How Does Treatment Self-Selection Affect Inferences About Political Communication?

Thomas J. Leeper*
Department of Government
London School of Economics and Political Science
thosjleeper@gmail.com

December 24, 2015

Ecological validity is often seen as vital to experimental research because experiments that provide too stylized or artificial an experience for participants may not speak to any real world phenomenon of interest. One concern related to ecological validity is treatment self-selection: if individuals in the real world self-select their own treatments, as they often do with political communications, how well does the experimentally identified sample average treatment effect estimate the effect of message exposure for those individuals who would — if given the choice — opt into or out of receiving that treatment? This study shows that treatment randomization masks effect heterogeneity between individuals that would have selected different messages if given the option. Yet such treatment selections are themselves complicated, revealing further challenges for realistically studying treatments prone to self-selection. The results have both methodological and substantive implications: evidence of effect heterogeneity raises questions about the appropriateness of random assignment experiments for studying political communication, and the results advance our understanding of citizens' selection into and responses to communications when, as they often do, have choice over what messages to receive.

*An earlier version of this paper was presented at the Midwest Political Science Association 2012 Annual Meeting, Chicago, IL and the International Society for Political Psychology 2012 Annual Meeting, Chicago, IL. Thanks are due to John Brehm, Jamie Druckman, Samara Klar, Jeff Mondak, Marco Steenbergen, and the editorial team at JEPS for feedback on various stages of the project, and to Jamie, Fay Lomax Cook, and Toby Bolsen for sharing survey data from the $t1$ panel wave.

Experiments are seen as a gold standard because randomized assignment to a treatment provides unparalleled internal validity, ensuring that differences in outcomes between groups are due to the treatment alone. While experiments are sometimes criticized for lacking external validity (see McDermott 2011; Shadish, Cook, and Campbell 2001), another important but less studied concern is the notion of ecological validity: i.e., the "realism" of the experiment. Ecological validity is vital to experimental design because studies that involve too stylized an experience for participants may not speak to any real world political phenomenon of interest.

One recent concern about ecological validity relates to the notion of treatment self-selection (Lau and Redlawsk 2006; Gaines and Kuklinski 2011*b*,*a*; Druckman, Fein, and Leeper 2012; Arceneaux and Johnson 2012). If in the real world individuals self-select their own "treatments" then randomized exposure may not constitute a realistic exploration of that phenomenon. Given this natural self-selection, how well does the experimentally identified treatment effect estimate the effect of message exposure for those individuals who would — if given the choice — opt into or out of receiving it? This is particularly important in the study of political communication, where real-world exposure to messages is often self-selected.

To answer this, the present research uses a population-based survey experiment to test how individuals respond to policy arguments are either randomly assigned or self-selected by participants. The findings show that treatment randomization masks effect heterogeneity across individuals inclined to select alternative messages. Furthermore, a pretreatment manipulation of attitude importance in the experiment modified participants' choice of treatments and led to variations in apparent treatment effects, suggesting that the inferences from choice-based experiments hinge on which respondents self-select each treatment. The results speak to research on selective exposure, motivated reasoning, and political communication, and more generally to the use of experiments for studying treatments characterized by real-world self-selection.

Studying the Effects of Self-Selected Political Communications

Arguments conveyed by media and political elites are widely seen as an important source of political information that citizens can use when forming preferences about politics (Disch 2011; Chong and Druckman 2007). Much of the evidence for the influence of elite communications comes from randomized experiments that expose participants to different messages and then measure information effects on outcomes such argument evaluations, opinions, issue importance, and information-seeking (see, for example, Ansolabehere et al. 1994; Arceneaux and Johnson 2012; Berinsky and Kinder 2006; Brewer 2003; Brewer and Gross 2005; Iyengar and Kinder 1987; Levendusky 2013; Miller and Krosnick 2000; Nelson, Clawson, and Oxley 1997; Petty and Cacioppo 1986; Tichenor, Donohue, and Olien 1970). Questions have been raised, however, about the extent to which such experiments provide an adequate empirical test of real-world information exposure given that, in Hovland's words: "In an experiment the audience on whom the effects are being evaluated is one which is fully exposed to the communication. On the other hand, in naturalistic situations with which surveys are typically concerned, the outstanding phenomenon is the limitation of the audience to those who expose themselves to the communication" (Hovland 1959, 9; Bennett and Iyengar 2008, 724). For example, Albertson and Lawrence (2009) use a field experimental design to study the effects of watching a television news program, but had to address the fact that some people encouraged to watch the program chose not to watch it and some who were not encouraged watched it anyway. Their solution is to use the randomized encouragement as an instrument for viewing the program and identify the effect of the program on compliers (i.e., the effect for the unobservable subset of individuals who only watch the program if compelled to do so). In cases such as this, it is unclear what the effect of the treatment would have been had it been applied to everyone (a consequence of encouragement non-compliance), nor the effect for those who chose to view the program or the effect for those who chose not to view

the program (in both cases because the CATE is an effect for those whose treatment uptake can be manipulated). In the face of practical challenges such as these, much experimental research continues to focus on the sample average treatment effect, which averages the effects for the two subgroups that choose to be treated and choose to be untreated, respectively (but see Levendusky 2012; Arceneaux and Johnson 2012).

Aside from the sample average treatment effect, Hovland encourages researchers to focus on these separate effects for "those who expose themselves to the communication" and those who do not. Why would we care about the treatment effect for these two groups? Consider the classic experiment by Nelson, Clawson, and Oxley (1997) in which laboratory participants were randomly assigned to watch a news story framing a hate rally in terms of either free speech or public order. While our interest may lie in knowing the effect of a hypothetical intervention in which all citizens receive the free speech news story, our interest in understanding citizens' interactions with media suggests that we also (and perhaps rather) want to know how that message affects different segments of the public. How are those who opt-in to free speech framed news affected by it? How are those who would rather opt-in to public order news affected by free speech news? The answers to these questions matter because much exposure is selective in this way. Indeed, despite early skepticism (Sears and Freedman 1967), it is now a relatively uncontroversial claim that citizens selectively expose themselves to political information (Bennett and Iyengar 2008; Bolsen and Leeper 2013; Stroud 2011; Kim 2007, 2009; Iyengar and Hahn 2009; Iyengar et al. 2008; Garrett 2009*b*,*a*; Garrett, Carnahan, and Lynch 2013; Feldman et al. 2011; Smith, Fabrigar, and Norris 2008) and a recent meta-analysis suggests that *political* selective exposure is especially prevalent (Hart et al. 2009). While individuals select messages based on prior attitudes, they also appear to engage in selective exposure according to ideology, habit, and topical interests (Bennett and Iyengar 2008; Prior 2007; Iyengar and Hahn 2009; Baum 2002). Given this selectivity, we should aim to understand how individuals are affected by treatments they actually encounter in the real world.

In what ways might we expect the effect of a political communication to differ across individuals inclined and disinclined to select it? There are two general possibilities. One is *homogeneous effects*: the effects of exposure to a political communication are the same for those who would choose to receive the message as for those who would choose an alternative message. In the case of such homogeneity, the SATE reflects the effects for both groups and there are no concerns about what can be learned from a randomized experiment.[1] The other possibility is *heterogeneous effects*: the effects of exposure to a political communication are different for those who would choose and those who would not choose a given message.

Why might effect heterogeneity exist? One possibility is suggested by research on motivated reasoning (Kunda 1990; Lodge and Taber 2000; Nickerson 1998; Taber and Lodge 2006). If an individual selects a message that reinforces their prior views (e.g., an individual supportive of a policy chooses a supportive argument), then exposure to a preferred message is likely to either have no effect due to ceiling constraints or to make that individual's opinion more extreme. By contrast, an individual for whom that message is dispreferred is likely to counterargue against this dispreffered message, such that exposure has no effect due to their motivated resistance to countervailing arguments or even has a backfire effect by which the individual becomes more negative in responses to a supportive message. Thus, homogeneous effects occur if messages are uniformly persuasive (which is unlikely given past evidence for polarization effects; Taber and Lodge 2006) or if a message has no effect on anyone. Heterogeneous effects across these groups seem more likely given that null effects for one group paired with either positive or negative effects for the other group produce a pattern of heterogeneity. Some past research suggests that individuals are likely to see attitude-congruent arguments as stronger and more effective than incongruent arguments (Ditto et al. 1998), which means that the effects of political communications on opinions for message selectors and non-selectors are likely point in opposite directions. It is therefore

[1] Another possibility is that even in the face of effect heterogeneity across individuals, message self-selection is as-if random so the groups of message self-selectors and non-selectors are indistinguishable from on another on average.

reasonable to expect effect heterogeneity, which means the sample average treatment effect may not be representative of the effects of exposure to political communications for either those inclined to choose the treatment or the remainder of the population.

Given the theoretical grounding in motivated reasoning, it is also worth considering whether this expected pattern of effect heterogeneity holds when individuals are more or less inclined to engage in attitude-congruent selective exposure. Attitude importance is one mechanism thought to affect the degree of attitude-congruent selective exposure (Holbrook et al. 2005; Taber and Lodge 2006; Leeper 2014), so when individuals believe an issue to be very important they are more likely to choose congenial messages. By contrast, under low importance, individuals are thought to make more balance information choices. High attitude importance should also increase information-seeking in general, as has been demonstrated in prior work (Krosnick 1990; Holbrook et al. 2005). Therefore when issue opinions are felt to be personally important, it is reasonable to expect greater effect heterogeneity. When issue opinions are felt to be less important, the increased likely of attitude-incongruent selection means that the audiences for alternative political communications are more similar and the effects of exposure to a given message should be, on average, more similar.

Experimental Design

I use a hybrid experimental design — combining randomized exposure and message self-selection — to test for this expected pattern of effect heterogeneity across those selecting into one of two alternative arguments about a policy issue. In the hybrid design, one-half of participants participate in a randomized experiment and the other half participate in an observational study involving treatment self-selection (Gaines and Kuklinski 2011*a*). The present research builds on past research by employing a hybrid design with three additional features: (1) a pretreatment manipulation of attitude importance that is meant to modify the degree to which respondents choose treatment messages consistent with their prior attitudes,

(2) a two-wave panel design that allows for a clean measure of prior opinions, which are a theorized cause of message selection, and (3) a nationally representative sample of adult participants.

I test for heterogeneous effects between message selectors and non-selectors in response to randomly assigned or self-selected information on the issue of "renewable energy portfolio" standards, which require electrical utilities to produce energy from renewable resources.[2] The survey took place in two-parts, with the first wave of data was collected in Summer 2010 (hereafter, $t1$) and involved collection of demographic covariates and baseline attitudes. The second, experimental wave was collected in Spring 2011 (hereafter, $t2$). The two-wave design is advantageous because it provides a clean measure of $t1$ opinion, avoiding accessibility or consistency biases into respondents' behavior during the $t2$ survey experiment and enables estimation of opinion effects using more precise within-subjects, pre-/post-treatment changes.[3]

Manipulations

The $t2$ experiment involved three manipulations: a manipulation of attitude importance, the direction of the energy policy argument, and whether that argument was self-selected by respondents or captively (i.e., randomly) received. Figure 1 provides a visual summary of the design, along with treatment group sample sizes and notation used in defining causal effects. The first manipulation involves modifying the apparent personal impact of the energy proposal. This serves as an instrument for respondents' information choices (which would otherwise be nonrandomly determined) and provides insight into the extent to which self-

[2]U.S. Department of Energy, 2003, "Analysis of a 10-percent Renewable Portfolio Standard," Office of Integrated Analysis and Forecasting of the Energy Information Administration; Chen, Cliff, Ryan Wiser, and Mark Bolinger, 2007, "Weighing the Costs and Benefits of State Renewables Portfolio Standards: A Comparative Analysis of State-Level Policy Impact Projections," Report prepared for the Permitting, Siting, and Analysis Division, Office of Electricity Delivery and Energy Reliability, U.S. Department of Energy.

[3]Given that the issue received almost no media coverage in the intervening months, there is no reason to believe that any subjects were substantially influenced between the earlier and present wave (an assumption that is testable by comparing present attitudes of those in the control condition against their responses in the earlier wave). On average, opinions became slightly more negative between the two waves.

Figure 1: Experimental Design and Treatment Group Sample Sizes

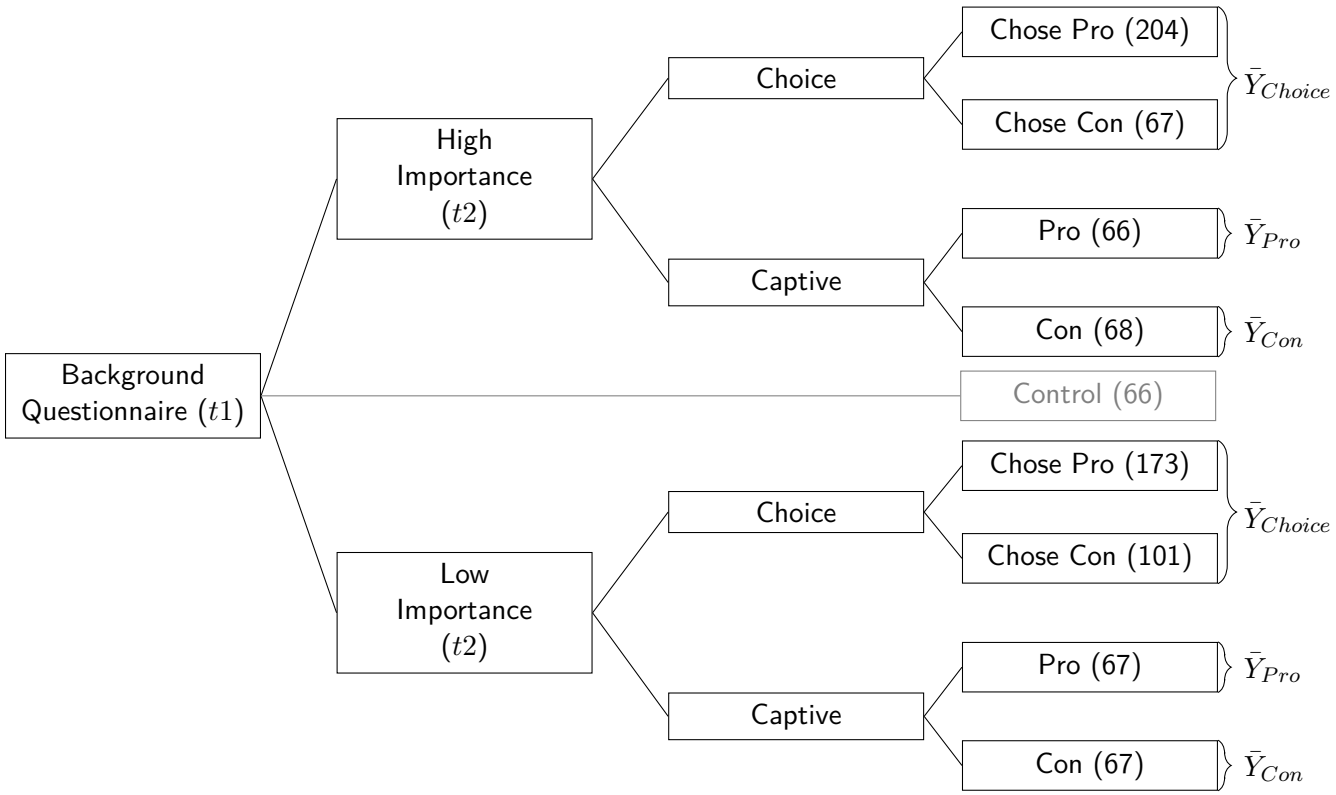

Differences in sample sizes within the captive conditions reflect random assignment. Differences in sample sizes between "Chose Pro" and "Chose Con" conditions reflect treatment self-selection. Choice conditions were intentionally oversampled. The $\bar{Y}$ values at right are meant to clarify what groups are used for estimating treatment effects.

selection behavior varies across contexts. The logic of the manipulation was that individuals who believe their self-interest is at-stake would display higher attitude importance and thus be more likely to choose a treatment message (if given the choice) that was congruent with their prior attitudes (Krosnick 1989; Boninger, Krosnick, and Berent 1995). Importance was manipulated to be high by telling respondents:

> A new law is currently moving through Congress that would require your electricity provider to purchase energy from renewable sources (e.g. wind and solar). This is relevant to you since it will influence your energy bills and the environment. The law would go into effect immediately.

Those in low importance condition read that:

> Some have proposed a bill that would require electricity providers to purchase energy from renewable sources (e.g. wind and solar). This is probably not directly relevant to you because Congress does not appear to be ready to act on the bill and even if they did it is unlikely to personally affect you.

A post-treatment manipulation check asked "How important to you personally is your opinion about this renewable energy restriction?" The results of the manipulation check confirm that attitude importance was manipulated. On a 0-1 scale, those in high importance conditions rated their attitude importance at 0.77 and those in low importance conditions rated their attitude importance at 0.69, a statistically significant difference ($p = 0.00$). The mean attitude importance reported by a control condition that received no manipulation and no information was 0.72. There was also no indication that the manipulation of importance affected information evaluations ($p = 0.15$), post-treatment opinions ($p = 0.42$), or within-subject opinion changes ($p = 0.34$).[4]

The second manipulation varied whether the information respondents read was supportive

[4]Three rounds of pilot tests were conducted with 80 participants each in order to determine the question and manipulation wordings. Participants were recruited from Amazon Mechanical Turk and compensated \$.20 each for their participation. In the final pretest, the high importance manipulation achieved a mean importance (on a 7-point scale) of 5.66 (SD=1.08), while the low importance manipulation achieved a mean importance of 5.2 (SD=1.44). This effect is small, but in the intended direction. It is also a larger effect than either of the previous pretests. The difficulty of manipulating attitude importance (Visser, Holbrook, and Krosnick 2007) makes the relatively small apparent effects in the small-n pretest seem reasonable.

(Pro) or opposed (Con) to the policy. The Pro message was entitled "Renewable Energy Rules Beneficial" and the Con message was entitled "Renewable Energy Rules Ineffective." Exact text of the messages is available in Appendix A. An effort was made to ensure that the informational content of the Pro and Con messages was near-identical. The difference between the two treatments is in the language chosen to describe the same basic facts. Pilot testing confirmed that messages were equally effective. The Pro message was rated as 5.08 (SD=1.38) on a 7-point effectiveness scale, while the Con message was rated 4.97 (SD=1.47) on the same scale. And, the Pro message produced attitudes (on a 7-point scale) of 5.63 (SD=1.64) and the Con message produced attitudes at 4.5 (SD=1.93), a difference that is in the intended direction.

The final manipulation involved how the informational treatments were assigned to respondents. One-third of the respondents were assigned to "captive exposure" conditions that involved simply reading either the Pro or Con argument (based on random assignment). The other two-thirds were presented with the headlines for each passage and told to choose one of them to read, which they were then given, directly emulating the process of selective exposure.[5]

Measures

After receiving the importance and informational manipulations, outcome measures were collected including evaluations of the information received, attitude toward the policy, subjective intentions to obtain further information, and a behavioral measure directly tapping willingness to receive additional information in the form of an email message.[6] All variables are coded to range from 0 to 1, with higher scores indicating more positive evaluations, higher policy support, or greater intention to seek information. The opinion question read, "Thinking about energy related restrictions, to what extent do you oppose or support requir-

[5]A control group exposed to no treatment was also included but is not analyzed here.

[6]While the primary interest is in evaluations and opinions, the information-seeking measures are thought to be affected by heightened attitude importance and may be further shaped by exposure to issue arguments.

ing electricity providers to purchase energy generated from renewable sources (e.g., wind, solar)?" and solicited responses on a seven-point scale from "strongly opposed" to "strongly support." As already mentioned, this item was also measured on the $t1$ survey, enabling estimation of treatment effects for opinion using both post-treatment and pre/post changes.

The argument evaluation question read, "How effective would you say the information you read was in making an argument about this energy-related restriction?" The subjective information-seeking question asked "How likely are you to seek more information about renewable energy requirements?" The behavioral measure asked "Can we send you an email with more information about renewable energy requirements?" Responses to the latter measure were coded as 1 if the respondent entered their email address and 0 otherwise. The two information-seeking measures correlate to some extent ($r = 0.44$). Treatment group means for all outcome measures are reported in Appendix B.

Estimation Strategy

Following from Gaines and Kuklinski (2011*a*), I estimate three different treatment effects for each outcome variable. The first is the familiar sample average treatment effect (SATE):

$$SATE = \bar{Y}_{Pro} - \bar{Y}_{Con} \tag{1}$$

where $\bar{Y}_{Pro}$ is the mean outcome value among those captively assigned to the Pro message and $\bar{Y}_{Con}$ is the mean outcome value among those captively assigned to the Con message.

As Gaines and Kuklinski point out, this SATE is a weighted average of effects of the Pro (versus Con) message for different observable subsamples: one for those who would choose the Pro message and one for those who would choose the Con message *if given the choice*.[7] The second effect is therefore the effect of the treatment on those who would choose it (the

[7]Because respondents in choice conditions were only allowed a choice between Pro and Con messages, it is necessary to focus on the difference between receiving these two messages as opposed to a treatment-control comparison.

Pro-Selector Effect, hereafter the PSE):

$$PSE = \frac{\bar{Y}_{Choice} - \bar{Y}_{Con}}{\hat{\alpha}} \tag{2}$$

where $\bar{Y}_{Choice}$ is the mean outcome value among all respondents assigned to the "choice" condition (see Figure 1) and $\hat{\alpha}$ is the proportion of these individuals choosing the Pro message.

The third effect captures the effect of receiving the Pro rather than Con message among those who would not choose the Pro message (i.e., the *Con-Selector Effect*, hereafter the CSE):

$$CSE = \frac{\bar{Y}_{Pro} - \bar{Y}_{Choice}}{1 - \hat{\alpha}} \tag{3}$$

In other words, the PSE represents the effect of the Pro message treatment versus the Con message for those who would opt for the Pro message *if given the opportunity* and the CSE represents the effect of the Pro message for those who would opt for the Con message *if given the opportunity*. The PSE and CSE are thus two average effects for different subsets of the population. These effects are causally identified given random assignment to the choice and captive arms of the experiment, which makes individuals in each arm identical in expectation. A concern with this strategy would be if being given a treatment choice per se (regardless of treatment chosen) affects the outcome (an exclusion restriction), but is not something that is easy to test empirically.

Effect homogeneity occurs when the PSE and CSE are similar, and are in turn similar to the SATE. Effect heterogeneity occurs when these quantities diverge. If homogeneous, there is little reason to be concerned about the ecological validity of a captive exposure experiment or the inferences that can be drawn from that experiment to different segments of the population. Without such heterogeneity, the SATE is a reasonable estimate of treatment effects for Pro selectors and Con selectors.

Respondents

Data were collected by Bovitz Research Group of Encino, CA, who provide an online panel of approximately one million respondents recruited through random digit dialing and empanelment of those with internet access. As with most internet survey samples, respondents participate in multiple surveys over time and receive compensation for their participation. A total of 879 respondents completed both waves and analysis is restricted to these respondents. A total of 885 respondents completed the first wave, suggesting there is little concern about sample attrition. The sample was drawn to demographically match the U.S. adult population and data are analyzed without weighting. Respondents had a median age of 49, and were 49.0% female, 75.0% white, and 98.9% had at least high school degrees and 80.9% had university degrees. The partisan composition was 39.2% Democrats and 30.7% Republicans.

Results

I begin by examining the information choices made by respondents in the choice conditions. To start, it is clear that the manipulation of attitude importance was effective: 88% of high importance respondents chose information congruent with their $t1$ opinion, while only 63% of low importance respondents chose information congruent with their $t1$ opinion (a statistically significant difference).[8] This means the odds of a high importance respondent choosing the message congruent with their $t1$ opinion was roughly 7:1, while the odds for a low importance subject were only about 1.5:1, yielding a odds-ratio of 4.25, which is a dramatic and significant effect.[9] These variations in attitude-congruent selective exposure are consis-

[8]Recall that 50% choosing congruent information would be considered "random" choice.

[9]The odds-ratio would be equal to one if there was no effect of importance. In this case, the 95% confidence interval for the odds-ratio is (2.68,6.75), suggesting a large effect of attitude importance on attitude-congruent selective exposure. This result also holds when looking at supporters and opponents separately: among those with high importance attitudes, 88% of supporters chose the Pro information and 85% of opponents chose the Con information. By contrast, under low importance, only 65% of supporters chose the Pro information and a mere 35% of opponents chose the Con information. Given that all respondents were presented with the

tent with expectations. As such, it is reasonable to expect that if there are heterogeneous effects of exposure to a Pro message among selectors and non-selectors, the pattern of that heterogeneity is likely to differ across individuals in the high and low importance groups.

Table 1 presents the main results, separately for the full sample in panel (a), the low importance condition in panel (b), and the high importance condition in panel (c). Because importance affected the degree of attitude-congruent selective exposure, separately estimating the effects in this way makes it possible to detect further effect heterogeneity when the pattern of treatment self-selection changes. This possible source of variation in selectivity has important methodological consequences because past work has implicitly assumed that selectors are *fixed types* who would always choose the same way. That assumption has not previously been tested.

Looking at panel (a), we see sample average treatment effects (SATEs) estimated from the captive conditions for each of the five outcome measures (argument evaluation, opinion level, opinion change, planned information-seeking, and requests for an email with issue-relevant information), alongside the estimated PSEs and CSEs. In substantive terms, these five SATEs indicate that, on average across the whole sample: (1) the Pro argument is seen as more effective than the Con argument, which makes sense given the supportive attitudes of the sample, (2) exposure to the Pro message increases support for the policy, (3) this increase in support holds when measured by $t2 - t1$ changes in opinion rather than $t2$ levels of opinion, (4) the Pro message insignificantly reduces self-reported plans to seek out issue-relevant information, and (5) the Pro message reduces the likelihood of requesting an email with more policy information.

two headlines in the same order — the Pro headline coming first — it is possible that the slight preference for the Pro information among those with low importance is attributable to primacy effects. It is also possible that low importance respondents might have opted to read nothing if they had been given the choice, meaning that lacking importance people disengage from political information entirely. Furthermore, given that respondents' opinions leaned positive ($\bar{x}_{t1}$=0.76, SD=0.01), breaking out the results in this way confirms that the aggregate rates of attitude-congruent selective exposure are not simply a result of the sample disproportionately choosing the Pro message. As a comparison, of those with neutral $t1$ opinions, 49% chose the Con message and 51% chose the Pro message, suggesting their choices were relatively arbitrary and confirming the pretest finding that neither the Pro headline or Con headline was more enticing to read.

Table 1: Treatment Effect Estimates, by Importance Condition

	SATE	PSE	CSE
Evaluation	0.14 (0.03)	0.21 (0.04)	-0.01 (0.08)
Opinion	0.08 (0.04)	0.12 (0.04)	0.01 (0.09)
Opinion Change ($t2 - t1$)	0.11 (0.03)	0.12 (0.04)	0.09 (0.08)
Information-Seeking	-0.02 (0.03)	-0.03 (0.04)	0.02 (0.08)
Email Request	-0.13 (0.06)	-0.20 (0.07)	0.04 (0.15)

(a) Full sample

	SATE	PSE	CSE
Evaluation	0.14 (0.02)	0.09 (0.05)	0.11 (0.09)
Opinion	0.14 (0.03)	0.04 (0.06)	0.07 (0.11)
Opinion Change ($t2 - t1$)	0.08 (0.03)	0.06 (0.06)	0.20 (0.10)
Information-Seeking	0.03 (0.03)	-0.05 (0.06)	0.05 (0.10)
Email Request	-0.12 (0.05)	-0.32 (0.10)	-0.17 (0.17)

(b) Low Importance Condition

	SATE	PSE	CSE
Evaluation	0.21 (0.03)	0.31 (0.05)	-0.20 (0.14)
Opinion	0.31 (0.03)	0.18 (0.06)	-0.08 (0.17)
Opinion Change ($t2 - t1$)	0.08 (0.03)	0.17 (0.05)	-0.07 (0.13)
Information-Seeking	0.05 (0.03)	-0.01 (0.04)	-0.03 (0.13)
Email Request	0.05 (0.05)	-0.09 (0.09)	0.36 (0.28)

(c) High Importance Condition

Note: Cell entries are estimated treatment effects, with bootstrapped standard errors in parentheses. The three estimated values of $\hat{\alpha}$ used in estimating the PSEs and CSEs are: 0.69 (full sample), 0.63 (low importance), and 0.75 (high importance).

The PSEs paint a largely similar story, with differences only in the magnitude but not direction of effects. The story is different for CSEs: while the Pro message still increases policy support for these individuals (as measured by changes over time), there are no other substantively sizable effects on exposure and none of the CSEs are statistically distinguishable from zero. In other words, there appears to be effect heterogeneity: those inclined to select the Pro message are affected in various ways, while there appears to be no effect for those who select the Con message.

We can now turn to panel (b) of Table 1, which displays the low importance condition. Recall that low importance diminished attitude-congruent selective exposure, such that the groups represented by the PSE and CSE are more similar to one another here than in the high importance condition, where the audiences for the two messages were more divided along attitudinal lines. The SATEs in this condition are very similar to those for the sample as a whole: the Pro message is seen as more effective, it increases policy support as measured by both $t2$ opinions and opinion changes over time, and decreases the rate of requesting a follow-up email, while there is no effect on intended information-seeking. Turning to the PSEs in column 2, we see a pattern generally consistent with the SATEs in effect direction but only one effect is statistically distinguishable from zero: exposure to the Pro message leads to a substantial decrease in requests for the follow-up email. The CSEs in the third column are also consistent with the SATEs in direction but only the effect on opinion changes (wherein exposure to the Pro message increases support) is statistically distinguishable from zero. Under low importance, there appears to be relatively little effect heterogeneity: effects for Pro-selectors and Con-selectors are fairly consistent with one another, meaning the SATEs provide inferences that seem to apply equally well to those preferring and not preferring the treatment.

Finally, panel (c) of Table 1 shows results for the high importance condition. Recall that under high importance, respondents were much more likely to self-select a treatment message congruent with their prior ($t1$) opinion than in the low importance condition (the

odds-ratio was 4.25). The consequence of this for inferences about the effects of the treatment messages should be immediately clear: while the SATEs in this condition mirror those for the low importance condition and the sample as a whole (with the exception of the email message outcome), the PSEs and CSEs differ considerably. Under high importance, there is a very large positive PSE on argument evaluation (meaning the Pro message was seen much more favorably than the Con message) and opinions were moved nearly 20% more supportive among Pro-selectors. The PSEs for the two information-seeking measures were not distinguishable from zero.

Looking at the CSEs in column 3, the effects on argument evaluation and opinion now are substantively negative, which indicates that the Pro message was seen as less effective than the Con message by these individuals and actually produced a backlash that made these respondents less supportive of the policy, though the large standard errors mean these effects are not statistically distinguishable from zero. Looking at the information-seeking measures, there is a substantively large positive effect of the Pro message on email requests but again the large standard errors mean this is difficult to distinguish from statistical noise. The high importance conditions therefore suggest a quite different pattern of causal effects than the low importance conditions, in both size and direction of effects for the Pro-selectors and Con-selectors.

These results are striking. As expected, the subsample effects vary widely between the high and low importance conditions. If one were to use these data to make inferences about the effects of exposure to a political argument, those inferences would differ substantially depending whether one looked only at the randomized SATEs or at effects derived from the hybrid design, and those effects would differ even further depending on the level of attitude-congruent selective exposure driving the participants' choices of message. Thus relatively simple experimental design decisions could lead us to infer that the difference between receiving the Pro and Con messages on opinions was any of the following:

– increased policy support (full sample, high importance, and low importance SATEs)

- increased support only among those inclined to receive the message (full sample or high importance PSEs)
- no effect on those inclined to receive it (low importance PSE)
- no effect on those disinclined to receive it (full sample, high importance, and low importance CSEs)
- possible backfire effect on those disinclined to receive it (high importance CSE)

If we were interested in the hypothetical, universal application of the Pro rather than Con message to all members of the population, the SATE would tell us that such an intervention would increase policy support. If we were instead interested to know how receiving the Pro rather than Con message affected those who chose the Pro message — as we would be, if we wanted to understand communication effects as they might actually occur in the real world — then our inference about that effect depends on whether the issue is seen as personally important or not. If it is personally important, then those choosing the Pro message would be made more supportive by that exposure; if the issue were less important, then the effect among Pro selectors would likely be smaller (because that group would consist of a more attitudinally diverse set of the public). The results strongly suggest effect heterogeneity: those preferring a message respond to it differently from those who disprefer it.[10]

Discussion

What do political communication researchers want to know? Is it what would happen if everyone was exposed to a political message? Or, is it instead what effect a message has on individuals who are exposed to it? Arguably it could be both, but while we know quite

[10]Similarly interesting is the fact that the SATEs are very similar for the high and low importance groups despite the underlying group-specific effects differing considerably. This means that the randomized experiment offers a very similar inference about the effects of exposure to a political communication even when the underlying group-specific treatment effects for individuals who prefer the Pro and Con messages are quite different. While social scientists frequently acknowledge that the SATE is necessarily an average of underlying individual-level effects and may not represent the effect for any given unit or group of units, these results show that the SATE can be quite misleading.

a bit about the former due to increasing use of randomized experiments to study political phenomena, we know surprisingly little about the effects political communications would have on those who do and do not receive them in the real-world. The present research shows that the SATE masks substantial effect heterogeneity among those that would be inclined to select distinct messages (i.e., self-select into the various treatment alternatives).

How does this change our understanding of political communications? The results showing heterogeneous effects bolster the argument — made by researchers dating as far back as Hovland (1959) — that we should try to understand how messages affect those that would actually receive the message in the real-world rather than just study the hypothetical effect of communicating the message to everyone, regardless of their information preferences. While the SATE might lead us to believe that a message increases policy support thereby inviting a policy implication of expanding provision of the message, that average effect can mask a reality in which those inclined to receive are affected while those disinclined are not even when they are forced to receive it (as was the case here). Or the treatment might have no effect on those who prefer it, while affecting only those who would never choose it of their own accord. The SATE can mislead us about who is affected and how much, and therefore misguide our understanding of the role of communications on the public. If an experiment shows a treatment effect but it actually occurs only for those who would never choose it, what have we learned about political communication? Taking the leap from SATE to practical implications without acknowledging this heterogeneity might be fundamentally problematic.

Yet even as the present research highlights the value of incorporating choice elements in an experiment, the results demonstrate that uncovering such heterogeneity is complicated by the complexity of individuals' choice behavior. While a self-selection experiment presents a face valid form of experimental realism, the design is no panacea for ecological validity concerns. Pre-choice manipulations — like the importance manipulation used here – may be useful for better understanding such dynamics; the use of alternative choice sets may also

clarify the conditional nature of treatment choices (see, for example, Leeper 2014). While the findings suggest incorporating self-selection into a randomized experiment can be fruitful, more research is needed on how best to study self-selection processes.

References

Albertson, Bethany L., and Adria Lawrence. 2009. "After the Credits Roll: The Long-Term Effects of Educational Television on Public Knowledge and Attitudes." *American Politics Research* 37(2): 275–300.

Ansolabehere, Stephen, Shanto Iyengar, Adam F. Simon, and Nicholas A. Valentino. 1994. "Does Attack Advertising Demobilize the Electorate?" *American Political Science Review* 88(4): 829–838.

Arceneaux, Kevin, and Martin Johnson. 2012. *Changing Minds or Changing Channels? Media Effects in the Era of Viewer Choice*. Chicago, IL: Chicago University Press.

Baum, Matthew A. 2002. "Sex, Lies, and War: How Soft News Brings Foreign Policy to the Inattentive Public." *American Political Science Review* 96(1): 91–109.

Bennett, W. Lance, and Shanto Iyengar. 2008. "A New Era of Minimal Effects? The Changing Foundations of Political Communication." *Journal of Communication* 58: 707–731.

Berinsky, Adam J., and Donald R. Kinder. 2006. "Making Sense of Issues Through Media Frames: Understanding the Kosovo Crisis." *The Journal of Politics* 68(3): 640–656.

Bolsen, Toby, and Thomas J. Leeper. 2013. "Self-Interest and Attention to News among Issue Publics." *Political Communication* In press(3): 329–348.

Boninger, David S., Jon A. Krosnick, and Matthew K. Berent. 1995. "Origins of Attitude Importance: Self-Interest, Social Identification, and Value Relevance." *Journal of Personality and Social Psychology* 68(1): 61–80.

Brewer, Paul R. 2003. "Values, Political Knowledge, and Public Opinion About Gay Rights: A Framing-Based Account." *Public Opinion Quarterly* 67(2): 173–201.

Brewer, Paul R., and Kimberly Gross. 2005. "Values, Framing, and Citizens' Thoughts about Policy Issues: Effects on Content and Quantity." *Political Psychology* 26(6): 929–948.

Chong, Dennis, and James N. Druckman. 2007. "Framing Public Opinion in Competitive Democracies." *American Political Science Review* 101(4): 637–655.

Disch, Lisa. 2011. "Toward a Mobilization Conception of Democratic Representation." *American Political Science Review* 105(01): 100–114.

Ditto, Peter H., James A. Scepansky, Geoffrey D. Munro, Anne Marie Apanovitch, and Lisa K. Lockhart. 1998. "Motivated Sensitivity to Preference Inconsistent Information." *Journal of Personality and Social Psychology* 75(1): 53–69.

Druckman, James N., Jordan Fein, and Thomas J. Leeper. 2012. "A Source of Bias in Public Opinion Stability." *American Political Science Review* 106(2): 430–454.

Feldman, Lauren, Edward W. Maibach, Connie Roser-Renouf, and Anthony A. Leiserowitz. 2011. "Climate on Cable: The Nature and Impact of Global Warming Coverage on Fox News, CNN, and MSNBC." *The International Journal of Press/Politics* , 2–31.

Feldman, Lauren, Natalie Jomini Stroud, Bruce Bimber, and Magdalena E. Wojcieszak. 2013. "Assessing Selective Exposure in Experiments: The Implications of Different Methodological Choices." *Communication Methods and Measures* (August): 37–41.

Gaines, Brian J., and James H. Kuklinski. 2011*a*. "Experimental Estimation of Heterogeneous Treatment Effects Related to Self-Selection." *American Journal of Political Science* 55(3): 724–736.

Gaines, Brian J., and James H. Kuklinski. 2011*b*. "Treatment Effects." In *Cambridge Handbook of Experimental Political Science*, eds. James N. Druckman, Donald P. Green, James H. Kuklinski, and Arthur Lupia. New York: Cambridge University Press , 445–458.

Garrett, R. Kelly. 2009*a*. "Echo Chambers Online? Politically Motivated Selective Exposure Among Internet News Users." *Journal of Computer-Mediated Communication* 14(2): 265–285.

Garrett, R. Kelly. 2009*b*. "Politically Motivated Reinforcement Seeking: Reframing the Selective Exposure Debate." *Journal of Communication* 59(4): 676–699.

Garrett, R. Kelly, Dustin Carnahan, and Emily K. Lynch. 2013. "A Turn Toward Avoidance? Selective Exposure to Online Political Information, 2004–2008." *Political Behavior* 35(1): 113–134.

Hart, William, Dolores Albarracín, Alice H. Eagly, Inge Brechan, Matthew J Lindberg, and Lisa Merrill. 2009. "Feeling Validated Versus Being Correct: A Meta-Analysis of Selective Exposure to Information." *Psychological Bulletin* 135(4): 555–588.

Holbrook, Allyson L., Matthew K. Berent, Jon A. Krosnick, Penny S. Visser, and David S. Boninger. 2005. "Attitude Importance and the Accumulation of Attitude-Relevant Knowledge in Memory." *Journal of Personality and Social Psychology* 88(5): 749–769.

Hovland, Carl I. 1959. "Reconciling Conflicting Results Derived from Experimental and Survey Studies of Attitude Change." *American Psychologist* 14(1): 8–17.

Iyengar, Shanto, and Donald R. Kinder. 1987. *News That Matters: Television and American Opinion*. Chicago, IL: The University Of Chicago Press.

Iyengar, Shanto, and Kyu S. Hahn. 2009. "Red Media, Blue Media: Evidence of Ideological Selectivity in Media Use." *Journal of Communication* 59(1): 19–39.

Iyengar, Shanto, Kyu S. Hahn, Jon A. Krosnick, and John Walker. 2008. "Selective Exposure to Campaign Communication: The Role of Anticipated Agreement and Issue Public Membership." *The Journal of Politics* 70(01): 186–200.

Kim, Young Mie. 2007. "How Intrinsic and Extrinsic Motivations Interact in Selectivity:

Investigating the Moderating Effects of Situational Information Processing Goals in Issue Publics' Web Behavior." *Communication Research* 34(2): 185–211.

Kim, Young Mie. 2009. "Issue Publics in the New Information Environment: Selectivity, Domain Specificity, and Extremity." *Communication Research* 36(2): 254–284.

Krosnick, Jon A. 1989. "Attitude Importance and Attitude Accessibility." *Personality and Social Psychology Bulletin* 15(3): 297–308.

Krosnick, Jon A. 1990. "Government Policy and Citizen Passion: A Study of Issue Publics in Contemporary America." *Political Behavior* 12(1): 59–92.

Kunda, Ziva. 1990. "The Case for Motivated Reasoning." *Psychological Bulletin* 108(3): 480–498.

Lau, Richard R., and David P. Redlawsk. 2006. *How Voters Decide: Information Processing in Election Campaigns*. New York: Cambridge University Press.

Leeper, Thomas J. 2014. "The Informational Basis for Mass Polarization." *Public Opinion Quarterly* 78(1): 27–46.

Levendusky, Matthew S. 2012. *Partisan News that Matters: How Partisan News Impacts American Politics*. Chicago, IL: The University of Chicago Press.

Levendusky, Matthew S. 2013. "Why Do Partisan Media Polarize Viewers?" *American Journal of Political Science* 57(3): 611–623.

Lodge, Milton, and Charles S. Taber. 2000. "Three Steps Toward a Theory of Motivated Political Reasoning." In *Elements of Reason: Cognition, Choice, and the Bounds of Rationality*, eds. Arthur Lupia, Mathew D. McCubbins, and Samuel L. Popkin. New York: Cambridge University Press , 183–213.

McDermott, Rose. 2011. "Internal and External Validity." In *Cambridge Handbook of Experimental Political Science*, eds. James N. Druckman, Donald P. Green, James H. Kuklinski, and Arthur Lupia. New York: Cambridge University Press.

Miller, Joanne M., and Jon A. Krosnick. 2000. "News Media Impact on the Ingredients of Presidential Evaluations: Politically Knowledgeable Citizens are Guided by a Trusted Source." *American Journal of Political Science* 44(2): 301–315.

Nelson, Thomas E., Rosalee A. Clawson, and Zoe M. Oxley. 1997. "Media Framing of a Civil Liberties Conflict and Its Effect on Tolerance." *American Political Science Review* 91(3): 567–583.

Nickerson, Raymond S. 1998. "Confirmation Bias: A Ubiquitous Phenomenon in Many Guises." *Review of General Psychology* 2(2): 175–220.

Petty, Richard E., and John T. Cacioppo. 1986. "The Elaboration Likelihood Model of Persuasion." *Advances in Experimental Social Psychology* 19: 123–205.

Prior, Markus. 2007. *Post-Broadcast Democracy: How Media Choice Increases Inequality in Political Involvement and Polarizes Elections*. New York: Cambridge University Press.

Sears, David O., and Jonathan L. Freedman. 1967. "Selective Exposure to Information: A Critical Review." *Public Opinion Quarterly* 31(2): 194–213.

Shadish, William R., Thomas D. Cook, and Donald T. Campbell. 2001. *Experimental and Quasi-Experimental Designs for Generalized Causal Inference*. Boston, MA: Houghton-Mifflin.

Smith, Steven M., Leandre R. Fabrigar, and Meghan E. Norris. 2008. "Reflecting on Six Decades of Selective Exposure Research: Progress, Challenges, and Opportunities." *Social and Personality Psychology Compass* 2(1): 464–493.

Stroud, Natalie Jomini. 2011. *Niche News: The Politics of News Choice*. New York: Oxford University Press.

Taber, Charles S., and Milton Lodge. 2006. "Motivated Skepticism in the Evaluation of Political Beliefs." *American Journal of Political Science* 50(3): 755–769.

Tichenor, P.J., G.A. Donohue, and C.N. Olien. 1970. "Mass Media Flow and Differential Growth in Knowledge." *Public Opinion Quarterly* 34(2): 159–170.

Visser, Penny S., Allyson L. Holbrook, and Jon A. Krosnick. 2007. "Knowledge and Attitudes." In *Public Opinion Research (SAGE Handbook)*. Thousand Oaks, CA: Sage Publications , 123–140.

Appendix A. Survey Material

Attitude Importance Manipulation:
High: A new law is currently moving through Congress that would require your electricity provider to purchase energy from renewable sources (e.g. wind and solar). This is relevant to you since it will influence your energy bills and the environment. The law would go into effect immediately.
Low: A few legislators in Congress have proposed a bill that would require electricity providers to purchase energy from renewable sources (e.g. wind and solar). This is probably not directly relevant to you because Congress does not appear to be ready to act on the bill and even if they did it is unlikely to personally affect you.

Information Choice Manipulation:
Assigned: [Randomly assigned to Pro or Con]

Choice: Before you proceed, please choose one of the following brief passages to read:
"Renewable Energy Rules Beneficial"
"Renewable Energy Rules Ineffective"

Article Text:
Pro: "Renewable Energy Rules Beneficial"
The proposed federal law would create uniform nationwide standards which is necessary since many states have not adopted renewable energy provisions. The new standards would require electricity utilities to produce between 10% and 30% of their energy from renewable sources especially wind power as well as potentially innovative new sources of energy. This in turn would reduce pollution. The impact on consumers is also affordable: adopting a nationwide standard would increase monthly electricity bills by only about 1%. Renewable standards therefore reduce reliance on fossil fuels for energy production without dramatically increasing costs to American consumers.

Con: "Renewable Energy Rules Ineffective"
The proposed federal law would intervene in state policies and regulate private businesses to create uniform nationwide standards, where up until now many states have not adopted renewable energy provisions. The new standards would drive down innovation by requiring energy utilities to adopt specific technologies (e.g. wind power) rather than directly targeting the reduction of polluting greenhouse gas emissions. The impact on consumers is also problematic: adopting a nationwide standard would increase monthly electricity bills by up as much as 4%. Renewable standards therefore increase the cost of energy through government regulation without directly addressing potential environmental impacts of energy production from fossil fuels.

Appendix B. Supplementary Descriptive Statistics

Table B1a. Argument Evaluations, by Treatment Condition

	Mean (SE)	N
High Choice Pro	0.75 (0.01)	204
High Choice Con	0.57 (0.03)	67
High Captive Pro	0.65 (0.03)	66
High Captive Con	0.47 (0.04)	68
Low Choice Pro	0.69 (0.02)	173
Low Choice Con	0.53 (0.03)	101
Low Captive Pro	0.68 (0.03)	67
Low Captive Con	0.58 (0.03)	67

Note: Cell entries are treatment group means with standard errors in parentheses and higher values indicating greater perceived argument effectiveness (scaled 0-1).

Table B1b. Argument Evaluations, by Importance and Choice Condition

	Mean (SE)	N
Choice High	0.70 (0.02)	271
Captive High	0.56 (0.03)	134
Choice Low	0.63 (0.02)	274
Captive Low	0.63 (0.02)	134

Note: Cell entries are treatment group means with standard errors in parentheses and higher values indicating greater perceived argument effectiveness (scaled 0-1).

Figure B1. Mean Argument Evaluations, by Importance and Choice Condition and Congruence with $t1$ Opinion

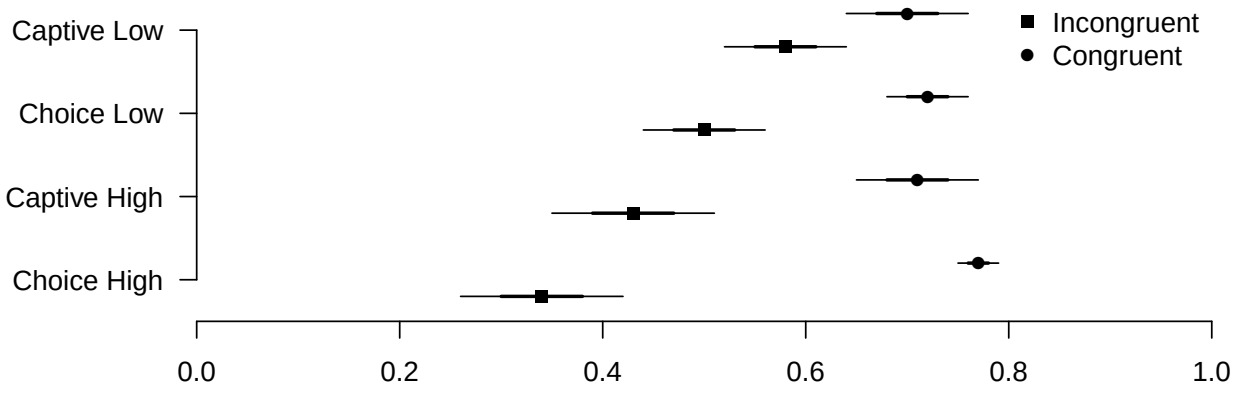

Note: Figure displays mean argument evaluations (and bars representing one and two standard errors of the mean) by importance and choice Condition, separately for arguments congruent or incongruent with respondents' $t1$ opinions.

Figure B2. Argument Evaluation Bias, by Importance and Choice Condition

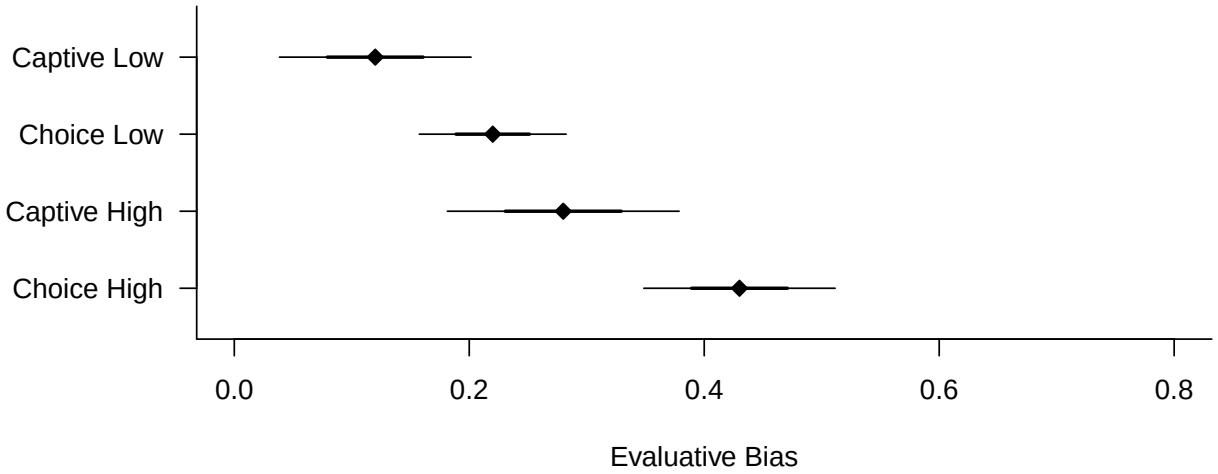

Note: Figure displays the bias toward seeing attitude-congruent messages as more effective than attitude-incongruent messages. Points represent difference-in-differences estimates along with bars representing one and two associated standard errors, based on the mean argument ratings displayed in Figure C1.

Table B2a. $t1$ **and** $t2$ **Opinions, by Treatment Condition**

	t1 Mean (SE)	t2 Mean (SE)	t2-t1 Mean (SE)	N
Control	0.78 (0.03)	0.65 (0.03)	-0.14 (0.03)	66
High Choice Pro	0.86 (0.01)	0.82 (0.01)	-0.04 (0.02)	204
High Choice Con	0.47 (0.04)	0.39 (0.04)	-0.08 (0.04)	67
High Captive Pro	0.76 (0.03)	0.70 (0.04)	-0.07 (0.03)	66
High Captive Con	0.75 (0.04)	0.58 (0.04)	-0.18 (0.04)	68
Low Choice Pro	0.82 (0.02)	0.74 (0.02)	-0.07 (0.02)	173
Low Choice Con	0.69 (0.02)	0.55 (0.03)	-0.14 (0.03)	101
Low Captive Pro	0.72 (0.04)	0.70 (0.04)	-0.02 (0.03)	67
Low Captive Con	0.78 (0.03)	0.65 (0.03)	-0.13 (0.03)	67

Note: Cell entries are treatment group means with standard errors in parentheses and higher values indicating greater support (scaled 0-1).

Table B2b. $t1$ **and** $t2$ **Opinions, by Importance and Choice Condition**

	t1 Mean (SE)	t2 Mean (SE)	t2-t1 Mean (SE)	N
Choice High	0.76 (0.02)	0.72 (0.02)	-0.05 (0.02)	271
Captive High	0.76 (0.02)	0.64 (0.03)	-0.12 (0.02)	134
Choice Low	0.77 (0.02)	0.67 (0.02)	-0.10 (0.02)	274
Captive Low	0.75 (0.02)	0.67 (0.03)	-0.08 (0.02)	134

Note: Cell entries are treatment group means with standard errors in parentheses and higher values indicating greater support (scaled 0-1).

Table B3a. Information Seeking, by Treatment Condition

	Subjective	Email	N
Control	0.62 (0.03)	0.41 (0.07)	66
High Choice Pro	0.68 (0.02)	0.54 (0.03)	204
High Choice Con	0.59 (0.04)	0.43 (0.07)	67
High Captive Pro	0.65 (0.03)	0.61 (0.05)	66
High Captive Con	0.67 (0.03)	0.59 (0.05)	68
Low Choice Pro	0.54 (0.02)	0.31 (0.05)	173
Low Choice Con	0.49 (0.03)	0.34 (0.07)	101
Low Captive Pro	0.54 (0.03)	0.25 (0.09)	67
Low Captive Con	0.56 (0.03)	0.52 (0.06)	67

Note: Cell entries are treatment group means with standard errors in parentheses and higher values indicating greater information (scaled 0-1) and likelihood of requesting email (0/1).

Table B3b. Information Seeking, by Importance and Choice Condition

	Subjective	Email	N
Choice High	0.66 (0.01)	0.52 (0.03)	271
Captive High	0.66 (0.02)	0.60 (0.03)	134
Choice Low	0.52 (0.02)	0.32 (0.04)	274
Captive Low	0.55 (0.02)	0.39 (0.05)	134

Note: Cell entries are treatment group means with standard errors in parentheses and higher values indicating greater information (scaled 0-1) and likelihood of requesting email (0/1).