

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

Thomas J. Leeper*

Abstract

Ecological validity is vital to experimental research because designs that are too artificial may not speak to any real-world political phenomenon. One such concern is treatment of self-selection: if individuals in the real-world self-select treatments, such as political communications, how well does the sample average treatment effect estimate the effects of message exposure for those individuals who would—if given the choice—opt-in to and out of receiving treatment? This study shows that randomization masks effect heterogeneity between individuals who would select different messages if given the choice. Yet, such selections are themselves complex, revealing additional challenges for realistically studying treatments prone to self-selection. The evidence of effect heterogeneity raises questions about the appropriateness of random assignment experiments for studying political communication and the results more broadly advance our understanding of citizens' selection into and responses to communications when, as they often do, have choice over what messages to receive.

Keywords: self-selection, effect heterogeneity, ecological validity, attitude importance

Experiments are seen as a gold standard because randomized assignment offers unparalleled *internal* validity. An important and less studied concern is *ecological* validity, or the “realism” of the experiment. Ecological validity is vital because studies that are too artificial may not speak to any real-world political phenomenon. In particular, the recognition that in the real-world individuals self-select their own “treatments” raises a concern that randomized exposure may not constitute a realistic exploration of political communication (Arceneaux and Johnson 2012; Druckman et al. 2012; Gaines and Kuklinski 2011a, b; Lau and

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 and reviewers 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.

*Department of Government, London School of Economics and Political Science, Houghton Street, London WC2A 2AE, United Kingdom, e-mail: thosjleeper@gmail.com

Redlawsk 2006). Given this self-selection, what can we learn from a randomized experiment about the effect of a message on those individuals who would—if given the choice—opt into or out of receiving it? To answer this, I use a population-based survey experiment to test how individuals respond to policy arguments that 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, with the results hinging on how issue importance drives different subsets of respondents to self-select each message.

EFFECTS OF SELF-SELECTED POLITICAL COMMUNICATIONS

Arguments conveyed by media and political elites are widely seen as an important source of information that citizens can use when forming preferences (Chong and Druckman 2007; Disch 2011). Evidence of this influence comes largely from experiments that expose participants to different messages and measure effects on argument evaluations, opinions, issue importance, information-seeking, etc. (see Ansolabehere et al. 1994; Arceneaux and Johnson 2012; Berinsky and Kinder 2006; Brewer and Gross 2005; Iyengar and Kinder 1987; Miller and Krosnick 2000; Nelson et al. 1997; Petty and Cacioppo 1986). Questions have been raised, however, about these kinds of studies 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” (Bennett and Iyengar 2008, 724; Hovland 1959, 9). A randomized experiment cannot identify the effect of a message for those who chose to view it or the effect for those who chose not to view it. Instead, it can identify only the sample average treatment effect (SATE), which averages the effects for the two subgroups that choose to be treated and choose to be untreated.

Hovland encouraged 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 effects for these groups? Consider the classic experiment by Nelson et al. (1997) in which participants were assigned to watch a story framing a rally in terms of either free speech or public safety. While we may care about the SATE in this case, our interest in understanding citizens’ interactions with media suggests that we also 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, it is widely known that citizens selectively expose themselves to information (Bennett and Iyengar 2008; Bolsen and Leeper 2013; Feldman et al. 2011; Garrett 2009a, b; Garrett et al. 2013; Iyengar and Hahn 2009; Iyengar et al. 2008; Kim 2007; Smith et al. 2008; Stroud 2011), and meta-analysis suggests that

political selective exposure is especially potent (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 (Baum 2002; Bennett and Iyengar 2008; Iyengar and Hahn 2009; Prior 2007). Given this selectivity, there is value in knowing how individuals are affected by treatments they actually encounter in the real world.

In what ways might the effect of a communication differ across individuals inclined and disinclined to select it? There are two possibilities. One is *homogeneous effects*: the effects of exposure to a communication are the same for those who would choose the message as for those who would not. In the case of such homogeneity, the SATE averages the effects for both groups and there are no limitation to what can be learned from a randomized experiment. Thus, homogeneity occurs if messages are uniformly influential or if a message has no effect on anyone.

The other possibility is *heterogeneous effects*: the effects of exposure to a message are different for those who would choose and not choose the message. Such effect heterogeneity might occur because of motivated reasoning (Kunda 1990; Taber and Lodge 2006), wherein individuals tend to select attitude-reinforcing messages and avoid attitude-incongruent arguments, as well as see attitude-congruent arguments as stronger and more effective than incongruent arguments (Ditto et al. 1998). As such, the effects of communications for message selectors and non-selectors on numerous outcomes—including argument evaluations, opinions, and willingness to acquire more issue-relevant information—are likely to point in opposite directions.

Given the prevalence of motivated thinking, it is reasonable to expect effect heterogeneity rather than homogeneity. To foreshadow and thereby situate that general expectation in the specific design used here, exposure to a message supportive of a policy should be expected to increase support for that policy. Yet, that increase in support—and attendant positive message evaluations and interest in the issue—should apply only to those inclined to receive the message and not to those disinclined to choose the message. Further, because issue importance is one mechanism thought to affect the degree of selective exposure (with high importance linked to greater attitude-congruent message exposure; Holbrook et al. 2005; Leeper 2014; Taber and Lodge 2006), higher importance should exacerbate this heterogeneity and lower importance should mitigate it (given the greater similarity between the audiences for different messages).

EXPERIMENTAL DESIGN

I use a design combining randomized exposure and message self-selection to test for this expected pattern of effect heterogeneity. One-half of participants are assigned to a randomized experiment and the other half participate in an observational study involving treatment self-selection (see Gaines and Kuklinski 2011a). Expanding past work, I also employ three additional features: (1) a pre-treatment manipulation

of issue importance to modify the degree of attitude-congruent selective exposure, (2) a two-wave panel design, and (3) a population-based sample of participants. The issue under investigation is so-called “renewable energy portfolio” standards, which require electrical utilities to produce energy from renewable resources.¹ The first wave of the study occurred in Summer 2010 (hereafter, *t1*) to measure demographics and baseline attitudes. The 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* experiment and enables estimation of opinion effects using within-subjects, pre-/post-treatment changes.

Respondents

Data were collected by Bovitz Research Group of Encino, CA, that 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 represent 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.²

Manipulations

The *t2* experiment involved three manipulations: issue importance, the direction of the policy argument, and whether that argument was self-selected or randomly assigned. [Figure 1](#) provides a visual summary of the design, with treatment group sizes and notation used in defining causal effects.

The first manipulation modified the personal impact of the energy proposal. This serves as an instrument for respondents’ information choices. The logic of the manipulation was that individuals who believe their self-interest is at-stake would display higher importance and be more likely to choose an opinion-congruent message. Importance was manipulated to be high by telling respondents:

¹U.S. Department of Energy, 2003, “Analysis of a 10% Renewable Portfolio Standard,” Office of Integrated Analysis and Forecasting of the Energy Information Administration; Chen, Cliff, Ryan Wisler, 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.

²These demographics roughly match Current Population Survey estimates of the U.S. adult population in 2011 (49% female, 75% white), though the sample was slightly older and better educated than the population as a whole (adult median age 46; 86% high school degrees; 36% university degrees).

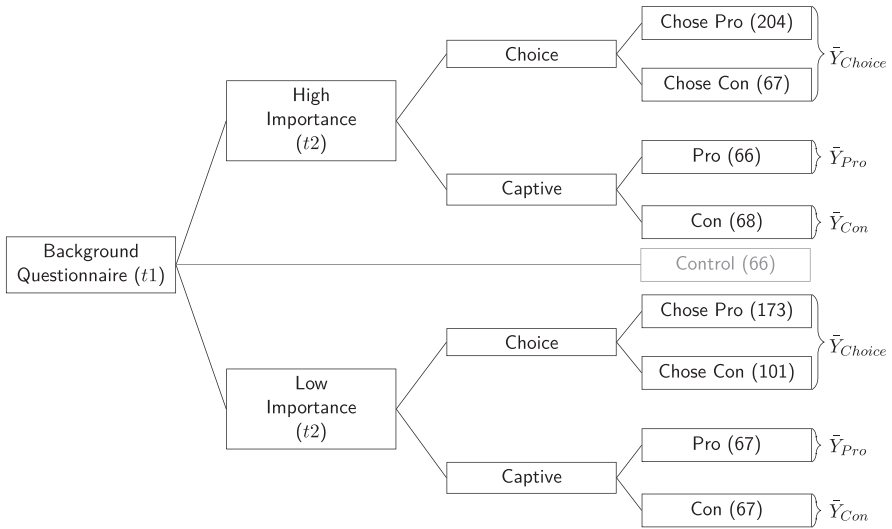


Figure 1

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.

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 the low-importance condition read:

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 manipulation check asked respondents “How important to you personally is your opinion about this renewable energy restriction?” and the results confirm that importance was manipulated. On a 0–1 scale, those in high-importance conditions averaged 0.77 and those in low-importance conditions averaged 0.69, a statistically significant difference ($p = 0.00$).³

³The mean 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 ($p = 0.34$). Three rounds of pilot tests were also conducted with 80 participants each in order to determine the question and manipulation wordings. Participants were recruited from Amazon Mechanical Turk and compensated \$0.20 each. In the final pre-test, the high-importance manipulation

The second manipulation presented participants with either a message supportive (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.”⁴ 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.⁵

The final manipulation involved how the informational treatments were assigned. One-third of the respondents were randomly assigned to read either the Pro or Con argument. The other two-thirds were presented with the headlines for each passage and told to choose one to read, which they were then given.⁶

Measures

Outcome measures included 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. 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 requiring 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 ($r = 0.44$).⁷

achieved a mean importance (on a seven-point scale) of 5.66 (SD = 1.08), while the low-importance manipulation achieved a mean importance of 5.2 (SD = 1.44). The difficulty of manipulating attitude importance (Visser et al. 2007) makes the relatively small apparent effects in the small-*n* pilot reasonable and thus seemed adequate for the main study.

⁴Exact text of the messages is available in Appendix A.

⁵Pilot testing confirmed that messages were equally effective. The Pro message was rated as 5.08 (SD = 1.38) on a seven-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 seven-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.

⁶A control group exposed to no treatment was also included but is not analyzed here.

⁷Treatment group means for all outcome measures are reported in Appendix B.

Estimation Strategy

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

$$\text{SATE} = \bar{Y}_{\text{Pro}} - \bar{Y}_{\text{Con}} \quad (1)$$

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.

The 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*.⁸ These effects are identified by the present design if we are willing to assume (1) that, given random assignment, the choice behavior of those in the choice conditions is on average identical to the unobserved choice behavior of those in the randomized conditions, and (2) the equivalence of potential outcomes for a randomly assigned message versus and the same self-selected message (i.e. that there is no effect of the assignment mechanism; an exclusion restriction). Under these assumptions, we can identify two further effects. One is the effect of the treatment on those who would choose it (the *Pro-Selector Effect*, or PSE):

$$\text{PSE} = \frac{\bar{Y}_{\text{Choice}} - \bar{Y}_{\text{Con}}}{\hat{\alpha}} \quad (2)$$

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 final effect captures the difference between receiving the Pro and Con messages among those who would not choose the Pro message (i.e., the *Con-Selector Effect*, or CSE):

$$\text{CSE} = \frac{\bar{Y}_{\text{Pro}} - \bar{Y}_{\text{Choice}}}{1 - \hat{\alpha}} \quad (3)$$

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*. Each represents the average treatment effect for distinct subsets of the population. Effect homogeneity occurs when the PSE and CSE are identical, and therefore match the SATE. Heterogeneity occurs when these quantities diverge, making the SATE reflective of only one or neither of the underlying subgroup effects.

⁸Because 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.

RESULTS

I begin by examining the information choices made by respondents in the choice conditions. Consistent with expectations, issue importance increased the degree of opinion-congruent selective exposure: 88% of high-importance respondents chose information congruent with their $t1$ opinion, while only 63% of low-importance respondents chose in this way (a statistically significant difference).⁹ This means that respondents are not *fixed types* who would always select in the same way (an assumption in past work). As such, it is reasonable to expect that if there are heterogeneous effects of exposure to the Pro message, the pattern of that heterogeneity is likely to be most clear in the high-importance groups where Pro selectors and Con selectors differ from each other most dramatically.

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). Looking at panel (a), we see 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 corresponding PSE and CSE estimates. In substantive terms, these five SATEs indicate that (1) the Pro argument is seen as more effective than the Con argument, which makes sense given the supportive leanings of the sample, (2) exposure to the Pro message increases $t2$ policy support, (3) this increase in support holds when measured by $t2 - t1$ opinion changes, (4) the Pro message insignificantly reduces plans to seek out information, and (5) the Pro message reduces the likelihood of requesting an informational email.

The PSEs paint a largely similar story, with differences only in effect magnitude. The story is different for CSEs: while the Pro message still increases their policy support over time, there are no other substantively sizable effects and none of the CSEs are statistically distinguishable from zero. In other words, there appears to be

⁹Recall that 50% choosing congruent information would be considered “random” choice. Or, said another way, the odds of a high-importance respondent choosing the opinion-congruent message was roughly 7:1, while the odds for a low-importance subject were only about 1.5:1, yielding an odds ratio of 4.25 (with a 95% confidence interval from 2.68 to 6.75). The odds ratio would be equal to one if there was no effect of importance. This result also holds when looking at supporters and opponents separately: among high-importance subjects, 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 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. 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 pre-test 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 (0.10)
Opinion change ($t_2 - t_1$)	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	0.10 (0.04)	0.09 (0.05)	0.11 (0.09)
Opinion	0.05 (0.05)	0.04 (0.07)	0.07 (0.11)
Opinion change ($t_2 - t_1$)	0.11 (0.05)	0.06 (0.06)	0.20 (0.10)
Information-seeking	− 0.01 (0.05)	− 0.05 (0.06)	0.05 (0.10)
Email request	− 0.27 (0.08)	− 0.32 (0.10)	− 0.17 (0.17)
(b) Low-importance condition			
	SATE	PSE	CSE
Evaluation	0.19 (0.05)	0.31 (0.05)	− 0.20 (0.14)
Opinion	0.12 (0.05)	0.18 (0.06)	− 0.08 (0.17)
Opinion change ($t_2 - t_1$)	0.11 (0.04)	0.17 (0.05)	− 0.07 (0.14)
Information seeking	− 0.02 (0.04)	− 0.01 (0.05)	− 0.03 (0.13)
Email request	0.02 (0.08)	− 0.10 (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).

effect heterogeneity: those inclined to select the Pro message are affected in various ways, while others are unaffected.

Panel (b) displays results for the low-importance condition. Recall that low-importance diminished attitude-congruent message choice, such that the groups represented by the PSE and CSE are more similar to one another here than in the high-importance condition. 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 (measured as $t_2 - t_1$), and decreases requests for the email, while there is no effect on information seeking. The PSEs (column 2) are consistent with the SATEs in direction but only the effect on email requests is statistically distinguishable from zero. The CSEs (column 3) are also consistent with the SATEs in direction, with the exception of the flipped sign on information seeking, but only the effect on opinion changes is statistically significant. These results point to a pattern of effect homogeneity, wherein the SATEs provide inferences that apply equally well to those preferring and not preferring the treatment.

Finally, panel (c) shows results for the high-importance condition, where respondents were much more likely to self-select a message congruent with their t_1 opinion. The consequence of this for inferences about the effects of the

treatment messages should be immediately clear: the SATEs in this condition mirror those for the low-importance condition (except for email requests where there is clearly no effect) and for the sample as a whole, yet the PSEs and CSEs differ considerably.

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 negative and not distinguishable from zero. The CSEs (column 3) show something quite different. The effects on argument evaluation and opinion, while not statistically distinguishable from zero given the large standard errors, would imply substantively negative effects. Similarly, the effect on email requests is difficult to distinguish from statistical noise, but points in substantively positive direction.

These results are striking. As expected, the effects vary widely between the high- and low-importance conditions and there is a clear pattern of effect heterogeneity between PSEs and CSEs in the high-importance conditions. If one were to use these data to make inferences about the effects of exposure to a political argument, those inferences could differ substantially depending on the specific effect estimate chosen for each set of the sample. We could 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 and high-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, the SATE would tell us that such an intervention would increase policy support. If we were instead interested in potentially distinct PSEs and CSEs, our inference would depend on whether the issue is personally important. Clearly, there is value in knowing all of these effects.

DISCUSSION

Do researchers want to know only what would happen if everyone was exposed to a political message? Or, are they also interested in what effect a message has on individuals who are exposed to it? Arguably it could be both, but we know

surprisingly little about the latter given the prevalence of randomized experiments in studies of processes defined by self-selection. The present research shows that the SATE masks substantial effect heterogeneity among those inclined and disinclined to select a given message. While the SATE might lead us to believe that a message increases policy support—thereby implying the desirability of universal message provision—that effect can mask a reality in which those inclined to receive it are affected while those disinclined are not (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. If an experiment shows a treatment effect but it actually occurs only for those who would never choose it, what have we learned? The present research shows that taking the leap from SATE to practical implications without acknowledging this heterogeneity is problematic. Yet, uncovering such heterogeneity is made complex because of individuals' choice behavior. While a self-selection experimental design offers a degree of experimental realism, it is an imperfect remedy for ecological validity concerns. Apparent treatment effects differed here depending on the importance participants attached to the issue at-hand. 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.

SUPPLEMENTARY MATERIALS

The appendix is available online as supplementary material at <https://doi.org/10.1017/XPS.2017.1>

REFERENCES

- Ansolabehere, Stephen, Shanto Iyengar, Adam F. Simon, and Nicholas A. Valentino. 1994. "Does Attack Advertising Demobilize the Electorate?" *American Political Science Review* 88(4): 829–38.
- 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(4): 707–31.
- 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): 640–56.
- Bolsen, Toby and Thomas J. Leeper. 2013. "Self-Interest and Attention to News among Issue Publics." *Political Communication* 30(3): 329–48.
- Brewer, Paul R. and Kimberly Gross. 2005. "Values, Framing, and Citizens' Thoughts about Policy Issues: Effects on Content and Quantity." *Political Psychology* 26(6): 929–48.

- Chong, Dennis and James N. Druckman. 2007. "Framing Public Opinion in Competitive Democracies." *American Political Science Review* 101(4): 637–55.
- Disch, Lisa. 2011. "Toward a Mobilization Conception of Democratic Representation." *American Political Science Review* 105(01): 100–14.
- 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): 430–54.
- 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* 17(1): 3–31.
- Gaines, Brian J. and James H. Kuklinski. 2011a. "Experimental Estimation of Heterogeneous Treatment Effects Related to Self-Selection." *American Journal of Political Science* 55(3): 724–36.
- Gaines, Brian J. and James H. Kuklinski. 2011b. "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, 445–58.
- Garrett, R. Kelly. 2009a. "Echo Chambers Online? Politically Motivated Selective Exposure Among Internet News Users." *Journal of Computer-Mediated Communication* 14(2): 265–85.
- Garrett, R. Kelly. 2009b. "Politically Motivated Reinforcement Seeking: Reframing the Selective Exposure Debate." *Journal of Communication* 59(4): 676–99.
- 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): 113–34.
- Hart, William, Dolores Albarraacín, 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): 555–88.
- 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): 749–69.
- 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.
- Kunda, Ziva. 1990. "The Case for Motivated Reasoning." *Psychological Bulletin* 108(3): 480–98.
- 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.
- 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): 301–15.
- 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): 567–83.
- 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.
- 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): 755–69.
- Visser, Penny S., Allyson L. Holbrook, and Jon A. Krosnick. 2007. "Knowledge and Attitudes." In *The SAGE Handbook of Public Opinion Research*, eds. Wolfgang Donsbach and Michael W. Traugott. Thousand Oaks, CA: Sage Publications, 123–40.