a division of labor, and focusing more on reasoning and process than on preference-policy congruence in evaluating such representation (Habermas 1996; Neblo 2005).

In both Madisonian and deliberative theory, then, citizens recognize that representation is more complex than simply assessing whether representatives have delivered the goods. Instead of merely assessing policy congruence, citizens judge the process, and whether the representative has a compelling rationale to justify her actions (Pitkin 1967: 222-3). Lest one think that such notions are merely naïve normative hopes, the scholarly literature on Congress widely reports that members of the U.S. Congress believe that they are actually held to this Madisonian /deliberative standard more than to policy congruence; they feel they have leeway to exercise their judgment as long as they can explain their decisions adequately (Fenno 1978:231, Kingdon 1989: 248).

On these alternative conceptions, what motivates representatives to exercise good judgment on their constituents' behalf is citizens' latent capacity to assess the quality of their representatives' judgments. When citizens possess the capacity to judge the reasoning of political actors, this can serve as

a credible threat to their future in office even if the capacity often remains latent (Arnold 1990: 68, Kingdon 1989: 31). If the process is working well, citizens do not need to second guess their representative constantly, and representatives typically make informed judgments that advance their constituents' interests (Bianco 1994, Habermas 1996: 341).³ Thus, "pop quiz" measures would be, at best, weakly and indirectly relevant to the health of a democracy on the Madisonian and deliberative accounts. The critical question on this view is, instead, citizens' willingness and capacity to learn about an issue if the occasion arises (Gilens 2001). Appropriate surveys, then, would measure this capacity by observing changes in knowledge that coincide with an occasion for citizens to learn (Fishkin and Luskin 2005, Gastil and Dillard 1999, Kuklinski et al 2001, Thompson and Bell 2006, Prior and Lupia, forthcoming). In the summer of 2006 we conducted a field experiment that created such an occasion by giving a random sample of citizens the opportunity to interact with their current member of Congress on an important and controversial issue (immigration reform) as a part of a small deliberative group.

Experimental design and participant recruitment

Twelve Members of Congress conducted either one or two online deliberative sessions with a random sample of their constituents. The number of participants in each session ranged from eight to thirty constituents. The topic of each session was immigration and border security policy, and the discussions lasted for thirty-five minutes. Constituents participated by typing comments or questions into the online discussion platform. The questions and comments were posted to a queue visible only to a moderator. The moderator, in turn, posted them to the whole group in roughly the order they were received.⁴ The member responded to the questions and comments through a telephone linked to a computer. Constituents listened to the Member's responses over their computer speakers, and also could choose to read the Member's responses via real time transcription. After thirty-five minutes, the Member logged off and constituents were directed to a chat room to have an open ended discussion about the member's responses, and immigration more generally. The chat lasted twenty-five minutes.

The Congressional Management Foundation, a non-profit, non-partisan organization recruited the Members to participate in the study. The group contained a diverse mix of Members. There were five

Republicans and seven Democrats, spread across geographical regions, two women, an African-American, and representatives of both parties' leadership. Knowledge Networks, an online survey research firm, recruited constituent participants from the corresponding congressional districts and administered the surveys. Each subject was asked to complete a pretest survey, and then was randomly assigned to one of three groups: a control group, an information only group, or a deliberative group. In the pretest, participants were told the date and time of the session(s) for their congressional district, and asked if they were willing and able to participate. Those who replied "no" were randomly assigned to either the information only condition or to the true control condition.

Constituents assigned to the information only group read background information on immigration policy based on Congressional Research Service and Congressional Budget Office reports, edited for brevity and reading level. Afterwards, they filled out a short background materials (BGM) survey. Those assigned to the deliberative group also read this information and took the BGM survey. In addition, they engaged in one deliberative session with their Member of Congress and the post session chat.⁵ Those in the true control group were not exposed to either the background information or to a deliberative session. A week following the sessions, KN administered a follow up survey to subjects in all treatment arms.⁶ That is, all constituents in a given congressional district, whether they were assigned to the treatment or to a control condition, received the follow up survey at the same time. KN also administered a final survey in November, after the 2006 congressional elections.

In total, we assigned 2275 constituents to the three experimental conditions, of which 1176 completed the follow up survey. A total of 438 subjects participated in an online deliberative session. There are two significant ways that the data collection deviated from an ideal randomized experiment. First, while subjects' *assignments* to treatment and control are truly randomized, in this sort of field experiment those assigned to the treatment have the ability to self-select out of the treatment; that is, we do not have the capacity to compel participants to show up for their deliberative sessions. To a non-trivial extent, subjects are effectively selecting themselves into the various conditions, and so the treatment cannot be considered fully exogenous. For example, in some cases, subjects assigned to the treatment

read the background materials, skipped the deliberative session, then completed the follow up survey. In this case, subjects effectively re-assigned themselves to the information only condition. This is the problem of non-random “compliance.” Second, some subjects from each of the three groups, including those who participated in a deliberative session, chose not to respond to the follow up survey, selecting themselves out of the dataset. This is the problem of non-random “attrition” (or non-response). Thus, we need an analytical method that can accommodate these realities of self-selection within our study.

Methods

The average treatment effect (ATE) is the causal parameter of interest for evaluating the treatment’s impact on a typical subject. In experimental research, the ATE is a *counterfactual*, comparing the participant’s response having received the treatment, to her response had she not received the treatment (Rubin cite). True random assignment ensures that selection into the treatment is uncorrelated (in expectation) with both the outcomes and the participants’ attributes. If the treatment is truly random, the ATE counterfactual is simply the average difference between treatment and control groups.

In practice, however, self-selection can introduce systematic differences between treatment and control groups. In this case, the appropriate statistical model for identifying the ATE depends on one’s assumptions about the nature of the selection processes (Cameron and Trivedi 2005). The statistical model we propose, the generalized endogenous treatment model (or GET model), explicitly models participants’ unobserved propensity to comply with the treatment (non-random compliance), to select themselves out of the dataset (non-random attrition), and allows for the levels of selection and the intensity of the treatment to vary across sites in the study. The GET model estimates the correlations among the equations modeling each respondent’s probability of taking up the treatments, responding to the follow up survey, and the dependent variables, all of which are taken to be endogenous. The GET model builds on Aakvik, Heckman and Vytlacil (2005 [AHV]) and Miranda and Rabe-Hesketh (2006).⁷

Formalization: we conducted the experiment **E** to assess the effects of the deliberative treatment on a set of outcomes, **O**. Cells in this experiment are defined by a partitioned vector $\mathbf{E}^i = (\mathbf{P}^i, \mathbf{Z}^i, \mathbf{X}^i, S)$ with elements of **E** defined as follows:

A vector $\mathbf{P} = (P_A=a, P_{B|Z}=b, P_{D|Z}=d, P_{R|Z}=r, P_N)$ of endogenous variables of length five indicates five separate, but potentially dependent, self-selected choices that affect participants' treatment compliance and response status. The filter variable P_A indicates whether the potential participant would participate in any part of the study. $P_A=0$ if the participant expressed both a willingness and an ability to do sessions at a specific time if assigned to the deliberative arm of the study. $P_A=1$ if the participant expressed only a willingness or ability to fill out the surveys as a part of the study.⁸ Potential participants who indicated they would not participate in any manner were dropped from the study.

The vector \mathbf{P} also contains variables indicating compliance with treatments and response on surveys, some of which are conditioned on randomized treatment assignment vector \mathbf{Z} described below. The variable $P_{B|Z}$ indicates the participant read the background materials and responded to the BGM survey ($P_{B|Z}=1$ if completed the BGM survey, 0 if not); the variable $P_{D|Z}$ indicates compliance with the deliberative session ($P_{D|Z}=1$ if complied with the deliberation, 0 otherwise); the variable $P_{R|Z}$ indicates response on the follow up survey ($P_{R|Z}=1$ if responded, 0 if not); and the variable $P_N=1$ if the participant responded to the November survey, and 0 otherwise. Collectively, we label the \mathbf{P} vector the "participation" variables. As we discuss next, $P_M, M \in \{B|Z, T|Z, R|Z, N\}$ are constrained to zero for some participants depending on their treatment arm assignments and their participation history in the experiment.

The exogenous vector $\mathbf{Z}' = (Z_D, Z_B, Z_C)$ of length 3, with $Z_i \in \{0,1\}$ indicates participants' randomized treatment arm assignment into one of four mutually exclusive conditions: deliberation ($Z_D=1$ if assigned to deliberation, 0 if not), background information only ($Z_B=1$ if assigned to background information only condition, 0 if not), true control ($Z_C=1$ if assigned to the true control condition, 0 if not).⁹ $\mathbf{Z}=\mathbf{0}$ if the participant who passed the filter ($P_A=a$) was randomly excluded from the study.

The elements of \mathbf{P} are constrained to zero for some participants in the study depending on their randomized assignments. $P_{B|Z}=0$ if the participant is assigned to the true control condition ($Z_C=1$); $P_{D|Z}=0$ if the participant is not assigned to the deliberation condition ($Z_D=0$). In addition, due to limited

resources, some participants were dropped from the study based on their participation history: the participant was not administered the follow up survey ($P_{R|Z}=0$) if she was assigned to either the deliberation or the information only condition ($Z_B=1 \mid Z_D=1$) and she participated in none of her assigned treatments ($P_{B|Z}=0$ & $P_{D|Z}=0$). Participants were not administered a November survey ($P_N=0$) if they failed to respond to the follow up survey ($P_{R|Z}=0$). In each case where a participation variable is constrained to zero, we say the participant is not eligible for that treatment or survey.

A vector \mathbf{X} of length j contains exogenous pre-treatment variables. All participants willing to participate ($P_A=a$) were administered a baseline survey that captured pretreatment data, including values of the outcome variables. Define the vector \mathbf{O}^{PT} of length K , a subvector of \mathbf{X} , to be the pretreatment values of the outcomes corresponding to those measured in \mathbf{O} . Define the subvector $\mathbf{X}_{\mathbf{O}}$ to be the remaining variables in \mathbf{X} . In this application, \mathbf{X} also includes a measured variable for overall political knowledge (see Delli Carpini and Keeter 1993). Note that \mathbf{X} does not include a constant. The variable $S=s$ indicates the site where the experiment took place, in this study defined as the congressional district in which the participants reside.

Identification of Structural Parameters in the GET Model: There are K separate (but potentially dependent) experimental outcomes in \mathbf{O} ; in this application, $K=7$. In estimating treatment effects from an experiment, we wish to compare potential outcomes from the experiment, $\Delta=E(O_{1k} - O_{0k})$, where O_{1k} is defined as the outcome that would be observed if the participant received the treatment, and O_{0k} is the outcome that would be observed if the participant did not receive the treatment (Imai 2006). For each participant, we only observe the outcome corresponding to their actual treatment status. Define the treatment outcome vector \mathbf{O} of length K to be the potential outcomes that are observed for a given participant on the follow up survey, that is, conditional on the treatment she actually received. In addition, the treatment outcome vector \mathbf{O} is observed if and only if $P_{R|Z}=1$.

For this application, the outcomes are dichotomous, $O_k \in \{0,1\}$. We take the realization of O_k to be a function of a latent index variable O_k^* . We wish to estimate a matrix $\theta_K=(\boldsymbol{\alpha},\boldsymbol{\beta})$ of structural parameters in the set of K regressions:

$$O_k^* = \alpha_{0k} + \alpha_{1k}P_D + \alpha_{2k}P_B + \alpha_{3k}O_k^{PT} + \beta_k X_{-0} + \omega_k$$

$$O_k = \begin{cases} 1 & \text{if } O_k^* \geq 0 \\ 0 & \text{if } O_k^* < 0 \end{cases}$$

where we model the probability $P(O_k=1|.)$ using the logit link function and distributional assumptions for ω_k that we state below.

Identifying the structural parameters θ requires accounting for the endogeneity of participants' ability and willingness to participate in the experiment ($P_A=a$), compliance with their assigned treatment ($P_{D|Z}=d, P_{B|Z}=b$), and attrition through non-response on the follow up ($P_{R|Z}=r$), since some participants indicated they were unwilling or unable to participate in a deliberative session, some did not comply with their treatment assignment (that is to read the background material or to read the material and engage in deliberation), and some did not respond to the follow up survey. These selection processes can be jointly correlated with the outcome variables, and hence make treatment and attrition endogenous to the model.

We develop the GET model to identify the structural parameter matrix θ by taking the decisions to comply ($P_A=a, P_{B|Z}=b, P_{D|Z}=d$) and the decision to respond ($P_{R|Z}=r$) as endogenous to the model. Each indicator of participation in \mathbf{P} is dichotomous, and we assume that each is the realization of a latent index process, where P_i takes its value given the index and a set of thresholds. To model \mathbf{P} we estimate five additional regressions.

$$F_A^* = \alpha_{0A} + \beta_A X_{-0} + \omega_A$$

$$F_{B|Z}^* = \begin{cases} \alpha_{0B|Z} + \alpha_{2B|Z}Z_D + \beta_{B|Z}X_{-0} + \omega_{B|Z} & \text{if not assigned to the true control condition} \\ 0 & \text{otherwise} \end{cases}$$

$$F_{D|Z}^* = \begin{cases} \alpha_{0D|Z} + \beta_{D|Z}X_{-0} + \omega_{D|Z} & \text{if assigned to the deliberation condition} \\ 0 & \text{otherwise} \end{cases}$$

$$F_{R|Z}^* = \begin{cases} \alpha_{0R|Z} + \alpha_{1R|Z}Z_B + \alpha_{2R|Z}Z_D + \beta_{R|Z}X_{-0} + \omega_{R|Z} & \text{if eligible for the follow up survey} \\ 0 & \text{otherwise} \end{cases}$$

$$F_N^* = \begin{cases} \alpha_{0N} + \beta_N X_{-0} + \omega_N & \text{if eligible for the November survey} \\ 0 & \text{otherwise} \end{cases}$$

and with

$$P_i = \begin{cases} 1 & \text{if } F_i^* > 0 \\ 0 & \text{if } F_i^* \leq 0 \end{cases}$$

In each case, we model $P(P=1|.)$ using the logit link, again giving distributional assumptions for ω_m , $M \in \{A, B|Z, T|Z, R|Z, N\}$ below.

Notice that the decisions to participate and subsequently to respond are qualitatively different in each arm of the study. To model these participation variables across the treatment arms of our study, we assume that these latent indices have the same scale for participants across the three arms of the study; randomizations ensures this is true in expectation.¹⁰ Where appropriate, the intercepts are conditioned on elements of the exogenous assignment vector \mathbf{Z} .

The GET model uses the behavioral indicators \mathbf{P} to estimate a common latent factor, η_1 , along with a vector of factor coefficients, λ_m , measuring participants' latent tendency to participate in all parts of the experiment.

$$\omega_m = \lambda_m \eta_1 + \varepsilon_m, \quad m \in \{A, B|Z, D|Z, R|Z, N\}$$

For identification, we assume $\eta_1 \sim N(0,1)$, $\varepsilon_m \sim N(0,1)$, and $\text{cov}(\eta_1, \varepsilon_m) = 0$ for all m . Simultaneously, the GET model includes the common factor η_1 as a latent control variable in the outcome equations.

$$\omega_k = \lambda_k \eta_1 + \varepsilon_k, \quad k = 1 \text{ to } 7$$

For identification, we assume $\lambda_k=1$ for $k=1$, and $\varepsilon_k \sim N(0,1)$, $\text{cov}(\eta_1, \varepsilon_k) = 0$ for all k , and

$\text{cov}(\varepsilon_m, \varepsilon_k) = 0$ for all m, k . By including η_1 in both the participation and in the outcome equations, the GET model in effect holds participants' "type" constant as a way to identify the structural parameters.¹¹

In addition, the common latent variable allows estimation of the correlations between all endogenous participation variables \mathbf{P} and the choice variables \mathbf{O} .¹²

$$\rho_{(O_k, P_m)} = \frac{\lambda_k \lambda_m}{\sqrt{(\lambda_k^2 + 1)(\lambda_m^2 + 1)}} \quad k \in \{1 \text{ to } 7\}, m \in \{A, B|Z, D|Z, R|Z, N\}$$

That is, the GET model assumes the decisions to select into the treatment and to select out of the study are correlated with the outcomes. In this sense, GET simultaneously accounts for both endogenous compliance and attrition, which are inherent properties of most data generated in a field experiment.

This GET approach assumes that participants' unobserved tendency to participate, or their "type," can be characterized by a latent variable (see AHV 2005).¹³ Since the GET model controls for participants' type using an estimated latent variable, we assume overlap on this latent variable across the treatment assignment arms \mathbf{Z} of the experiment. This assumption requires that the selection processes measured in \mathbf{P} are not deterministic, so that some participants with a low expectation of participating happen to participate in all aspects of the study, and some with a high expectation do not participate.

Finally, the GET model is flexible enough to account for hierarchical dependencies in the data. Thus, both the distribution of participant type and the treatment effects can vary across sites of the experiment. To capture level 3 variation in the behavioral types of participants across congressional districts, we can decompose η_1 as a function of a level 3 latent variable η_2 , fixed exogenous regressors, and a random component:

$$\eta_1 = \gamma_1 \eta_2 + \gamma_{-1} \mathbf{X}_\gamma + \varepsilon_{\eta_1}$$

where $\eta_2 \sim \mathcal{N}(0, \sigma_{\eta_2}^2)$, $\varepsilon_{\eta_1} \sim \mathcal{N}(0, \sigma_{\varepsilon_{\eta_1}}^2)$, and $\text{cov}(\eta_2, \varepsilon_{\eta_1}) = \mathbf{0}$. To use a fixed effects approach to decompose η_1 across sites of the experiment, simply constrain γ_1 to zero and include district (site) dummies with no constant in \mathbf{X}_γ . In this application, we conducted two sessions with an expert, as opposed to a member of Congress, in March of 2007, and to test the hypothesis that these sessions were different from the member-based sessions, we set $\gamma_1 = \mathbf{0}$ and \mathbf{X}_γ equal to a level 3 dummy variable that equals one for the two expert sessions and zero for the member sessions. In a similar fashion, we can define the structural estimates of the coefficients for the two treatment variables, α_{1k} and α_{2k} , to be random coefficients that vary across congressional districts, or simple interactions with dummy variables indicating each district.

As AHV note (2005: 28), factor models, such as GET, and matching models (see Rosenbaum and Rubin 1985) have a "close affinity." Given η_1 , the matching conditions that require the treatment and outcomes to be orthogonal are satisfied. The latent variable approach solves the limitation that matching can only make treatment and outcome orthogonal conditional on observable variables, since the latent variable model conditions on both observed and unobserved propensities to take up the treatment. In

addition, the GET model extends the factor model in AHV by also accounting for selective attrition and multisite dependencies.

Retrieving Treatment Effect Estimates from the Structural Parameters: To estimate treatment effects, we compare average responses across the arms of the study: those who participated in a deliberative session, those who merely read the background information but did not participate in deliberation, and those who served as true controls. Within the following cells of \mathbf{E} we potentially observe the follow up response on the outcomes \mathbf{O} , that is, $P_{R|Z=r}$ is due only to self-selection.¹⁴

- A. *Assigned to deliberation condition:* The set of participants i assigned to the deliberation condition $\mathbf{E}'=(P_A=1, P_{B|Z}=b, P_{D|Z}=d, P_{R|Z}=r, P_N=n, Z_D=1, Z_B=0, Z_C=0, \mathbf{X}=\mathbf{x}, S=s)$ and complied on the BGM and deliberation $\mathbf{P}'=(1,1, 1, r, n)$ or complied on deliberation but did not comply on the BGM $\mathbf{P}'=(1,0, 1, r, n)$ or complied on the BGM but did not comply on deliberation $\mathbf{P}'=(1,1, 0, r, n)$
- B. *Assigned to information only condition:* The set of participants i assigned to the information only condition $\mathbf{E}'=(P_A=a, P_{B|Z}=b, P_{D|Z}=0, P_{R|Z}=r, P_N=n, Z_D=0, Z_B=1, Z_C=0, \mathbf{X}=\mathbf{x}, S=s)$ and completed the BGM and indicated would do the deliberative session if assigned to deliberation $\mathbf{P}'=(1,1, 0, r, n)$ or completed the BGM and indicated would not deliberate if assigned $\mathbf{P}'=(0,1, 0, r, n)$
- C. *Assigned to the true control condition:* The set of participants i assigned to the true control condition $\mathbf{E}'=(P_A=a, P_{B|Z}=0, P_{D|Z}=0, P_{R|Z}=r, Z P_N=n, D=0, Z_B=0, Z_C=1, \mathbf{X}=\mathbf{x}, S=s)$ and indicated would do the deliberative session if assigned to deliberate $\mathbf{P}'=(1,0, 0, r, n)$ or indicated would not do the deliberative session if assigned $\mathbf{P}'=(0,0, 0, r, n)$

These observations are in the sets, $E_c, c = 1$ to $7, E_1 = \{i: \mathbf{E} = (1,1,1, r, n, 1,0,0, \mathbf{x}, s)\}$,

$E_2 = \{i: \mathbf{E} = (1,0,1, r, n, 1,0,0, \mathbf{x}, s)\}$, $E_3 = \{i: \mathbf{E} = (1,1,0, r, n, 1,0,0, \mathbf{x}, s)\}$,

$E_4 = \{i: \mathbf{E} = (1,1,0, r, n, 0,1,0, \mathbf{x}, s)\}$, $E_5 = \{i: \mathbf{E} = (0,1,0, r, 0,1,0, \mathbf{x}, s)\}$,

$E_6 = \{i: \mathbf{E} = (1,0,0, r, n, 0,0,1, \mathbf{x}, s)\}$, $E_7 = \{i: \mathbf{E} = (0,0,0, r, n, 0,0,1, \mathbf{x}, s)\}$. The vector of effect

parameters $\Lambda = [E(\mathbf{O} | i \in \{E_1 \cup E_2\}; \mathbf{O}^{PT}) - E(\mathbf{O} | i \in \{E_3 \cup E_4 \cup E_5\}; \mathbf{O}^{PT}), E(\mathbf{O} | i \in \{E_1 \cup E_2\}; \mathbf{O}^{PT}) - E(\mathbf{O} | i \in \{E_6 \cup E_7\}; \mathbf{O}^{PT}), E(\mathbf{O} | i \in \{E_3 \cup E_4 \cup E_5\}; \mathbf{O}^{PT}) - E(\mathbf{O} | i \in \{E_6 \cup E_7\}; \mathbf{O}^{PT})]$ defines

comparisons among blocks of experimental cells E_C identified under the GET model assumptions. In words, $\Delta = [\text{Deliberative Group v. True Controls; Deliberative Group v. Information Only Group; Information Only Group v. True Controls}]$.

For these comparisons, we can define both mean and distributional treatment effect parameters (derived in AHV 2005). To simplify notation, define:

$$\mu_k^0 = \alpha_{0k} + (\alpha_{2k} \times 0) + E(X)' \beta_k$$

$$\mu_k^1 = \alpha_{0k} + \alpha_{1k} + (\alpha_{2k} \times 0) + E(X)' \beta_k$$

$$\mu_k^2 = \alpha_{0k} + \alpha_{1k} + (\alpha_{2k} \times 1) + E(X)' \beta_k$$

These are the expected values of each equation's linear component for the true controls, the information only group, and the deliberative group (respectively), for those who did not get the response correct on the background materials survey. The average treatment effect (ATE) mean parameter estimates the average effect of the treatment for all participants in the experiment. The ATE among two arbitrary groups for an outcome equation k is,

$$\Delta_k^{ATE}(\cdot) = \int [\Phi(\mu_k^i + \lambda_k \eta_1) - \Phi(\mu_k^j + \lambda_k \eta_1)] \phi(\eta_1) d\eta_1, \text{ for } i \neq j, i, j \in \{0, 1, 2\}$$

The marginal treatment effect (MTE) mean parameter estimates the treatment effect for participants, given their underlying propensity to participate, and is,

$$\Delta_k^{MTE}(\cdot, \omega_k^*) = \frac{\int [\Phi(\mu_k^i + \lambda_k \eta_1) - \Phi(\mu_k^j + \lambda_k \eta_1)] \phi(\omega_k^* + \eta_1) \phi(\eta_1) d\eta_1}{\phi\left(\frac{\omega_k^*}{\sqrt{2}}\right)}, \text{ } i \neq j$$

The average treatment effect for the treated (ATT) is the effect of the treatment among those who actually received the treatment, and is

$$\Delta_k^{ATT}(\cdot, P_{D|Z} = 1) = \frac{\int [\Phi(\mu_k^i + \lambda_k \eta_1) - \Phi(\mu_k^j + \lambda_k \eta_1)] \Phi(\alpha_{0D|Z} + \beta_{D|Z} X - \sigma + \eta_1) \phi(\eta_1) d\eta_1}{\Phi\left(\frac{\alpha_{0D|Z} + \beta_{D|Z} X - \sigma}{\sqrt{2}}\right)}, \text{ } i \neq j$$

The average treatment effect (ATE) distributional parameter estimates the probability of some positive improvement from the treatment, and is:

$$E[1(\Delta_k^{ATE} > 0) | \cdot] = \int \left[\Phi(\mu_k^1 + (\lambda_{1R}^1 + \lambda_{2R}^1)\eta_1) (1 - \Phi(\mu_k^0 + (\lambda_{1R}^0 + \lambda_{2R}^0)\eta_1)) \right] \phi(\eta_1) d\eta_1, \quad 1 \in \{1, 2\}$$

where 1 is the indicator function. The ATE and MTE parameters are identified even when there is strong selection into treatment and control, since the latent variables accommodate correlations among the treatment compliance, attrition, and outcome equations. In addition, note that the GET model is a flexible generalization of generalized linear models, including simple logit and a single site selection model.

Estimation: We estimate the structural parameters using Bayesian MCMC methods implemented in WinBUGS (called from R2WinBUGS). We standardize the exogenous variables and assign independent diffuse t-distribution priors for the parameters, with 2 degrees of freedom and scale of 10 for equation constants and 2.5 for all other parameters (Gelman et al 2007). Our sample is large enough for the posterior distributions to be robust to alternative priors. We estimate three chains using randomized starting values. The chains in the models implemented below typically converge after a burn in period of 500 iterations. For the results presented below, we draw a total of 2000 simulations from the model, discard the first 1000 (which includes the 500 burn in iterations), and then retain one in every three draws for a total of 1000 draws from the posterior distribution. We use these posterior draws to estimate standard errors of all estimates, including functions of the structural parameters such as correlations coefficients and mean and distributional treatment effect parameters.

Data and Expectations

Table 1 summarizes the variables we use to specify the model. The table is divided into five sections, and the columns indicate whether the variable is measured prior to the study (pretreatment or exogenous) or as a part of the study (endogenous). The first section lists the experimental outcome variables for this study (\mathbf{O} and \mathbf{O}^{PT}). In this paper, we model the probability that a respondent gives a correct answer on a battery of seven items that indicate the respondent's knowledge about immigration policy. Since these responses are dichotomous, the columns give the average probability that a

respondent in our study (across all treatment arms) gives a correct response.¹⁵ We measured the pretreatment values of these items in the baseline survey and the post treatment values in the follow up survey. In general, between 20% and 40% of respondents could correctly answer a given policy question prior to the experiment, and the unconditional probability of a correct answer increases slightly across all questions on the follow up survey. Notably, nearly all of the sample (84% in the baseline survey) could correctly answer the question regarding the citizenship status of a child born in the US to an illegal alien. This level leaves little room for improvement on the follow up survey.

The second section presents the variables that indicate respondents' randomized assignments into the three treatment arms of the study: those assigned to read background material on immigration and then participate in a deliberative session (the "deliberative condition" Z_D); those assigned to read the background material but not participate in deliberation (the "information only" condition, Z_B); and those who were given access to neither the background reading materials nor to a deliberative session (the "true controls", Z_C). Note that participating in deliberation required both that the participant set aside a significant amount of time (one hour), and also to be available at the scheduled time. Thus, in order to have a large enough sample of deliberative participants, we assigned slightly more than half of the participants who made it through the initial filter to the deliberative condition.

Of course, as in any field experiment, assignment to the treatment arms of this study in no way assures that participants actually comply with the treatment, nor can participants be compelled to respond to surveys. The \mathbf{P} variables in Table 1 indicate participants' actual participation in the study, and the GET model uses these variables to measure their unobserved propensity to participate in the study.¹⁶ Since these variables are all dichotomous, the cells give the compliance and response rates. In the baseline survey we informed participants the specific dates and times when the deliberative session(s) would be held for their congressional district. Among those who reported they could not or would not participate in deliberation, we asked if they still would be willing to do surveys as a part of the study. Those who passed this second filter (P_A) were randomized into either the information only group or to the true control

condition. Though we measured this variable prior to any treatment, it directly measures participants' compliance with the deliberative treatment, so we treat it as endogenous.

Thirty-six percent of those assigned to the deliberative group actually participated in their assigned deliberative session ($P_{D|Z}$), and 53% among those assigned either to the information only group or to the deliberative group read the background materials and filled out the post-reading survey ($P_{B|Z}$). Among all arms of the study, 71% responded to the follow up survey ($P_{R|Z}$), and 70% completed the post-election survey administered in November of 2006 (P_N). Given emergent budget constraints, we excluded from the follow up survey participants who were assigned to the information only group or to the deliberative group who did not complete their assigned treatment. Similarly, we only administered the November survey to participants who successfully completed the follow up survey. The response rates are all conditional on the eligible sample for the respective survey.

Expectations: A number of studies assert that citizens tend to be knowledge “generalists,” in that those with a high general knowledge of politics also tend to know more about specific policy topics (e.g., Delli Carpini and Keeter 1996: 270, Gilens 2001). In the model reported below we condition all outcomes on a factor measure of participants' general political knowledge, which we label η_2 . We measure participants' general political knowledge (η_2) in the baseline survey using the “Delli Carpini and Keeter five” (G_k , for $k=1$ to 5), employing a standard item response model (Delli Carpini and Keeter 1993). We simultaneously enter η_2 as an exogenous regressor in each of the seven outcome equations and five participation equations in the full GET model.¹⁷ Luskin (1990) and Nadeau and Niemi (1995) examine demographic determinants of political knowledge, and on the basis of their findings we condition on race, gender, and employment. To get large enough samples for each district, our survey vendor, Knowledge Networks, subcontracted with two other online survey vendors. Since the panels at Knowledge Networks are highly representative, while those of the other firms are less so, we condition on the participant's panel. Finally, and crucially, we control for whether subjects' correctly answered each of the immigration knowledge items on the baseline survey.

We predict that respondents who participate in the treatment condition, all else equal, will tend to have a higher knowledge of immigration policy, and hence a higher probability of a correct response on each of the seven knowledge items, when compared to the information only and true control conditions. There are three possible mechanisms for such effects. First, compared to those in the information only and true control conditions, those in the deliberative condition might receive information in the course of the online deliberation. To test whether the sessions themselves are directly informative, we include a dichotomous variable in the model for each question indicating whether the Member happened to give the answer to the question during the deliberative session.

Second, it is possible that respondents in the deliberative condition will pay closer attention to the background material since they know they are preparing for a discussion with their member of Congress (Kuklinski et al 2001, Hutchings 2001, Luskin 1990). To test this mechanism, we include two interaction variables multiplying the (logged) time spent on the background materials by treatment received. If those in the deliberative group who spent more time with the background material are more likely to get the questions correct, this would indicate that the prospect of interacting with their Member induced greater motivation to learn from the background materials, and more effortful processing.

Third, anticipating an interaction with their member of Congress, subjects might attend more to information external to the experiment, such as media reports (immigration was intensely covered during this time) or from talking to others. To test this, we estimated a series of additional regressions within the GET model to evaluate whether there are treatment effects for participants' attention to media reporting on immigration and the elections; their attitudes toward a duty to be informed; and the number of people external to the study with whom they discussed immigration.

Finally, the GET model can account for heterogeneity of participants and treatment effects across congressional districts (Nadeau and Niemi 1995). In the models reported below, we found no heterogeneity across congressional districts using either fixed or random effect intercepts and coefficients. In two of the districts, however, we used an expert on immigration policy rather than a current member of Congress as a way to test whether citizens respond differently to experts than to Members. We conducted

these expert sessions in March of 2007. In each of the outcome equations below, we include the indicator for expert session (S), and interact it with the treatment the participant received.

Results

Recall that the level 2 latent variable η_1 induces a correlation between the respondent's compliance with the treatment, the respondent's compliance with the follow up, and the respondent's probability of correctly answering each immigration policy item. Thus, the estimated correlations among the equations test for non-random selection both into and out of the study. If compliance and attrition were truly random, then there would be no statistical correlation among the compliance, response, and outcome equations, and the corresponding factor loadings would be indistinguishable from zero.

We find a high degree of selection in both compliance and attrition, and so reject the null hypothesis that participation in the experiment is random. In particular, Table 2 demonstrates a strong negative correlation between selection into the treatment conditions (compliance with the deliberative session and with reading the background materials) and the probability that a respondent correctly answers questions 1, 2, 6, and 7. In addition, there is a strong negative correlation between responding to the follow up survey and these same questions. This means that those who had lower baseline knowledge of immigration policy (controlling for general knowledge) were more likely to select themselves deeper into the study, and to respond to the follow-up survey. Thus, proceeding as if the treatment were exogenous would retrieve biased estimates of the ATE. Below, we compare the GET results with simple logit estimates of the ATE and note their differences.

Notice that there is no correlation between the probability of a correct response and whether the participant was unwilling or unable to participate in the deliberative session, nor with any of the other participation variables (results not reported). This finding indicates that participants who are filtering themselves out of the deliberative treatment up front are doing so for mostly exogenous reasons (e.g., a previous commitment at the time of the deliberative session). Thus, participants who were randomized into the deliberative arm of the study are similar to those who tend to participate in surveys in general,

which implies that the average treatment effects we report below are for a sample comparable to those found in other knowledge surveys (e.g., Delli Carpini and Keeter 1993).

Estimates of the Average Treatment Effects: Figure 1 summarizes the ATE mean parameter (Δ^{ATE}) results. There are two different treatment arms: the information plus deliberation condition and the information only condition. Given this, there are three treatment comparisons across seven questions: deliberative group compared to true controls, deliberative group compared to the information only group, and information only group compared to the true controls. Each of these average treatment effect comparisons are shown in Figure 1 as a vertical bar, with the center box indicating the inter-quartile range of the estimate and the whiskers indicating 90% confidence intervals for the effect parameters.¹⁸

Note first that the estimates for the ATE comparing the deliberative group to the true controls are all positive and substantial. For each item save one, the deliberative treatment increases the probability of a correct answer by between 20% and 45%. On the other item, question 7, the ATE was significant for the deliberative group to true control comparison, but substantively close to zero. Nearly everyone in the study (84%) could answer the item correctly on the baseline survey, so there was little room for improvement. In addition, the ATE comparing the deliberative group to the information only group is positive and significant for questions 1, 4, 6, and 7, and close to significant on questions 2 and 3. The information only group tended to score better than the true controls on questions 2, 5, 6, and 7.

As noted above, the GET model correctly identifies the ATE's by explicitly modeling the correlation between the various forms of self selection and the outcome responses. Since we find that selection into and out of the study is non-random, taking the treatment variable as exogenous is likely to yield biased estimates. We re-estimated the outcome equations taking participation in the two treatment conditions as exogenous, controlling for the exogenous variables (including the pretreatment response). We find (results not shown) that ignoring selection causes one to dramatically underestimate the ATE comparing the discussion group to the true controls by 15% to 58% (as a proportion of the GET estimate) depending on the item, save again for the last item, where the two models yield similar estimates. This

result makes sense since those who select themselves into deliberation, after controlling for their general political knowledge, tend to have less factual knowledge overall about immigration policy.

Figure 1b reports the distributional average treatment effects (Δ^{DATE}), which is the probability that there is some increase in the chance of a correct answer from participating in one or the other of the treatments. These changes in probabilities are significant for each treatment effect comparison. Save for question 7, the effect of deliberation compared to the true controls ranged from about a 25% chance of improvement on question 5 to a 50% chance of improvement on question 6. The comparisons among the other combinations of treatment groups are parallel but of equal or smaller magnitude.

Figure 2 displays the marginal treatment effects, which is the effect of taking up the treatment for a participant at a fixed probability of complying. Figure 2a shows the MTE for a typical respondent with a low probability of participating in the study, 2b shows a typical respondent who is on the cusp between participating and not participating, and 2c shows the MTE for a typical respondent with a high probability of participating. These figures indicate that the MTE is relatively constant across all participant types and across all questions, and consequently the MTE is roughly equal to the ATE. This suggests that the benefit of the treatment is similar for all participants, whether they have a latent tendency to participate or a latent tendency not to participate. That is, those who are more interested in participating in the study, and perhaps more motivated to participate in politics more generally, do not necessarily have a higher capacity for gaining policy knowledge (Prior and Lupia forthcoming). Figure 2d shows the ATT mean parameters, which indicate that there is possibly a slightly higher benefit from the treatment among those who actually participate in the treatment, but the ATT parameter estimates are not significantly different from the ATE estimates, consistent again with the proposition of equal effects among participant types.

Figure 3 shows the effect of general political knowledge on the probability of a correct response on each of the seven questions, holding constant the treatment actually received. Like in other studies (e.g., Delli Carpini and Keeter 1993, 1996:270), we find that participants are “generalists” of a sort, that is, those with a higher general knowledge of politics also tend to know more about immigration, with some variability across the questions.¹⁹ In addition, we find that participants with a higher general

knowledge are more likely to participate in deliberation by 20% ($p < 0.05$). We do *not* find, however, that those with higher general knowledge typically respond better to the treatments than those with low general knowledge (Gilens 2001, Jerit et al. 2006). For questions 1 and 2, those with higher general knowledge appear to gain more from the deliberative treatment than from merely being exposed to the background information. But this pattern is reversed for question 6, and there is no difference between the treatment groups for questions 3, 4, 5 and 7. This suggests, as with the MTE, average participants, even those with low general political knowledge, learn from the deliberative sessions just as well as those with high general knowledge.

Similarly, we found few consistent differences in the probability of a correct responses across race, gender or employment for any of the questions. White respondents were ($p < 0.05$), slightly less likely to answer question 3 (by 6%) and question 5 (by 7%); female respondents were slightly less likely to answer question 1 correctly (by 7%) and question 2 correctly (by 2%) and more likely to answer question 3 correctly (by 4%); and those who are employed full time are less likely to answer questions 3 (by 4%) 4 (by 4%) and 7 (by 5%) correctly. We found that the KN panelists were statistically less likely to answer questions 1 (by 16%), 2 (by 4%), 5 (by 9%) and 7 (by 1%) correctly, which is consistent with the proposition that the KN panels are demographically better representative than the non-KN panels.²⁰

We found very few significant effects at level 3, comparing participant types and treatment effects across the congressional districts. After controlling for the exogenous variables in the model, we find no statistically significant differences in the size of the effects for the deliberative group and the information only group across the congressional districts, which indicates that the sessions were equally informative. We also found no heterogeneity in the responses across the congressional districts, which suggests that constituents are equally motivated to learn about the issue and discuss the issue across the sites. In addition, we found no differences between the March 2007 expert sessions compared to the member sessions held in the summer of 2006. This latter finding suggests that citizens are able to learn from interacting with an expert just as much as their Member of Congress.

Finally, we included several variables in an attempt to disentangle the mechanism by which the treatments lead to an improvement in policy knowledge. The first mechanism posits that respondents who participate in deliberation are exposed to information from their Member (or the expert) during the course of the discussion. To test this mechanism, we included a dichotomous variable in each question response equation, indicating whether that item was mentioned in the respondent's session. Here we found substantial effects, with an increased chance of a correct answer on question 1 (by 27%), question 4 (by 21%), question 5 (by 40%), and question 7 (by 3%), all statistically significant at $p < 0.05$.

The second mechanism posits that those who are given the opportunity to deliberate with their member (or, a policy expert) tend to be more motivated to learn the background material. We test this second mechanism by interacting the treatment the participant actually received (deliberative or information only) with the log of the time she spent reading the background materials and taking the BGM survey. If the interaction coefficient is significantly larger for the deliberative group compared to the information only group, this would suggest the marginal productivity of time spent on the background materials was greater for the former. For the question responses considered separately the point estimates for the several of the interaction terms for the treatment group were positive and large, although not statistically significant. We re-estimated the GET model using the sum of correct responses as a single outcome variable, using a binomial distribution with denominator of seven and a probit link function, to test whether effects could be observed for the aggregated responses. Here we find that increasing time spent on the background materials from one standard deviation below the mean to one above the mean increased the probability of answering a question correctly by 21% for those in the deliberative group. There was no effect of time spent for those in the information only group. This suggests the marginal productivity of time with the background materials was substantially higher for the deliberative group compared to the information only group, which is consistent with the proposition that the former group was more motivated to learn the material.

The third mechanism hypothesizes that the prospect of deliberating induces participants to pay closer attention to media reports on immigration (abundant in the summer of 2006) and other extra-study

information sources. Participants with the opportunity to interact with their Member of Congress might have felt a greater motivation to learn about the topic in preparation for the session. To test this we conducted an additional series of regressions using the GET model (results available upon request). We observe statistically significant and substantively large positive treatment effects on following media reporting on immigration (measured on the BGM survey); feeling the duty to stay informed about politics (measured on the follow up survey); following the 2006 elections in the news and knowing which party controlled both the House and the Senate after the elections (all three from the November survey); and for the number of people outside the study with whom respondents discussed immigration (measured separately on the BGM survey and the November survey). While space does not permit developing these supplemental models and reporting the results in full here, the size and direction of the results lend strong support to the idea that those in the deliberative group were more motivated to attend to information about immigration policy (and politics in general) outside the context of the study.

Discussion

In what is perhaps the most prominent analysis of political knowledge in survey research, Delli Carpini and Keeter (1996:10) argue that the knowledge relevant to a citizen's capacity for effective democratic citizenship is the "range of factual information about politics that is stored in long term memory." This stored factual information is "the kind of information that gives citizens... the ability to think and act with greater autonomy and authority." We argue that such a collection of static, stored facts is largely orthogonal to the expectations for democratic citizens embedded in Madisonian democratic theory. Instead, similar to Prior and Lupia (forthcoming) and Zaller (2003), we view the kind of competence democratic citizens need is the capacity to learn about policies when a appropriate occasion arise.

In our experiment, we find that when the average citizen is given the opportunity to interact with their member of Congress – a highly direct occasion for exercising accountability – she has a strong tendency to gain knowledge that is useful to make accountability effective. The ATE estimates in our study are all large and positive. With all due respect to H.L. Mencken, the typical survey measures of static political knowledge underestimate the competence of citizens.

We do observe that those participants with a higher general knowledge of politics do have a better overall grasp of immigration policy (as found in Delli Carpini and Keeter 1993: 1184 and 1996: 270). Delli Carpini and Keeter interpret this shift in the intercepts of the outcome equations as a threat to the fundamental principle of equality among citizens (1996: 271), in that those who sit at the top of society's knowledge pyramid – and who are privileged in so many other ways – also know more about specialized policy topics. In our experiments, however, the capacity to become informed in response to a political event is largely orthogonal to general knowledge (Gilens 200, Jerit et al. 2006) as well as to respondents' motivation to participate in the experiment (as observed in the MTE and ATT parameter estimates). Since this equality is with respect to citizen's capacity to become informed, the equality is evident in the dimension of intelligence that is relevant to Madisonian democracy. We argue that the pessimistic normative conclusions that dominates this literature may follow from an inadequate theory of democracy, and a correspondingly inadequate method for data collection and analysis (Achen 1975), rather than from an inadequate democratic public.

Conclusion

Of course, citizens do not ordinarily get the opportunity to interact with their Member of Congress, or to participate in deliberative experiments. Indeed, to our knowledge, we are the first in political science to involve current members of Congress in field experiments. But our argument suggests that democracy might be able to make do with modest levels of interaction between citizens and their representatives. To a point, opportunity matters more than uptake. In our experiment, we measure the typical citizen's capacity to learn about public policy, and this capacity is most relevant to Madisonian and deliberative representation, both in theory (Pitkin 1967: 222-3, Madison 1961, Habermas 1996) and in practice (Fenno 1978: 231, Kingdon 1989:248, Arnold 1990: 68). In addition, we observe large treatment effects from these online interactions, on a magnitude similar to or perhaps larger than those found in face-to-face deliberative experiments (Fishkin and Luskin 2005). This implies that with the expansion of online technologies, and the increasing penetration of Internet access in society, the emergence of digital government practices such as online town-halls may pave the way for citizens to more regularly interact

with their member of Congress and other publicly accountable actors, should they so choose (Lupia, forthcoming). This would afford more regular opportunities for citizens to exercise their apparent capacity to learn about politics and policies, and more robustly participate in democratic politics.

References.

- Aakvik, Arlid, James J. Heckman, and Edward J. Vytlačil. 2005. "Estimating Treatment Effects for Discrete Outcomes when Responses to Treatment Vary: An Application to Norwegian Vocational Rehabilitation Programs." *Journal of Econometrics* 125 (March): 15-51.
- Achen, Christopher H. 1975. "Mass Political Attitudes and the Survey Response." *The American Political Science Review* 69 (Dec.): 1281-1231.
- Arnold, R. Douglas. 1989. *The Logic of Congressional Action*. New Haven, Conn.: Yale University Press.
- Cameron, A. Colin, and Pravin K. Trivedi. 2005. *Microeconometrics: Methods and Applications*. New York: Cambridge University Press.
- Converse, Philip E. 1964. "The Nature of Belief Systems in Mass Publics." Chapter in *Ideology and Discontent*. David E. Apter, Ed. London: The Free Press of Glencoe.
- Delli Carpini, Michael X. and Scott Keeter. 1993. "Measuring Political Knowledge: Putting First Things First." *American Journal of Political Science* 37 (Nov.): 1179-1206.
- Delli Carpini, Michael X., and Scott Keeter. 1996. *What Americans Know about Politics and Why It Matters*. New Haven, Conn.: Yale University Press.
- Esterling, Kevin M. 2004. *The Political Economy of Policy Expertise: Information and Efficiency in American National Politics*. Ann Arbor: University of Michigan Press.
- Fishkin, James S. and Robert C. Luskin. 2005. "Experimenting with a Democratic Ideal: Deliberative Polling and Public Opinion." *Acta Politica* 40 (Sept.): 284-298.
- Fenno, Richard F., Jr. 1978. *Homestyle: House Members in their Districts*. Boston, Mass.: Little, Brown and Co.

- Frangakis, Constantine E., and Donald B. Rubin. 2002. "Principle Stratification in Causal Inference." *Biometrics* 58 (March): 21-29
- Gastil, John and James P. Dillard. 1999. "Increasing Political Sophistication Through Public Deliberation." *Political Communication* 16 (3): 3-23.
- Gelman, Andrew, Aleks Jakulin, Maria Grazia Pittau, and Yu-Sung Su. 2007. "A Default Prior Distribution for Logistic and Other Regression Models." Typescript, Columbia University Department of Statistics.
- Gilens, Martin. 2001. "Political Ignorance and Collective Policy Preferences." *The American Political Science Review* 95 (June): 379-396.
- Graber, Doris A. 1994. "Why Voters Fail Information Tests: Can the Hurdles be Overcome?" *Political Communication* 11 (4): 331-346.
- Habermas, Jürgen. 1984. *The Theory of Communicative Action, Volume One: Reason and the Rationalization of Society*. Thomas McCarty, trans. Boston, Mass: Beacon Press.
- Hamilton, Alexander, James Madison, and John Jay. 1961. *The Federalist Papers*. Clinton Rossiter, ed. New York, NY: New American Library.
- Horiuchi, Yusaku, Kosuke Imai, and Naoko Tanguchi. 2007. "Designing and Analyzing Randomized Experiments: Application to a Japanese Election Survey Experiment." *American Journal of Political Science* 51 (July): 669-687.
- Hutchings, Vincent L. 2001. "Political Context, Issue Salience, and Selective Attentiveness: Constituent Knowledge of the Clarence Thomas Confirmation Vote." *Journal of Politics* 63 (Aug.): 846-868.
- Imai, Kosuke. 2006. "Statistical Analysis of Randomized Experiments with Nonignorable Missing Outcomes. Princeton University, Department of Political Science, typescript.
- Jerit, Jennifer, Jason Barabas, and Toby Bolsen. 2006. "Citizens, Knowledge, and the Information Environment." *American Journal of Political Science* 50 (April): 266-282.
- Kingdon, John W. 1989. *Congressmen's Voting Decisions*. 3rd Edition. Ann Arbor, Mich.:

University of Michigan Press.

- Kuklinski, James H., Paul J. Quirk, Jennifer Jerit, and Robert F. Rich. 2001. "The Political Environment and Citizen Competence." *American Journal of Political Science* 45 (April): 410-424.
- Lodge, Milton, Kathleen M. McGraw, and Patrick Stroh. 1989. "An Impression-Driven Model of Candidate Evaluation." *American Political Science Review* 83 (June): 399-419.
- Luskin, Robert C. 1987. "Measuring Political Sophistication." *American Journal of Political Science* 31 (Nov.): 856-899.
- Luskin, Robert C. 1990. "Explaining Political Sophistication." *Political Behavior* 12 (Dec.): 331-361.
- Lupia, Arthur. 2006. "How Elitism Undermines the Study of Voter Competence." *Critical Review* 18 (1-3): 217-232.
- Lupia, Arthur. Forthcoming. "Can Online Deliberation Improve Politics? Scientific Foundations for Success." Chapter in *Online Deliberation: Design, Research, and Practice*, Todd Davies and Beth Noveck, Eds.
- Miranda, Alfonso, and Sophia Rabe-Hesketh. 2006. "Maximum Likelihood Estimation of Endogenous Switching and Sample Selection Models for Binary, Count, and Ordinal Variables." *The State Journal* 6:285-308.
- Mondak, Jeffrey J. 2001. "Developing Valid Knowledge Scales." *American Journal of Political Science* 45 (Jan.): 224-238.
- Nadeau, Richard, and Richard J. Niemi. 1995. "Educated Guesses: The Process of Answering Factual Knowledge Questions in Surveys." *Public Opinion Quarterly* (Autumn): 323-346.
- Pitkin, Hanna Fenichel. 1967. *The Concept of Representation*. Berkeley, Calif.: University of California Press.
- Prior, Markus, and Arthur Lupia. Forthcoming. "Money, Time, and the Ability to Answer Political Knowledge Questions." *American Journal of Political Science*.
- Rosenbaum, Paul R., and Donald B. Rubin. 1985. "Constructing a Control Group Using Multivariate

Matched Sampling Methods that Incorporate the Propensity Score.” *The American Statistician* 39 (Feb.): 33-38.

Stokes, Donald E. and Warren E. Miller. “Party Government and the Saliency of Congress.” *Public Opinion Quarterly* 26 (Winter): 531-546.

Thompson, Nicola and Derek R. Bell. 2006. “Articulating Political Knowledge in Deliberation.” *Contemporary Politics* 12 (Sept.-Dec.): 287-300.

Zaller, John. 2003. “A New Standard of News Quality: Burglar Alarms for the Monitorial Citizen.” *Political Communication* 20(2): 109-130.

Table 1. Descriptive Statistics					
	Pretreatment		Endogenous		N
	Mean	SD	Mean	SD	
Outcome Responses (Immigration Policy Knowledge)					
<i>Question 1:</i> About how many illegal immigrants currently reside in the U.S.? [12,000,000] $(\sigma_1^{PT}, \sigma_1)$	0.318	0.466	0.349	0.477	1265
<i>Question 2:</i> About how many illegal immigrants come into the U.S. each year? [500,000] $(\sigma_2^{PT}, \sigma_2)$	0.179	0.383	0.247	0.432	1265
<i>Question 3:</i> About what fraction of illegal immigrants in the U.S. are from Mexico? [1/2 – 2/3] $(\sigma_3^{PT}, \sigma_3)$	0.222	0.416	0.275	0.447	1265
<i>Question 4:</i> Under current law, it is a felony to reside illegally in the United States. [No] $(\sigma_4^{PT}, \sigma_4)$	0.281	0.450	0.402	0.490	1265
<i>Question 5:</i> Under current law, do companies who want to employ non-citizen immigrants have to prove that doing so will not hurt the employment of U.S. citizens? [Yes] $(\sigma_5^{PT}, \sigma_5)$	0.262	0.440	0.325	0.469	1265
<i>Question 6:</i> Under current law, are illegal immigrants who have lived in the U.S. for five years or more eligible to apply for citizenship? [No] $(\sigma_6^{PT}, \sigma_6)$	0.384	0.487	0.481	0.500	1265
<i>Question 7:</i> If an illegal immigrant has a baby while living in the U.S. does that child automatically become a U.S. citizen? [Yes] $(\sigma_7^{PT}, \sigma_7)$	0.839	0.368	0.882	0.322	1265
Assignment Indicators (Z)					
Deliberative Condition (Z _D)	0.529	0.499			2275
Information Only Condition (Z _B)	0.175	0.380			2275
True Control (Z _C)	0.295	0.456			2275
Participation Indicators (P)					
Willing to Do Surveys but not Deliberative (P _A)			0.271	0.445	2275
Participated in Deliberative Group (P _{D Z})			0.364	0.481	1204
Completed BGM Survey (P _{B Z})			0.529	0.499	1603
Completed Follow Up Survey (P _{R Z})			0.712	0.453	1652
Completed November Survey (P _N)			0.702	0.458	1348
General Political Knowledge Indicators (η_2)					
Job Currently Held by Dick Cheney (G ₁)	0.868	0.338			2275
Branch that Determines Law's Constitutionality (G ₂)	0.749	0.433			2275
Majority Required to Override Veto (G ₃)	0.528	0.499			2275
Current Majority Party (G ₄)	0.701	0.458			2275
Which Party More Conservative (G ₅)	0.770	0.421			2275
Exogenous Control Variables (X)					
Knowledge Networks Panelist (X ₁)	0.356	0.479			2275
White (X ₂)	0.810	0.393			2275
Female (X ₃)	0.693	0.461			2275
Full Time Work(X ₄)	0.401	0.490			2275
BGM Exposed Correct Answer for O ₁ (X ₅)	0.500	0.500			848

Length of Time on BGM and BGM Survey (X_6)	0.272	0.140	438
Session Exposed Correct Answer for O_k (average)	1.759	3.561	848
Expert Session (S)	0.087	0.281	2275

Table 2. Estimated Correlations between Participation and Outcome Equations, ρ_{Q_k, P_m} for $k \in \{1 \text{ to } 7\}$ and $m \in \{A, B Z, D Z R Z\}$ (Standard errors in parentheses)				
	Probability that Respondent is Willing/Able to Deliberate [$\lambda_A \hat{=} 0.04$]	Probability that Respondent Completes BGM Survey (if assigned) [$\lambda_{B Z} \hat{=} -1.8^*$]	Probability that Respondent Participates in Deliberation (if assigned) [$\lambda_{D Z} \hat{=} -1.6^*$]	Probability that Respondent Completes the Follow Up Survey (if eligible) [$\lambda_{R Z} \hat{=} -4.0^*$]
P(Question 1 Response is Correct) [$\lambda_1 \hat{=} 1$]	-0.029 (0.050)	-0.615* (0.014)	-0.599* (0.018)	-0.686* (0.004)
P(Question 2 Response is Correct) [$\lambda_2 \hat{=} 0.5^*$]	-0.018 (0.033)	-0.392* (0.093)	-0.382* (0.091)	-0.437* (0.104)
P(Question 3 Response is Correct) [$\lambda_3 \hat{=} 0.2$]	-0.008 (0.018)	-0.172 (0.121)	-0.167 (0.118)	-0.191 (0.135)
P(Question 4 Response is Correct) [$\lambda_4 \hat{=} 0.3$]	-0.010 (0.021)	-0.230 (0.115)	-0.224⁺ (0.112)	-0.256 (0.128)
P(Question 5 Response is Correct) [$\lambda_5 \hat{=} 0.2$]	-0.007 (0.018)	-0.163 (0.120)	-0.159 (0.117)	-0.182 (0.134)
P(Question 6 Response is Correct) [$\lambda_6 \hat{=} 0.6^*$]	-0.020 (0.036)	-0.427* (0.092)	-0.416* (0.090)	-0.477* (0.103)
P(Question 7 Response is Correct) [$\lambda_7 \hat{=} 0.5$]	-0.015 (0.032)	-0.341⁺ (0.166)	-0.332⁺ (0.162)	-0.380⁺ (0.185)
*p<0.05, ⁺ p<0.10 Notes: 1) Correlations between the outcome equations and the probability of responding to the November survey are not shown. 2) Estimated factor coefficients are in square brackets. The symbol $\hat{=}$ indicates an estimate.				

Figure 1a: Average Treatment Effects (Mean Parameters, ΔATE)

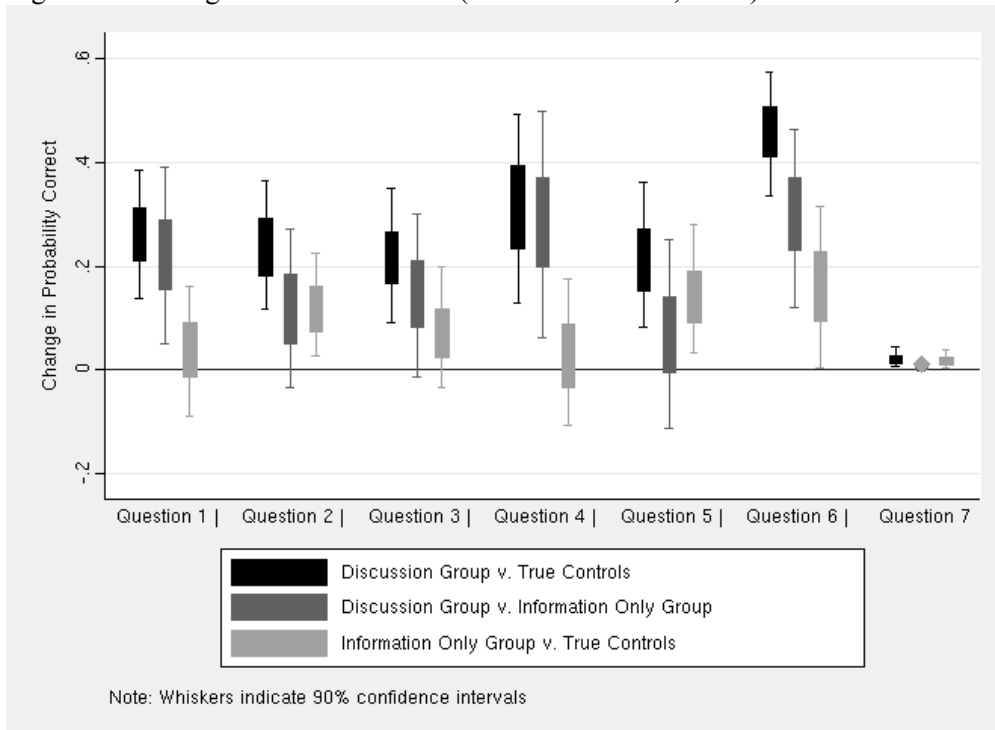


Figure 1b: Average Treatment Effects (Distributional Parameters, $\Delta DATE$)

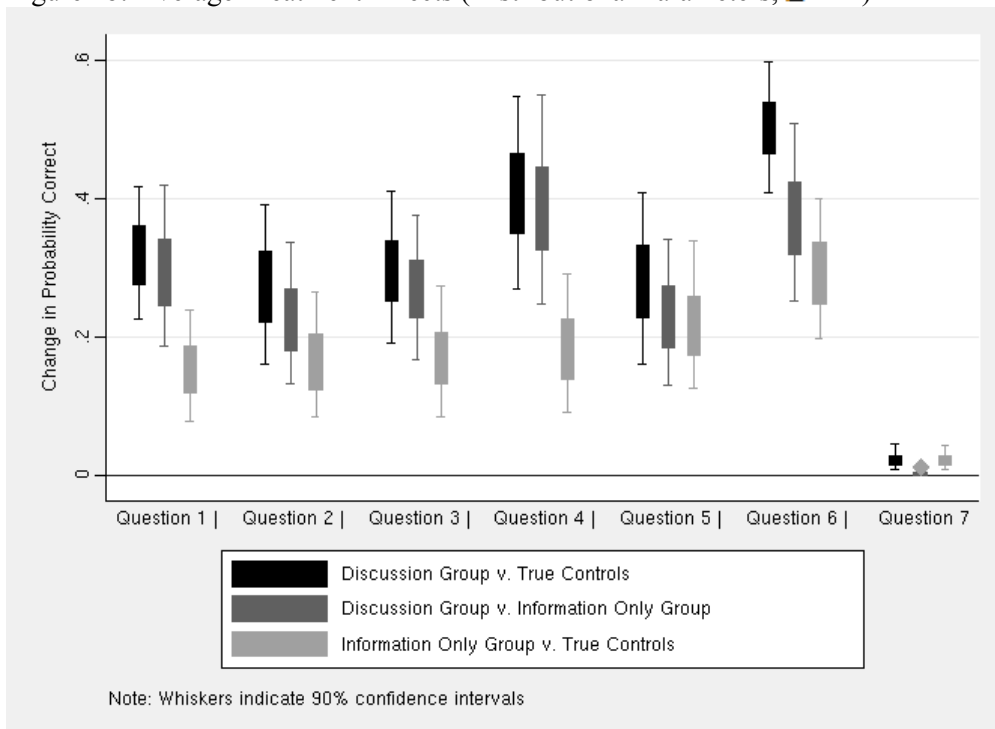


Figure 2a. Marginal Treatment Effects (Δ_{MTE})
Among those Least Likely to Participate

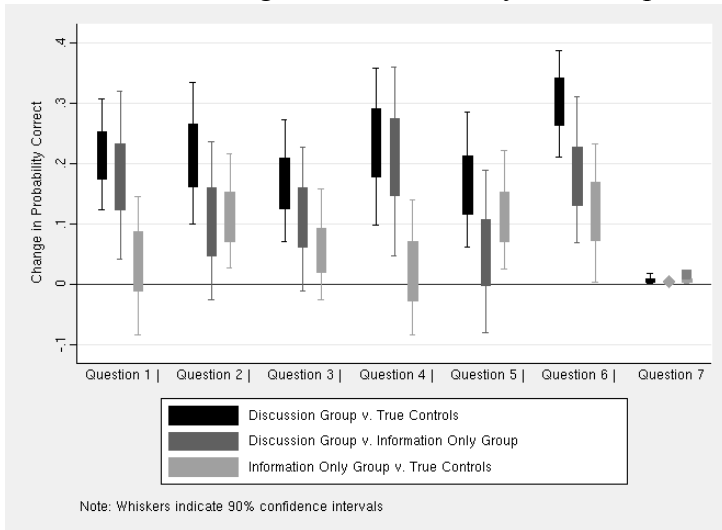


Figure 2b. Marginal Treatment Effects (Δ_{MTE})
Among those on the Cusp of Participating

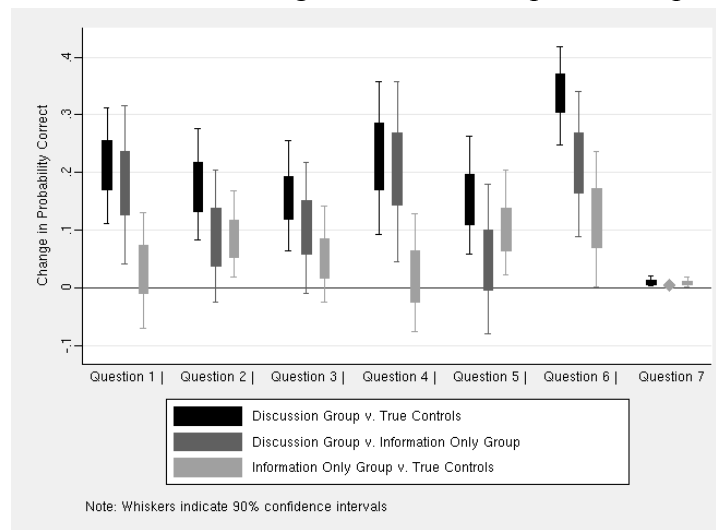


Figure 2c. Marginal Treatment Effects (Δ_{MTE})
Among those Most Likely to Participate

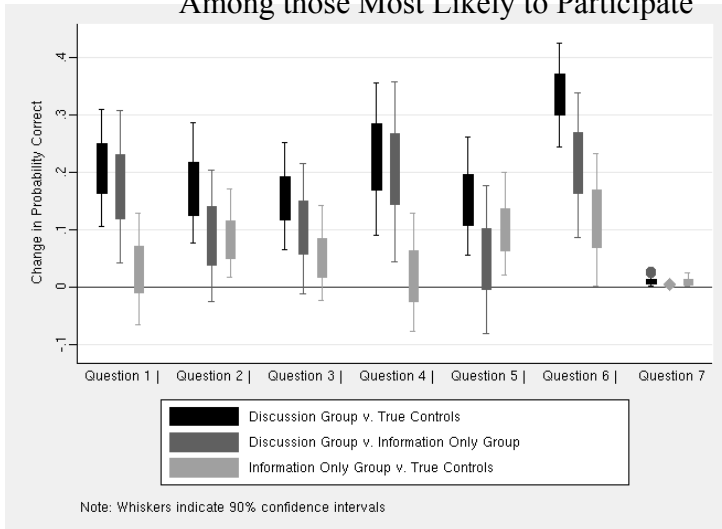


Figure 2d. Average Treatment Effects for the Treated (Δ_{ATT})

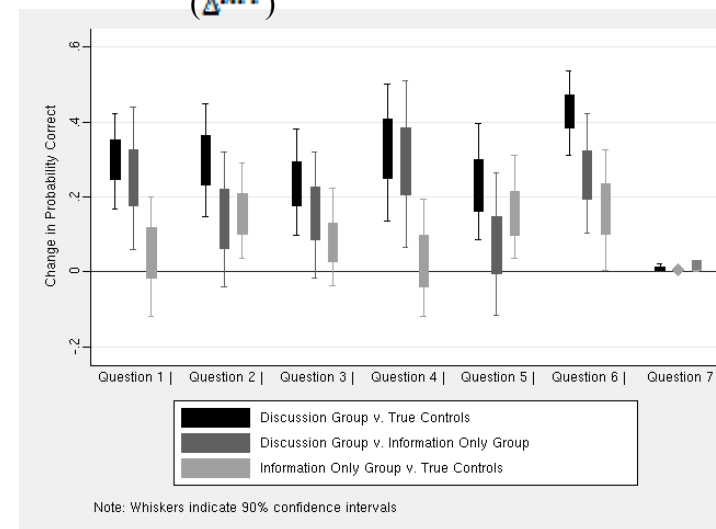
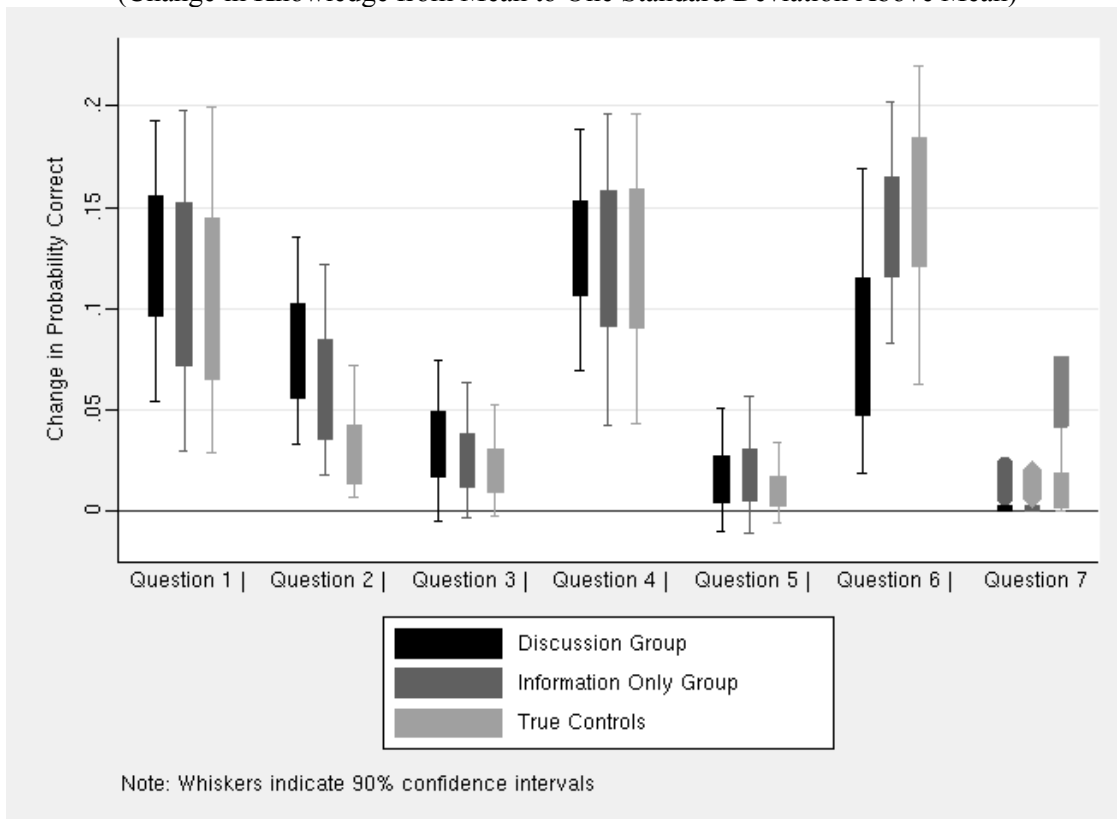


Figure 3: General Knowledge Effects, Conditional on Participation
(Change in Knowledge from Mean to One Standard Deviation Above Mean)



¹ We are using the term “deliberation” in a primarily descriptive, rather than a necessarily approbative sense. We make no claims that our sessions were high quality deliberation; only that subjects got relevant background reading and came together to exchange reasons for and against policy reform. Thompson (forthcoming) argues persuasively for a partial separation between conceptual and evaluative criteria for deliberation. (See also Neblo, forthcoming.) For our larger project we do attempt to code for the deliberative “quality” of each session (Steiner et. al. 2004; Neblo 2000).

² There were two planned sessions with the neutral expert. In addition, one Member backed out of participating too late to cancel the session. The expert substituted in this case as well. However, given the theoretical mechanisms for information gains we discuss below, this case is more appropriately grouped with the sessions involving Members.

³ For a similar argument regarding media coverage, see Zaller 2003.

⁴ The moderator very lightly screened questions if they were patently offensive or vulgar, completely incoherent, or closely duplicated the content of a previous question. Other than duplication, the need to screen was quite rare. In addition, no constituent got to have a second question/comment posted until all others in the queue had theirs posted.

⁵ We also conducted two sessions in March of 2007 where participants interacted with a policy expert rather than a member.

⁶ Due to budget constraints, we did not send the follow up survey to participants assigned to either the information only condition or to the deliberative condition and who did not read the background materials or participate in deliberation.

⁷ Another approach to estimate structural parameters, given the simultaneous problems of non-compliance and non-response is principal stratification (Frangakis and Rubin 2002), which Imai discusses extensively and extends in a series of papers (Imai 2006 and Horiuchi, Imai and Tanguchi 2007). This method identifies the average treatment effect for the treated, but not the average treatment effect.

⁸ Notice that potential participants were informed of the specific times when the session(s) for their district would occur, so presumably some, but not all, potential participants with $P_A=1$ were unable to participate in the session because of exogenous time and scheduling conflicts.

⁹ Note a complication that this randomization was not uniform across sites. Some districts had a relatively small KN panel and so for some (but not all) sites $Z_D=1$ if $P_A=1$

¹⁰ Randomization is not perfect in our application; some sites were fully randomized and some were not. Thus the assumption that the latent indices are the same scale across the treatment arms of our study is somewhat strong. We find that the common latent index assumption is highly robust to using alternative scales for the index (results not reported).

¹¹ AHV (2005) allow the treatment effect to differ across participant types by estimating λ_k as an expanded function of the (endogenous) treatment variables $\lambda_k = \lambda_{1k} + \lambda_{2k}P_{D|Z} + \lambda_{3k}P_{B|Z}$. This in effect estimates both the main effect of participant type on the outcome, and the interaction between the participant type and the treatment she actually received. In addition, this expansion of λ_k identifies the correlation between the individual's treated and untreated conditions. Empirically, we did not find that the expanded factor coefficients different from each other, and so in the estimates below we assume only a common treatment effect.

¹² To retrieve the correlations among the participation variables using this equation, substitute m^i for m for k , and to retrieve the correlations among the outcomes, substitute k^i for k for m .

¹³ Throughout we have used the normal distribution to structure all latent variables and error terms. This assumption is not required and can easily be relaxed. Note that the variances of the outcomes also can be estimated in GET when they are identified. In this application, none of the variances are identified, so the variance equations of the outcome models are suppressed.

¹⁴ We deterministically do not observe the response ($P_{R|Z}=0$ because of the study constraints) for those assigned to the deliberative condition and did not comply on either the BGM or deliberation $\mathbf{E}'=(P_A=1, P_{B|Z}=0, P_{D|Z}=0, P_{R|Z}=0, P_N=n, Z_D=1, Z_B=0, Z_C=0, \mathbf{X}=\mathbf{x}, S=s)$; those assigned to the information only

condition and did not comply on the BGM ($P_A=a$, $P_{B|Z}=0$, $P_{D|Z}=0$, $P_{R|Z}=0$, $P_N=n$, $Z_D=0, Z_B=1$, $Z_C=0$, $\mathbf{X}=\mathbf{x}$, $\mathbf{S}=\mathbf{s}$); and a residual group that was randomly excluded from the study ($\mathbf{Z}=\mathbf{0}$). Because of this, we can only identify comparisons between the cells listed E_C above.

¹⁵ There were five response options plus “Don’t Know,” so the baseline guess rate is 17%.

¹⁶ In the parlance of the statistical experimental literature, these variables indicate the unobserved attribute of whether the participant is a “complier” or “never taker.” In the GET model, this underlying propensity is taken as an unobserved, continuous attribute. Since we restricted respondents’ access to the background materials and the sessions, our sample cannot include either “always takers” or “defiers” for the deliberative condition.

¹⁷ We get very similar results if we simply include an additive general knowledge scale using the five items, or if we dichotomize the additive scale into “high” and “low” knowledge scores.

¹⁸ The estimated change in probabilities assume the participant is average in all respects except we assume he or she is from a KN panel, did not get the question correct on the baseline survey, and if in a deliberative group, was exposed to the correct answer during the session.

¹⁹ This is controlling for whether the respondent gave a correct response on the pretest. Those who gave a correct pretest response were 20% to 60% more likely to get the item correct on the follow up survey, across all questions ($p<0.01$).

²⁰ In addition, KN panelists were much more likely to participate in the study (all effects significant at $p<0.05$): 40% more likely to read the background materials take the survey, 37% more likely to participate in discussions if assigned to the deliberative arm, 42% more likely to respond to the follow up survey, and 87% more likely to respond to the November survey (the last finding likely driven by the eligibility rule and the fact that KN panelists are more likely to take up their assigned treatments).