

The Deliberative Voter*

Robert C. Luskin, University of Texas at Austin

Kyu S. Hahn, University of California at Los Angeles

James S. Fishkin, Stanford University

Shanto Iyengar, Stanford University

It is widely agreed, and not too much of a simplification to consider, that the influences on vote choice form a block-recursive system, roughly a causal chain. More distally, there are, at the individual level, all the sociodemographic variables constituting social location and, at the aggregate level, a variety of political and other environmental factors. More proximately, there are psychological variables, including party identification, policy distances, candidate images, and the like. The sociodemographic and environmental variables affect vote choice largely through the psychological ones. This was what Campbell, Converse, Miller, and Stokes (1960) called the “funnel of causality.” This causal set-up presents a choice: we can focus, more telescopically, on the “reduced form,” explaining vote choice in terms sociodemographics and the environment, or we can focus, closer-up, on the psychology of the vote. Here we choose the latter.

At this psychological level, the variables commonly used to explain vote choice are familiar. The righthand side of the equation typically contains party identification and variables capturing policy, candidate, and retrospective considerations. The findings from such models are also relatively consistent. The greatest influence on vote choice, these days, belongs to candidate considerations. Policy considerations, sadly, are much feebler. Sadly, because it is policy on which voters should, normatively, be basing their decisions. Party and retrospective judgments are only crude approximations, candidate personalities and personal attributes a distraction.

Underlying this pattern is a pervasive and chronic inattention to policy matters (Converse 1964, Luskin 1987, Delli Carpini and Keeter 1996, Kinder xxxx, Price 1999). It is hard to take much account of something one knows and has thought little about. If more voters knew and thought more about policy issues and the candidates’ positions thereon, policy would

presumably loom larger in their decision making. Granted, a number of studies suggest otherwise—that the well informed vote almost as much on the basis of personalities and almost as little on the basis of policies as the poorly informed (xxxx). The perverse inferences drawn from this perverse result are that voting on the basis of candidate personalities must be a sensible approach, since even the well informed do so, and that greater knowledge may not actually make much difference to voters' decision making. (For contrary results and opposing arguments, see Luskin and Globetti xxxx.)

All these analyses rest on cross-sectional survey data. The present study addresses the issue more directly, with a quasi-experiment, specifically an online Deliberative Poll (DP). The manipulation in a DP is designed to get a random sample to talk, learn, and think much more than they would ordinarily do (which is, in most cases, not much) about some electoral or policy choice(s) (for typical results and overviews see, for example, Luskin, Fishkin, and Jowell 2002 and Fishkin and Luskin 2005). Most DPs have been face-to-face, but a number of recent instances, including the present one, have been online (Luskin, Fishkin, and Iyengar 2006; Iyengar, Luskin, and Fishkin 2006). The online deliberations—by voice, not text—lack some of the intimacy, vividness, and force of their face-to-face counterparts but tend to produce similar, if somewhat weaker results (see Luskin, Fishkin, and Iyengar 2006).

The present DP concerned the choice between George W. Bush and John Kerry in the 2004 U.S. presidential election.¹ The discussions and briefing materials centered on some of the policy issues prominent in the campaign and on Bush's and Kerry's positions on them. The results therefore enable us to say something the effects of talking, learning, and thinking more about the candidates, particularly about their policy positions, on the choices voters make and

the criteria on which they rely in making them. We expect deliberation to shift voters' criteria toward policy as opposed to candidate considerations.

Design

Our online DP recruited independent random samples for treatment and control groups, sent the treatment group carefully balanced briefing materials discussing some of the issues prominent in the campaign and describing Bush's and Kerry's positions on them, and had them deliberate online in moderated small groups and ask questions, also online, about the issues and the candidates' positions on them. The participants completed an online questionnaire both when first recruited and at the end of the process.

Sampling

We randomly invited 790 voting-age Americans from the Knowledge Networks (KN) national panel to participate in the online DP. The KN panel, approximately equivalent to a standard RDD sample (Chang and Krosnick 2003; Dennis 2001), only much larger, numbers more than 100,000 Americans between the ages of 16 and 85. The panel members are recruited by RDD sampling, then given web TVs and internet access in exchange for their completing a brief questionnaire each week. For this study, we sweetened the usual KN deal by providing those respondents who lacked one (7.4% of the treatment group) with a free computer, rather than just a web TV, in addition to a two-month ISP trial membership. Those who already had a home computer and internet access were instead offered a financial incentive of \$16 per small group session attended. To fill out the treatment group, we also invited another xxxx KN panel members who were veterans of one or more previous national online DPs. Since a previous study showed little difference in deliberation's effects on veteran versus novice participants

(Luskin, Fishkin, and Iyengar 2006), we pool the two here. By way of control group, we randomly invited another 1390 KN panel members to complete the same pre- and post-deliberation questionnaires, and at the same times, as the treatment group, without participating in any of the deliberations in-between. Everyone, in both the treatment and the control group, was offered a \$6 for completing each survey.

In all, 326 of xxxx invitees initially agreed to participate, of whom 285 completed the pre-deliberation survey, 265 completed the post-deliberation survey, 252 completed both, and 189 completed both and attended at least one small group session. In the control group sample, 333 of 1390 initial invitees completed the pre-deliberation questionnaire, 929 completed the post-deliberation questionnaire, and 275 completed both.

Briefing Materials

Invitees who agreed to participate were sent briefing materials prepared by MacNeil/Lehrer Productions prior to the commencement of the online discussions.² The briefing materials examined the value of preemption in the war on terror, the question of “multi-” versus “unilateralism” in the war in Iraq, the proper balance between “protecting civil liberties” and “increasing homeland security,” whether spreading democracy to other countries should be a priority in the war on terror, the emphasis that should be placed on protectionism (including policies restricting outsourcing) versus free trade, and whether the Bush tax cuts should be made permanent or at least partially reversed. The arguments were carefully balanced, simply sketching what each side would say about each issue, without ever taking sides.

Small Group Discussions

The participants were asked to attend five weekly, hour-long online meetings, extending from September 13 through October 16. Each selected a convenient meeting time, from a menu of fifteen possibilities, and was assigned to a small groups accordingly.³ The groups ranged in size from 10 to 27, averaging 20. Attendance was imperfect, with 89 of 295 people who agreed to participate (roughly 30%) failing to attend any meeting at all. We lump these 98 with the control group. Among those attending at least one session, 18 (6.1%) attended only one, 22 (7.5%) attended two, 24 (8.1%) attended three, 50 (16.9%) attended four, and 95 (32.2%) attended all five. The discussions were voice- rather than text-based. Lotus Sametime software permitted the moderator to award the floor by “passing around the microphone.” Moderators were instructed to encourage everyone to participate and to prevent anyone from hogging the floor. Toward the end of each discussion session, the small groups formulated questions about the candidates’ policy positions.

Questionnaire

Both treatment and control group members completed an online questionnaire, both before and after deliberation (at what we shall call T1 and T2). The post-test questionnaire contained the same questions as the pretest, plus a few more. We asked about vote intention, overall candidate evaluation, the respondent’s and the candidates’ policy positions, candidate traits, and retrospective evaluations of the state of the country, among other things.

The Small Group Questions

The moderators transmitted questions formed in the small groups to the *Online Newshour with Jim Lehrer*, which collaborated with a team at the Center for Deliberative Democracy at Stanford University to answer them, using credible public documents and

sometimes consulting policy experts. The answers to each small group's questions were posted on its website prior to its next meeting.

Knowledge Gains

This, like all Deliberative Polls, was intended to stimulate learning and thought. Thought is difficult to measure, but learning, strongly correlated with thought, can be gauged by examining the participants' performance, before and after deliberation, on knowledge items of two kinds.

One kind concerns factual information about the candidates themselves or related to policy issues distinguishing them. Under this heading we have questions asking: (1) whether Kerry voted for, voted against, or did not vote on the resolution authorizing President Bush to go to war with Iraq; (2) whether India, Pakistan, Sri Lanka, or Indonesia harbored the most Al Qaeda and Taliban fighters; (3) whether it was true that Iraq was directly involved in the attacks on the World Trade Center and the Pentagon on 9-11-2001; (4) whether it was true that large quantities of weapons of mass destruction had been found in Iraq; (5) whether, "on average, prescription drugs cost more in Canada than in the U.S."; (6) whether George W. Bush was drafted but never went to Vietnam, was a decorated officer serving in the Vietnam War, was ineligible to serve in the military because of a medical deferment, or served in the Texas Air National Guard; (7) whether John Kerry was drafted but never went to Vietnam, was a decorated officer serving in the Vietnam War, was ineligible to serve in the military because of a medical deferment, or served in the Massachusetts Air National Guard; (8) whether a major destination for white collar jobs that have gone to other countries is South Africa, Japan, Brazil, or India; and (9) whether the number of Americans killed in Iraq since the war began was closest to 100, 500, 1,000, or 10,000. The correct answers are (1) voted for, (2) Pakistan, (3)

true, (4) true, (5) true, (6) served in the Texas Air National Guard, (7) was a decorated officer serving in the Vietnam War, (8) India, and (9) 1,000. The index is the proportion answered correctly. Cronbach's alpha is .785.⁴

The second kind of knowledge item makes use of the respondents' placements of the candidates on NES-style seven-point policy scales. We have six of these, asking respondents to place themselves, Bush, and Kerry on dimensions stretching (1) from cutting services in exchange for lower taxes to increasing services even at the cost of higher taxes, (2) from intervening militarily on our own to doing so only with international approval and cooperation, (3) from pursuing free trade to protecting U.S. industries, (4) from ensuring constitutional rights to finding every terrorist, (5) from adopting a government plan to cover all medical and hospital expenses for everyone to leaving health insurance in the hands of individuals and private insurance companies, and (6) from legalizing same-sex marriage to amending the constitution to prohibit it. We take Bush to be on the side of lower taxes, unilateral intervention, free trade, finding every terrorist, ensuring homeland security, leaving health insurance in the hands of individuals and private insurance companies, and amending the constitution to prohibit same-sex marriage and Kerry generally to be on the opposite side. On three issues—free trade, constitutional rights, and same-sex marriage—where Kerry was arguably at or very near the midpoint, we also count the midpoint as correct for him. This leads to twelve placement-based knowledge items (two candidates times six issues). Again the index is the proportion answered correctly. Cronbach's alpha is .864.

We can also pool these two sorts of knowledge items into a single index by averaging the factual and placement-based indices. That yields a combined scale with a Cronbach's alpha of .865.

Table 1 shows the percentages answering these questions correctly, both before and after deliberation. Even in the control group, the average percentage of factual items answered correctly increases by roughly 4%, as does the percentage of placement items answered

(Table 1 about here)

correctly. Both increases, though modest, are highly significant. They represent the learning that occurs during and as a result of the real-world campaign. In the treatment group, the corresponding increases, of 8.1% on the factual index and 5.9% on the placement index, are distinctly larger. These increases, too, are highly significant. So, more importantly, are the differences between the increases in the treatment and control groups: deliberation increased learning above and beyond what the campaign would have accomplished. This suggests that whatever changes we see in the participants reactions to the candidates reflect the sort of learning and thinking we were seeking to stimulate.

Vote Intention

Neither this learning nor the thought that presumably accompanied it, however, produced much change, either net or gross, in vote intention. We may set aside the handful of participants intending to vote for Nader or other “third” candidates, who numbered only three at T1 and four at T2. That leaves 188 participants, for whom Table 2 shows the results. Before

(Table 2 about here)

deliberation, Bush and Kerry tied at 43.6% apiece, with 12.8% undecided. After deliberation, Kerry edged ahead, at 47.9% to Bush’s 44.7%, with only 7.4% undecided. Neither the Bush-Kerry T2 difference, nor the T1-T2 change in either’s percentage, however, is statistically significant. Only 4.9% of those initially intending to vote for Bush and only 2.4% of those

initially intending to vote for Kerry (between them, comprising only 3.2% of the sample as a whole) switched sides.

Candidate Evaluation

These last results suggest some very faint shift toward Kerry. But vote intention may be extremely sticky, especially in such an emotional, polarizing campaign as 2004's. Perhaps there was a greater shift in the underlying evaluations of the candidates. To examine this possibility, we compare the ratings of Bush and Kerry on NES-style 0-to-100 "feeling thermometers." We subtract the Kerry rating from the Bush rating, yielding a measure that ranges from -100 (maximally positive about Kerry and maximally negative about Bush) to +100 (the reverse). Table 3 shows the results, along with those for the Bush percentage of the

(Table 3 about here)

two-candidate vote intentions (that is, eliminating the undecideds and the handful of respondents saying they wanted to vote for Nader or someone else). In the feeling thermometer comparison, as in vote intention, the control group shifts a hair, and the treatment group a slightly thicker hair, toward Kerry. Again the treatment-control difference therefore suggests some very slight pro-Kerry effect. But none of this is statistically significant.

Influences on Candidate Evaluation

Despite having moved neither vote intentions nor candidate evaluations very far, deliberation may have shifted the grounds on which the participants evaluated the candidates. The relevant variables include:

Policy. This is simply the difference of the Euclidean distance between the respondent's policy positions and his or her perceptions of Kerry's, on the one hand, and the Euclidean

distance between his or her policy positions and his or perceptions of Bush's, on the other hand. We use the six NES-style seven-point policy scales, concerning taxation, multilateralism, free trade, constitutional rights, health care, and same-sex marriage. Since we subtract the distance from Bush from the distance from Kerry, voters who are closer to Bush receive positive scores, consistent with the scorings of the other variables. Finally, we rescore the difference to run, theoretically, from -1 (perfect agreement with Kerry's perceived positions and maximum disagreement with Bush's) to +1 (the reverse).

Candidate traits. Here we use questions asking how well the phrases "sincere," "intelligent," "strong leader," and "thoughtful" describe each candidate. We average the linearly scored responses, subtract the Kerry score from the Bush score, and then rescale the differenced score to range from -1 (saying that these positive terms all describe Kerry "extremely well" and Bush "not well at all") to +1 (the reverse). Cronbach's Alpha is .928.

PID (Democrat, Republican). Party identification was measured by the stem question of the classic sequence: "Generally speaking, do you think of yourself as a Republican, a Democrat, an Independent, or what?" We create two dummy variables to pick out the Democrats and Republicans, taking independents as the "omitted category."

Retrospective Evaluation. This is based on a single question asking whether "things in general ... in the U.S. over the past year" have been going "very well, pretty well, neither well nor poorly, pretty poorly, or very poorly"? Responses were scaled to range between 0 ("very poorly") and 1 ("very well").

A simple model to explain post-deliberation candidate evaluations is

$$\begin{aligned} \text{EVAL} = & \gamma_0 + \gamma_1\text{TREAT} + \gamma_2\text{TRAIT} + \gamma_3\text{POLICY} + \gamma_4\text{RETRO} + \gamma_5\text{DEM} \\ & + \gamma_6\text{REP} + \gamma_7\text{TRAIT}*\text{TREAT} + \gamma_8\text{POLICY}*\text{TREAT} \end{aligned}$$

$$+ \gamma_9 \text{RETRO} * \text{TREAT} + \gamma_{10} \text{DEM} * \text{TREAT} + \gamma_{11} \text{REP} * \text{TREAT} + u,$$

where EVAL is the Bush thermometer rating minus the Kerry thermometer rating, TREAT is a dummy variable scored 1 for treatment group members and 0 for control group members, TRAIT is the difference between the respondent's mean rating Bush on the personality traits and his or her mean rating of Kerry on the same traits (linearly translated to the -1 to +1 interval), POLICY is the Euclidean distance between the respondent's own positions and his or her perceptions of Kerry's positions minus the Euclidean distance between his or her own positions and his or her perceptions of Bush's positions (also linearly translated to the -1 to +1 interval), DEM and REP are dummy variables scored 1 for Democrats and Republicans, respectively, and 0 for everyone else, the γ 's are the coefficients, and u is the usual disturbance or error term.⁵ Table 4 describes the regressors statistically.

The coefficients for TRAIT, POLICY, RETRO, DEM, and REP ($\gamma_2, \gamma_3, \dots, \gamma_6$) are those variables' effects in the control group. In the treatment group, TRAIT's, POLICY's, RETRO's, DEM's, and REP's effects are $\gamma_2 + \gamma_7, \gamma_3 + \gamma_8, \gamma_4 + \gamma_9, \gamma_5 + \gamma_{10}$, and $\gamma_6 + \gamma_{11}$, respectively. The differences that being in the treatment group makes to TRAIT's, ELECT's, and POLICY's effects are thus $\gamma_7, \dots, \gamma_{11}$. The treatment effect is $\gamma_1 + \gamma_7 \text{TRAIT} + \gamma_8 \text{POLICY} + \gamma_9 \text{RETRO} + \gamma_{10} \text{DEM} + \gamma_{11} \text{REP}$. To the extent that $\gamma_7, \dots, \gamma_{11}$ are nonzero, deliberation's effect depends on the voter's perceptions of the candidates' personalities, policy positions (in relation to his or her own), retrospective judgments, and party affiliation.

As the variables are scored, all the "main effects," with the sole exception of DEM's, should be positive. In the control group, undergoing the same experiences as the electorate as a whole, we expect TRAIT's effect to be largest, and POLICY's to be smallest. But we also expect the treatment to increase POLICY's effect and, less certainly, to reduce TRAIT's and

RETRO's. What should happen to the effects of the party dummies is less clear. Enough knowledge and thought should reduce party's effect, controlling for policy, to zero. A smaller increase, however—enough to make party a more effective heuristic but not enough to obviate the need for such heuristics—might increase it. Thus we expect $\gamma_7, \dots, \gamma_9$ to be positive but have no clear expectation about the signs of γ_{10} and γ_{11} .

The ordinary least squares estimates are presented in Table 5. In the control group, TRAIT does indeed carry far more clout than POLICY. Its coefficient estimate is huge and hugely significant. POLICY's is much smaller and only marginally significant ($p < .10$).

(Table 5 about here)

RETRO also has a modest but statistically significant effect. DEM decreases Bush's edge significantly, by about eight degrees, and REP increases it, not quite significantly, by about three. DEM's effect is noticeably larger in magnitude, but the difference in magnitude is not statistically significant ($p < \text{xxxx}$).

In the treatment group, the effect of policy soars, while that of personality declines. The POLICY coefficient increases from 17.7 to 65.5. The TRAIT coefficient decreases, less dramatically, from 89.8 to 75.9. The effect of being a Democrat (but not of being a Republican) increases a bit, although the increase only borders on statistical significance. The effect of retrospective evaluations scarcely budges.

To provide some perspective on these numbers, Table 5 also reports the "first differences" in the expectation of EVAL. The column headed E_t offers first differences based on each variable's theoretical range—from its minimum possible value to its maximum possible value, i.e. from -1 to +1 for TRAIT, POLICY, and RETRO and from 0 to 1 for REP and DEM. For example, a change in TRAIT from -1 (Kerry rated as positively as possible and Bush as

negatively as possible on all traits) to +1 (the reverse) could be expected to increase EVAL by about 179 degrees in the control group and by about 152 degrees in the treatment group. A change in POLICY from -1 (maximally distant from Bush and at the same point on every issue as Kerry) to +1 (the reverse) could be expected to increase EVAL by 35 degrees in the control group but by over 131 degrees in the treatment group.

Now, the policy effects are not quite as large as they look, because POLICY does not actually occupy its whole theoretical range. RETRO and TRAIT are frequently -1 or +1, but POLICY never reaches either extreme. Its observed minimum is actually only -.667, its observed maximum only .833. The column headed E_e therefore adds the fuzzier but more revealing first differences based on the variables' "effective" or "greatest likely" ranges, arbitrarily but not unreasonably defined as two standard deviations below to two standard deviations above the mean. This measure shows that even the DP participants weigh the candidates' personalities most heavily. Unlike the control group, however, they also weigh policy.

The first difference for the treatment effect, too long an expression to fit in its cell, is $2.90 - 13.9*\text{TRAIT} + 47.8*\text{POLICY} - 2.85*\text{RETRO}$. Its size and sign are thus mostly determined by the size and sign of POLICY. For participants equidistant between Bush and Kerry on policy (POLICY = 0) but revering Bush while reviling Kerry (TRAIT = 1) and thinking things in general have been going neither well nor poorly (RETRO = 0), deliberating decreases the preference for Bush by about 11 degrees. On the other hand, for participants seeing Bush and Kerry as equally appealing (TRAIT = 0) and thinking things in general have been going neither well nor poorly (RETRO = 0) but agreeing decidedly more with Bush on policy (POLICY = say .5, well short of even its observed maximum), deliberating increases

Bush's edge by roughly 27 degrees. Symmetrically, for participants agreeing decidedly more with Kerry than Bush ($POLICY = -.5$), while still finding them equally appealing and thinking things have been going neither well or poorly ($TRAIT = RETRO = 0$), deliberating *decreases* Bush's edge by a similar amount (slightly smaller, given the positive intercept). So the deliberation had an appreciable effect on the preferences of voters who saw themselves as sufficiently closer to Bush than to Kerry, or vice versa, in the policy space.

At the same time, this formula suggests the reason deliberation's effect on the *aggregate* preference for Bush versus Kerry was so modest: *on average*, the voters saw them as roughly equidistant from themselves. The mean (net) estimated treatment effect, across the sample, is $2.90 - 13.9 * \bar{T} + 47.8 * \bar{P} - 2.85 * \bar{R}$, where \bar{T} , \bar{P} , and \bar{R} denote the post-deliberation sample means of TRAIT, POLICY, and RETRO. The mean effect thus depends largely on \bar{P} , which was only .020 (on its -1 to +1 scale).

Discussion

Deliberation, in this Deliberative Poll, did not seem to affect the distribution of votes or even the underlying evaluations of the candidates very much. The only other DP about a general election, in Britain in 1997, saw much larger—indeed very large—effects. Why not here? It may have been because 2004 was such a bitterly fought contest, with so many voters dug in, immovably, on one side or the other. Or it may have been the relative pallor of online deliberation. Or insufficiently lengthy deliberation. Or any of a large number of other elements of the design. Or it could simply be that this was an election in which the actual distribution of votes was not far from the full-information distribution—an election in which much greater and more widespread deliberation would not have made much difference. That is bound to happen sometimes.

The deliberation in this DP did seem, however, to affect voters' decision-making criteria in a normatively desirable direction, increasing the influence of policy and decreasing the influence of personality. This does not appear to be simply a matter of priming. The event's focus on policy could conceivably have shifted the participants' attention to what they already knew and thought about policy issues, even without their learning more or thinking harder. But re-estimating the model of Table 5 separately for respondents high and low on T2 knowledge, produces a POLICY*TREAT coefficient estimate that is highly significant in the high T2 knowledge stratum and not even close to significant in the low T2 knowledge stratum. It is the people who take away the most—who are deliberating the most seriously—for whom the treatment has this conditioning effect.⁶

In many ways, these results echo those from a similar online DP conducted during and about the Democratic presidential nomination process (Iyengar, Luskin, and Fishkin 2006). There too candidate personalities tower over policy considerations in the control group, but policy considerations greatly narrow the gap in the treatment group. But there is also an interesting contrast. Here, in the general election, deliberation not only increases policy's effect but decreases personality's. In the nomination process it only increases the former. No doubt there are differences between primary and general elections that are responsible. The policy differences between candidates tend to be more muted in nomination campaigns, within-party affairs. Voters, even well-informed voters must perforce rely more on candidate considerations.

Table 1
Knowledge Gains

	Control				Treatment				Difference	
	T1	T2	T2 - T1 (s.e.)	<i>p</i>	T1	T2	T2 - T1 (s.e.)	<i>p</i>	Mean (s.e.)	<i>p</i>
Factual Information	.646	.688	.043 (.009)	.000	.749	.830	.081 (.012)	.000	.039 (.015)	.011
Placement Information	.551	.591	.040 (.013)	.000	.646	.705	.059 (.011)	.000	.019 (.017)	.133
Total Information	.591	.634	.042 (.009)	.000	.693	.759	.066 (.009)	.000	.024 (.013)	.030

NOTE: The *p*-values here are one-tailed.

Table 2
Vote Intention, Before and After Deliberation

		T1			
		Bush	Undecided	Kerry	Total
T2	Bush	91.5% (75)	29.2% (7)	2.4% (2)	44.7% (84)
	Undecided	3.7 (3)	33.3 (8)	3.7 (3)	7.4 (14)
	Kerry	4.9 (4)	37.5 (9)	93.9% (77)	47.9 (90)
Total		43.6 (82)	12.8 (24)	43.6 (82)	100 (188)

Entries are percentages (within columns), with raw n 's in parentheses. The bottom row and rightmost column give the T1 and T2 marginals, respectively.

Table 3**Vote Intention and Candidate Evaluation**

	Control			Treatment			Treatment – Control		
	T1	T2	T2-T1	T1	T2	T2-T1	T1	T2	T2-T1
EVAL	7.26	6.52	-.740	4.87	2.27	-2.60	-2.39	-4.26	-1.86
VOTE	.525	.521	-.004	.500	.487	-.013	-.025	-.034	-.009

NOTE: VOTE runs from -1 to +1, EVAL from -100 to +100.

Table 4
Description of Regressors

	Control			Treatment			Treatment – Control		
	T1	T2	T2-T1	T1	T2	T2-T1	T1	T2	T2-T1
TRAIT	.056*	.024	-.034***	.022	-.005	-.018	-.035	-.028	.016
POLICY	-.006	.002	.003	.014	.012	-.004	.020	.017	.001
RETRO	-.080**	-.060*	.019	-.050	-.099**	-.048	-.030	.039	-.066**
DEM									
REP									

*, **

Table 5
Explaining Candidate Evaluation

	Coeff.	s.e.	<i>p</i>	E_t	E_e
Const.	6.53	2.24	.004		
TREAT	2.90	3.48	.405	--	--
REP	3.17	3.46	.180	3.17/2.75	--
DEM	-7.79	3.23	.008	-7.79/-15.5	--
TRAIT	89.8	4.62	.000	179.5/151.8	--
POLICY	17.7	10.9	.052	35.5/131.2	14.5/53.8
RETRO	8.95	3.61	.007	17.8/12.1	--
TREAT*TRAIT	-13.9	7.40	.031	--	--
TREAT*POLICY	47.8	17.7	.004	--	--
TREAT*RETRO	-2.84	5.53	.304	--	--
TREAT*REP	-0.42	5.49	.939	--	--
TREAT*DEM	-7.67	5.13	.136	--	--
Adj. R^2	.903				
<i>n</i>	459				

NOTE: *p* values are one-tailed, excepting those for the coefficient estimates fore TREAT, TREAT*REP, and TREAT*DEM.

*We are grateful to Alice Siu and Jennifer McGrady for research assistance.

¹The briefing materials, discussion, and questionnaire gave little attention to Nader (in keeping with his slender prospects), although the questionnaire did list him in the question asking respondents whom they intended voting for.

²Available upon request.

³This is a departure from face-to-face Deliberative Polling, where the small groups are randomly assigned.

⁴All Cronbach's alphas are computed at T2 (post-deliberation).

⁵The Euclidean distance, for concreteness' sake from Kerry, is $[\sum(K_{ij} - R_{ij})^2/n_{iK}]^{1/2}$, where R_{ij} and K_{ij} are the i th respondent's own position and placement of Kerry on the j th issue, and the summation is over the n_{iK} (≤ 6) issues on which the respondent both takes a position and ascribes one to Kerry. Similarly, the TRAIT variable averages the traits on which the respondent rates each candidate, then subtracts the mean Kerry rating from the mean Bush rating. Respondents who rate either of the candidates on none of the traits, place themselves on none of the issues, or place either of the candidates on none of the issues are treated as missing data.

⁶We use T2 knowledge rather than the T2 - T1 knowledge gain because those who already know a lot learn the most but cannot show it. By contrast, those who emerge knowing a lot have generally learned a lot, either observably, if they initially knew only a little, or unobservably, if they already knew a lot. Indeed, observed T2 knowledge can be shown, under very comfortable assumptions, to be more highly correlated with true knowledge gain than is

observed knowledge gain (see Luskin, Fishkin, and Jowell 2002). High and low knowledge are defined as $K > \bar{K}$ and $K \leq \bar{K}$, where K is the average of the proportions of factual and placement based knowledge items answered correctly and \bar{K} is the sample mean of K (.72). The estimated coefficients for POLICY*TREAT are 27.8 and 87.1 for the low and high knowledge strata respectively. The estimated standard errors are 21.2 and 29.8, and the p -values are .096 and .004.