

Estimating the causal effects of cognitive effort and policy information on party cue influence ¹

Ben M. Tappin ² & Ryan T. McKay ³

This is a working paper and has not yet been peer-reviewed.

Date posted: February 14, 2021

Abstract

Why do party cues influence public opinion? A long-standing and influential theory holds that party cues function as heuristics, stand-ins for the lack of policy information and motivation to engage in effortful thinking that characterizes the average person. A key prediction follows that the influence of party cues would diminish if only people were to possess more information about policy, a greater propensity for effortful thinking, or both. This prediction has escaped decisive empirical testing to date, leaving in its wake a string of mixed results. Here, we characterize the challenges that limit previous tests, and we report on two large, novel experiments designed to overcome these challenges. Our experiments indicate that exposure to substantive policy information causally attenuates the influence of party cues, but engagement in effortful thinking per se does not. Our results provide new evidence, and have diverse implications, for the heuristic theory of party cue influence.

¹ We are grateful to Jamie Druckman for his helpful comments and suggestions, and to the Cogito Foundation for funding this research. All remaining errors are ours. The appendix, data and analysis code for this paper are available online <https://osf.io/374sf/>.

² Massachusetts Institute of Technology. Email benmtappin@gmail.com.

³ Royal Holloway, University of London.

Understanding the causes of the public’s policy opinions is a long-standing goal of political science. Among the most well-established phenomena therein is that the public are influenced by cues from political parties about which policies to support and which to oppose (for reviews see (Bullock 2011, 2020; Tappin 2020)). However, a clear understanding of *why* the public are influenced by “party cues” when forming their policy opinions remains elusive. This question has taken on new impetus in recent times, in light of evidence that public compliance with COVID-19 policies may have been deleteriously influenced by party cues (Bisbee and Lee 2020). A long-standing and influential theory holds that people rely on party cues because they lack substantive information about policy, the motivation to think for themselves in order to form an opinion, or both. In other words, the party cue functions as a heuristic; a stand-in for access to substantive policy information and engagement in cognitive effort. A key prediction of this theory is that the influence of party cues would diminish if only people were to possess more information about policy, a greater propensity for effortful thinking, or both.

There have been numerous empirical tests of this prediction in the preceding three decades. However, the results of many such tests are difficult to interpret in light of the study designs used. For example, the dominant study design operationalizes awareness of policy information and propensity for effortful thinking via observed covariates. This makes it difficult to distinguish between the distinct causal effects of policy information and cognitive effort on party cue influence; and, in addition, to distinguish their effects from confounding bias caused by other plausible moderators of party cue influence, such as pretreatment exposure to the party cue, or partisan motivation. A small handful of studies obviate these challenges via randomization of the relevant variables, but omit important control groups and make other design and analytic choices that constrain what is possible to learn from the results.

In this paper, we characterize the challenges to both causal identification and empirical generalization that limit previous research on this topic. We then report the results of two large,

novel experiments designed to overcome these challenges. In our experiments, we randomly and independently assigned both exposure to policy information and engagement in effortful thinking, as well as party cues, and measured American adults' opinions regarding 20 contemporary U.S. policy issues. We observe high levels of compliance with both the information and effort treatments.

Our primary finding is that exposure to substantive policy information causally attenuates the influence of party cues on average, but engagement in effortful thinking per se does not—neither by itself, nor when paired with exposure to policy information. Subsequent analyses characterize the heterogeneity in these results across policy issues and notable demographic subgroups. We also find evidence that people exposed to party cues behave qualitatively similarly to those exposed only to substantive policy information, consistent with the notion that party cues provide a valid informational basis for people's opinion formation. Our results provide new evidence, and have diverse implications, for the long-standing and influential theory that party cues function as heuristics.

Theory and causal identification

A large body of evidence indicates that the average American lacks information regarding politics in general, and policy in particular (Delli Carpini and Keeter 1996; Somin 2016); holds policy opinions that are unstable and ideologically incoherent (Converse 2006; Kinder and Kalmoe 2017); and has little interest in thinking about politics and current affairs (Prior 2018). At the same time, a similarly large body of evidence indicates that party cues reliably influence people's policy opinions—that is, on average people reliably update their opinions toward in-party positions and away from out-party positions when these are learned (Bullock 2020).

Set against this empirical backdrop, a long-standing and influential theory for *why* people's policy opinions are influenced by party cues holds that the cues function as heuristics,

stand-ins for the lack of policy-relevant information, motivation to engage in effortful thinking about policy, or both, that characterizes the average person (Downs 1957; Kam 2005; Lau and Redlawsk 2001; Lupia and McCubbins 1998; Mondak 1993; Popkin 1995). In short, the party cue is used *in lieu* of awareness of substantive policy information and engagement in cognitive effort. This theory has often been interpreted as encouraging for normative theories of democracy, because it suggests that, despite their lack of substantive policy information and motivation, voters can nevertheless “learn what they need to know” in order to form policy opinions and make decisions that approximate those they would have done were they better informed (Hobolt 2006; Lupia 1994; Lupia and McCubbins 1998; Mondak 1993).

A key prediction that follows from the heuristic theory of party cue influence is that, in the counterfactual world where people possessed more information about policy, a greater propensity for effortful thinking, or both, the influence of party cues would be diminished. In other words, there exist *causal interactions* between exposure to the party cue on the one hand, and access to substantive policy information and engagement in cognitive effort on the other (Acharya, Blackwell, and Sen 2018; Imai, Tingley, and Yamamoto 2013). The results of numerous studies offer tests of this key prediction.

Tests using covariate designs

In an influential study on the policy issue of food irradiation technology, Kam (2005) estimated that party cue influence was weakest among those people who were more politically “aware”, operationalized via the number of correct answers to a series of factual questions about politics, and who thus likely possessed greater amounts of policy-relevant information. This is largely consistent with the heuristic theory. Less consistent with the theory was her corresponding result with people’s self-reported “need for cognition”, defined as the propensity to “engage in and enjoy thinking” (p. 130, (Cacioppo and Petty 1982)), which did not appear to moderate

their reliance on the party cue. A recent direct replication of Kam’s design, however, found that adding more items to the need for cognition scale—thereby reducing measurement error—produced the opposite result expected under the heuristic theory: the influence of the party cue was in fact *stronger* among people with a greater need for cognition (Bakker and Lelkes 2018).

The designs of these studies join numerous others in operationalizing the key theoretical variables—exposure to policy relevant information, and engagement in effortful thinking—using observed covariates. Typically, a series of factual questions about politics, and some variant of self-reported need for cognition or correlate thereof, respectively. Taken together, the findings from this body of work are relatively mixed with respect to the interactions predicted by the heuristic theory. Some studies estimate that higher values of the covariates are associated with *stronger* party cue influence, inconsistent with the theory (Bakker and Lelkes 2018; Bakker, Lelkes, and Malka 2020; Slothuus and de Vreese 2010; Vössing 2020); other studies estimate that higher values of the covariates are associated with *weaker* influence, consistent with the theory (Anduiza, Gallego, and Muñoz 2013; Barber and Pope 2019; Boudreau and MacKenzie 2014; Kam 2005; Mondak 1993; Pannico 2020); while still other studies find little evidence of a reliable moderating relationship in one direction or the other (Bullock 2011; Ehret, Van Boven, and Sherman 2018; Gilens and Murakawa 2002; Kam 2005; Mérola and Hitt 2016; Petersen et al. 2013; Slothuus 2010; Tappin 2020).

Beyond their mixed findings, however, a more fundamental challenge in interpreting the results of these studies is that the covariate design is poorly equipped to identify the predicted causal interactions between exposure to the party cue on the one hand, and access to policy information and engagement in cognitive effort on the other. Figure 1 illustrates why. It shows the predicted causal interactions between exposure to the party cue and awareness of policy information (party cue → policy opinion ← policy information), and exposure to the party cue and engagement in cognitive effort (party cue → policy opinion ← cognitive effort).

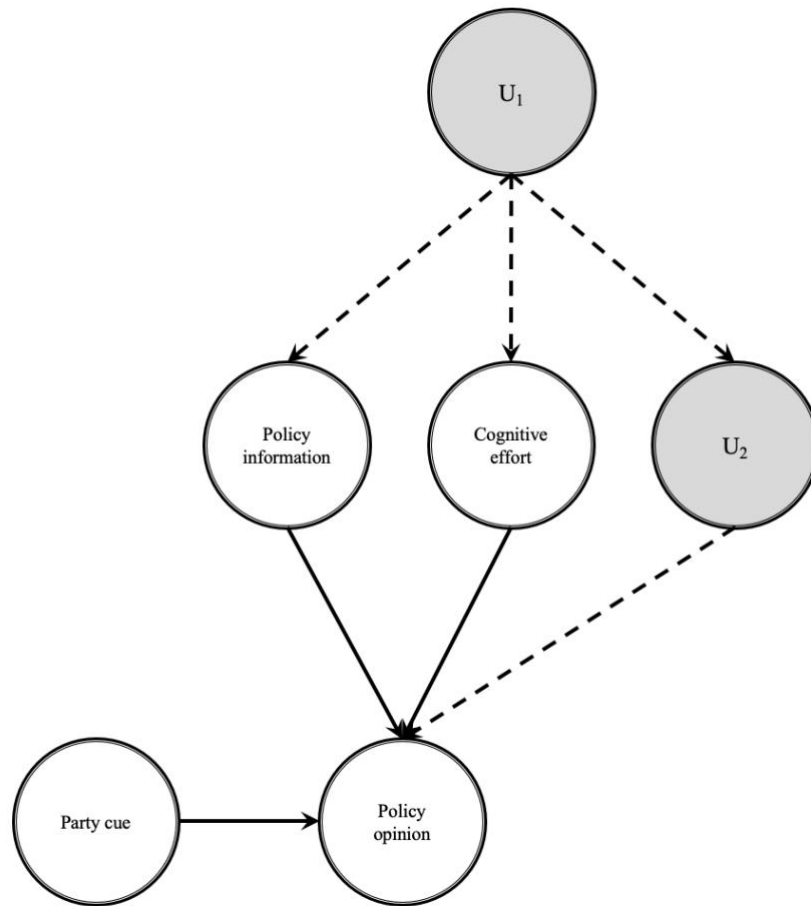


Figure 1. The identification challenge when using observed covariates to operationalize awareness of substantive policy information and engagement in cognitive effort. Here the party cue is assumed to be randomly assigned. Dashed edges are confounding paths.

The challenge for identification of these causal interactions is that there exist plausible *common causes* of both policy information and cognitive effort, indicated by U_1 . General cognitive ability is an example of one such common cause. General cognitive ability is positively associated with self-reported interest in politics (Prior 2018), as well as need for cognition (Cacioppo and Petty 1982; Gottfried et al. 2017). Thus, people with higher cognitive ability will tend to have greater (prior) exposure to policy-relevant information, as well as a higher propensity to engage in effortful thinking. Similarly, another plausible common cause is interest in politics per se, because people with greater interest will tend to possess a higher propensity to engage in effortful thinking about political issues in particular, as well as having greater prior exposure to policy-relevant information. Following the same logic, a third common cause could be peer group incentives to appear politically informed (Marshall 2019). There may be other common causes that typically remain unobserved or not considered. This common cause confounding has two implications for identification of the predicted causal interactions in the covariate design.

First, it renders it deeply challenging to distinguish *between* the predicted causal interactions. For example, observing an interaction between exposure to the party cue and self-reported need for cognition could reflect the respective causal interaction (party cue \rightarrow policy opinion \leftarrow *cognitive effort*), but it could also be spurious; arising instead because of the alternative causal interaction (party cue \rightarrow policy opinion \leftarrow *policy information*). Distinguishing between these causal interactions is not simply theoretical navel-gazing. On the contrary, it may be highly desirable if one's goal is to most effectively intervene in the system to reduce (or enhance) the influence of party cues on public opinion (VanderWeele and Tchetgen Tchetgen 2014). One strategy for distinguishing the interactions is statistical adjustment; that is, measure both covariates and estimate the interactions simultaneously. Alternatively, measure and adjust for all unobserved common causes (U_1).

Unfortunately, to credibly execute the adjustment strategy requires strong and implausible assumptions regarding the measurement of the key variables. For example, the covariates of exposure to policy information and propensity for cognitive effort are both psychological constructs, and thus not directly observable. This means they will likely always be measured with some nontrivial amount of error. Measurement error poses an acute challenge to the adjustment strategy because of the high probability of residual confounding; that is, confounding that remains in spite of statistical adjustment, owing to imperfect measurement of the confounders (Westfall and Yarkoni 2016). Similarly, adjusting for U_1 would require knowledge of, and the ability to measure, every possible common cause of both exposure to policy information and propensity for cognitive effort—an implausible prospect.

Difficulty distinguishing between the distinct causal interactions predicted by the heuristic theory is the first implication of the common cause confounding illustrated in Figure 1. The second, and perhaps more detrimental, implication is the difficulty in distinguishing these interactions from yet other interactions predicted by alternative theories.

For example, the magnitude of party cue influence estimated in any given study will depend in part upon people's prior knowledge about where the relevant parties stand on the policy issue in question (Gaines, Kuklinski, and Quirk 2007; Slothuus 2016). If the party positions are clearly defined and largely common knowledge, many people will have already been “treated” by the party cue. Thus, the study's estimate of influence will be biased downwards. This causal interaction between the party cue and the likelihood of pretreatment, indicated by U_2 , is illustrated in Figure 1 (party cue \rightarrow policy opinion $\leftarrow U_2$). Crucially, the likelihood of pretreatment is not uniform over people; rather, it will tend to be higher among those who are interested in and follow politics, and, by extension, among those with greater cognitive ability. Thus, the likelihood of pretreatment is confounded with exposure to policy information and engagement in cognitive effort via the common causes described earlier (U_1).

Absent rigorous statistical adjustment for U_1 or U_2 , this means that one could observe either or both of the causal interactions predicted by the heuristic theory *even if neither existed*.

Another causal interaction that plausibly confounds those predicted by the heuristic theory is the interaction between exposure to the party cue and partisan motivation (Bakker, Lelkes, and Malka 2020; Petersen et al. 2013). In particular, people who experience stronger social incentives to toe the party line—partisan motivation (U_2)—are expected to be more strongly influenced by the party cue (party cue \rightarrow policy opinion $\leftarrow U_2$). And, importantly, such incentives are likely to be correlated with the common causes described earlier (U_1). For example, it is plausible that people who are interested in and follow politics tend to be embedded in more politically partisan social networks than people who are less interested in politics; and, thus, face stronger social incentives to follow party cues. Absent rigorous statistical adjustment for U_1 or U_2 , the existence of this causal interaction therefore biases to an unknown extent *against* identifying the causal interactions predicted by the heuristic theory.

Many studies that offer tests of the causal interactions predicted by the heuristic theory using the covariate design do not attempt to adjust for the potential confounding described in this section. Even if they had, it is not clear whether the interactions could be credibly identified owing to the strong and implausible measurement assumptions required. The presence of confounding may in fact help to explain the mixed and inconsistent results across studies. For example, there is large heterogeneity between studies in policy issue stimuli (Guisinger and Saunders 2017; Tappin 2020), and the extent of confounding by variables such as likelihood of pretreatment and partisan motivation plausibly differs across policy issues. (Some issues are more publicly salient, for example, increasing the likelihood of pretreatment.) Thus, depending on the policy issue, estimates of the interaction between the party cue and policy information or cognitive effort could be biased in either direction, and to an unknown extent.

Given the hard identification challenges faced by the covariate-design, studies that randomize exposure to policy information or engagement in cognitive effort can provide more informative tests of the causal interactions predicted by the heuristic theory.

Tests using experimental designs

We are unaware of any studies that randomize engagement in cognitive effort alongside exposure to party cues; but a number of studies do randomize policy information alongside exposure to party cues, and measure policy opinions as an outcome (Agadjanian 2020; Boudreau and MacKenzie 2014; Bullock 2011; Ciuk and Yost 2016; Cohen 2003; Nicholson 2011; Pannico 2020; Peterson 2019).⁴ However, the majority of these studies do not randomize the *extent* of exposure to policy information—for example, by including a no-information control group—but rather only the *type* of information that is available (Bullock 2011; Ciuk and Yost 2016; Cohen 2003; Nicholson 2011; Pannico 2020; Peterson 2019); whether the policy is characterized as consistent with liberal or conservative values, for instance. This design precludes testing the key prediction of the heuristic theory that the influence of party cues will be diminished when substantive policy information is available (versus not), or when the extent of relevant information is increased.

Of the remaining studies that do randomize the extent of substantive policy information alongside party cues, the results are somewhat mixed regarding the key prediction of the heuristic theory. The estimates of Boudreau and MacKenzie (2014) offer qualitative evidence that the influence of party cues is diminished when policy information is available—consistent

⁴ A handful of studies randomize information alongside party cues but the measured outcome variable is candidate support, not policy opinions (Boudreau, Elmendorf, and MacKenzie 2019; Peterson 2017; Rahn 1993; Riggle et al. 1992).

with the prediction. However, they do not report a quantitative test of the causal interaction.⁵ Agadjanian (2020), on the other hand, does report a quantitative test. He finds that exposure to ~120 words describing the potential negative implications of either an infrastructure or trade policy bill does not reliably diminish the influence of the party cue among Republicans or Democrats, respectively. This is inconsistent with the key prediction of the heuristic theory.

Nevertheless, it remains unclear how well this result would generalize to other policy issues in light of the substantial between-issue variation in party cue effects (Tappin 2020). For example, Boudreau and MacKenzie (2014) included nine different policy issues in their study. For some of these issues, their estimates (qualitatively) suggested that the influence of the party cue was diminished in the presence of policy information; for others, it remained similar in size; while for some the influence of the party cue was *larger* under information.⁶ This suggests the possibility of substantial between-issue variation in the causal interactions predicted by the heuristic theory. Accordingly, providing a generalizable and informative test of the theory's predictions demands a much larger sample of policy issues (Yarkoni 2019).

Experiments randomizing effort, information, cues and issues

We conducted two large survey experiments in which we randomly and independently assigned both exposure to policy information and engagement in effortful thinking, as well as party cues, and measured American adults' opinions regarding 20 contemporary U.S. policy issues. The experiments covered a diversity of policy issues, corresponding to the broad policy areas of

⁵ See Figure 1 in their paper. As far as we could tell, they only report quantitative tests of differences between conditions; not quantitative tests of the difference-in-differences between conditions, which are necessary to test the causal interaction.

⁶ This was determined by consulting tables A12 and A13 in their appendix.

immigration, education, healthcare, foreign policy, taxation, the economy, and democratic governance, among others. The policy issues are reported in Table 1.

The basic design of the two experiments was the same. Respondents entering the survey answered a brief demographic questionnaire, before being asked for their opinion regarding five of the policy issues. Respondents gave their opinion on each issue in a sequential order (i.e., answering one issue before moving onto the subsequent survey page for the next issue), and the order of presentation was randomized. The five policy issues seen by each respondent were drawn randomly from a larger set of 10 issues per experiment. On each issue, the respondent was randomly assigned to receive a party cue (versus no cue); substantive policy information (versus no information); and an inducement to effortful thinking about the policy issue (versus no inducement). In other words, we implemented a factorial design with four factors, randomizing the party cue, policy information, inducement to cognitive effort, and policy issue. Policy opinions were always given on a seven-point Likert scale.

The set of 10 policy issues in experiment 1 was drawn from the battery of policy questions used in Phase 1 of the *Nationscope* surveys (Democracy Fund and UCLA 2020), largescale public opinion surveys of Americans' political attitudes that began in 2019. The set of 10 issues in experiment 2 was drawn from the website *isidewith.com*, an online encyclopedia that provides Americans with information about the policy positions of dozens of U.S. political candidates and leaders. The encyclopedia professes to be non-partisan and unaffiliated with any candidate, party, corporate sponsors, investors, or interest groups. The two sets of policy issues were selected on the basis of pilot studies and prior work (Tappin 2020) to balance three criteria: a detectable influence of the party cue, the contemporary relevance of the policy issue, and to maximize the diversity of policy issues included in the experiments reported here. The selection strategy is reported in detail in the Appendix.

Table 1. Policy issues in the two experiments.

Issue label	Issue text	Experiment
abortion_insurance	Allow employers to decline coverage of abortions in insurance plans	1
abortion_waiting	Require a waiting period and ultrasound before an abortion can be obtained	1
china_tariffs	Impose trade tariffs on Chinese goods	1
estate_tax	Eliminate the estate tax	1
health_subsidies	Subsidize health insurance for lower income people not receiving Medicare or Medicaid	1
immigration_system	Shift from a more family-based to a more merit-based immigration system	1
immigration_wire	Require proof of citizenship or legal residence to wire money to another country from the USA	1
public_option	Provide the option to purchase government-run insurance to all Americans	1
trans_military	Allow transgender people to serve in the military	1
vouchers	Provide tax-funded vouchers to be used for private or religious schools	1
bank_audit	Should the Federal Reserve Bank be audited by Congress?	2
donations	Should corporations, unions and non-profit organizations be allowed to donate to political parties?	2
foreign_aid	Should the U.S. increase or decrease foreign aid spending?	2
lobbyist_ban	Should there be a 5-year ban on White House and Congressional officials from becoming lobbyists after they leave the government?	2
nato	Should the U.S. remain in NATO?	2
pensions	Should pension payments be increased for retired government workers?	2
single_payer	Do you support a single-payer healthcare system?	2
surveillance	Should local police increase surveillance and patrol of Muslim neighborhoods?	2
tax_breaks	Should the President offer tax breaks to individual companies to keep jobs in the U.S.?	2
work_visas	Should the U.S. increase or decrease the amount of temporary work visas given to high-skilled immigrant workers?	2

The party cue treatments consisted in respondents being informed of the positions of other party voters (experiment 1, hereafter the party-*voter* cue), or Barack Obama and Donald Trump (experiment 2, the party-*elite* cue) on the policy issue in question.

The party-*voter* cue treatment consisted in respondents learning the distribution of policy support among a subset of Republican and Democratic Party voters who answered the Nationscape surveys in 2019. In particular, respondents learned the policy position adopted by the majority of the most politically interested, knowledgeable and partisan voters from each party. The 10 issues were chosen in such a way that the majority positions of these Republican and Democratic Party voters were always in opposition. That is, if a majority of Republican voters supported the policy then a majority of Democrat voters opposed the policy, and vice versa. The party-*elite* cue treatment consisted in respondents learning the policy positions of Barack Obama and Donald Trump, according to their public statements and voting record, documented by *isidewith.com*. Again, we chose the 10 issues such that their positions were always in opposition. The delivery of the party cue treatments is illustrated in the Appendix.

The policy information treatment consisted in respondents being exposed to a paragraph of background information relevant to the policy issue, as well as arguments for and against implementing the policy. The information paragraphs were typically between 120 and 140 words in length, and were written from scratch using media and other sources for the purposes of the experiment. They are reported in full in the Appendix. To illustrate, however, this was the information treatment for the *foreign_aid* issue:

U.S. foreign aid is aid given by America to other countries to support global peace, security, and development efforts, and provide humanitarian relief during times of crisis. In 2017, U.S. foreign aid allocations totalled \$48 billion (about 1.2% of total government spending). This came out to roughly \$144 per American for the year.

Supporters of decreasing foreign aid spending argue that foreign aid can weaken the accountability of foreign governments to their people – while receiving aid, those governments have less reliance on obtaining taxes from their populaces, and so less incentive to improve their economic wellbeing. Supporters of increasing foreign aid spending argue that it gives the U.S. considerable international influence (soft power) and that it is morally right to provide financial assistance to countries who are far less wealthy.

The cognitive effort treatment consisted in respondents being asked to explain in their own words their reasons for holding the opinion that they did about the policy issue. They were given an open-ended response box in which to write their explanation. Respondents were informed that their explanations were not anonymous, but might be read by another person to be recruited in a later survey, and that this person could be asked to evaluate the quality of the respondent's reasoning. The logic of this treatment derives from a large body of work on the effects of accountability on human reasoning. In reviewing this work, Lerner and Tetlock (1999) identify that unbiased, effortful thinking tends to increase when one is accountable to an audience whose views are unknown, but who are perceived to be interested in evaluating *accuracy* and *process*—rather than outcomes (e.g., see especially p. 259 and p. 263).

Unlike the party cue and policy information treatments, the target of randomization in this treatment (cognitive effort) is an unobserved psychological variable, and can therefore only be encouraged rather than directly manipulated. Thus, it is particularly important to assess and observe compliance with this treatment to ensure a diagnostic test of the causal interaction predicted by the heuristic theory. As we report below, there appeared to be very high levels of compliance with the cognitive effort treatment in our experiments.

Respondents were U.S. adults recruited via the survey provider *Prolific*.⁷ Samples recruited using this provider are not representative of the general U.S. population—for example, they tend to be more Democrat, among other demographic differences. (Sample characteristics are reported in the Appendix.) However, this does not significantly affect the ability of our design to test the causal interactions predicted by the heuristic theory, for two reasons. First, the estimates relevant for testing said interactions are all causal quantities obtained via randomization in our design. Recent work indicates that causal estimates obtained in convenience samples of U.S. adults track well with those obtained in national samples of U.S. adults (Coppock 2019; Mullinix et al. 2015). Second, there are no obvious scope conditions in the heuristic theory that should lead us to expect that only a national/probability (versus typical convenience) sample of U.S. adults would conform to the predictions of the theory. Rather, the predictions are expected to apply broadly to those who take surveys and to those who do not. In other words, samples of respondents recruited via *Prolific* are fit-for-purpose with respect to the aims of our research design (Coppock and McClellan 2019).

We aimed to recruit 3,000 respondents in each of our two experiments, equating to 15,000 respondent-issue observations in each experiment, and, therefore, an expected total of 6,000 respondents and 30,000 observations across the two experiments. To mitigate any risk of differential attrition by treatment biasing our estimates, we implemented a planned double-sampling strategy; re-contacting all respondents who dropped out of the experiments partway through the policy rating segment, regardless of treatment assignment, to obtain their missing outcomes (Coppock et al. 2017; Gomila and Clark 2020). Dropout was low in absolute terms: across the two experiments, a total of 6,288 respondents started the survey, and 148 (~2%) of these dropped out during the policy rating segment. Furthermore, double-sampling was highly

⁷ <https://www.prolific.co/>

effective: within one week of the initial survey being fielded, we recovered the missing outcomes of 145 (~98%) of those who dropped out. Accordingly, there is negligible risk of bias in our estimates caused by differential attrition.⁸ Experiment 1 was fielded July 16th, 2020 and Experiment 2 was fielded November 12th, 2020.

Results

Compliance with the effort and information treatments

We begin by assessing compliance with our cognitive effort and policy information treatments. This is important for determining the diagnosticity of our tests of the causal interactions predicted by the heuristic theory. High compliance equals more diagnostic tests; low compliance equals less diagnostic tests. We examine two outcome variables to assess compliance. The first variable is the time taken (in seconds) for the respondent to “submit” their answer to the policy question shown on the survey page. To the extent that respondents are reading the policy information and engaging in effortful thinking about the policy issue, the time-to-submit should be greater for policy questions assigned to the information and effort treatments, respectively. The second variable is the number of words written in the open-ended response box provided in the cognitive effort treatment. To the extent that respondents are complying with the encouragement to explain the basis of their policy opinion, the average number of words written should be considerably greater than zero.

⁸ In addition, one respondent (of the 148) who dropped out did so prior to the loading of the treatment on their survey page. Consequently, their missing outcome cannot have been caused by attrition due to treatment assignment (and so we did not attempt to re-contact this respondent). This reduces the number of respondents classed as potential differential dropouts to just two (i.e., 147 minus 145).

Figure 2A shows the mean time-to-submit in each condition, and Figure 2B shows the mean number of words written in the conditions that included the cognitive effort treatment. These estimates are the population-level condition means from Bayesian multilevel models fitted to the observation-level data from both experiments. The outcome variables in both models were log-transformed prior to fitting due to skew, and the estimates in Figure 2 are the exponents of the condition means. The models allow the condition means to vary across the 20 different policy issues, thus properly accounting for the between-issue variation in condition means.⁹ The prior distributions on the model parameters are vague and weakly-informative, allowing the data to “speak for itself.” The models were fitted using the package *brms* (Bürkner 2017) in the statistical software R, and are written formally in the Appendix along with greater detail regarding model specification and diagnostics.

The estimates in Figure 2 indicate high levels of compliance with both the policy information and cognitive effort treatments. As per Figure 2A, the mean time-to-submit is approximately three times larger in the condition where (only) policy information was

⁹ We omitted respondent- and experiment-varying condition means from these models to facilitate estimation. As a function of our design, respondents had five or fewer observations each. Thus, allowing the (eight) condition means to vary across respondents translates to estimating many more parameters than there are data points to fit the model, which can cause problems for identifiability. Similarly, because there are only two experiments, there are only two observations with which to estimate the between-experiment variance parameter—again, causing problems for identifiability of the model (in the absence of setting an arbitrarily precise prior on the parameter). In any case, the impact of omitting experiment-varying effects is likely negligible here, because the main difference between the experiments is the sample of policy issues (and their respective cues and information), which *are* modelled as varying effects.

provided, compared with the control condition. This is compelling evidence that, on average, respondents were reading the information. Similarly, the mean time-to-submit is dramatically larger in conditions that were assigned the cognitive effort treatment—approximately ten times larger than the control condition. This is compelling evidence that, on average, respondents were complying with the encouragement to explain the basis of their policy opinion, and engaging in cognitive effort while doing so. This conclusion is corroborated by the estimates in Figure 2B, which show that respondents wrote between thirty-three and thirty-nine words on average when prompted to explain the basis of their policy opinion. To illustrate what an explanation with these many words looks like, below is a randomly sampled response to the cognitive effort treatment that was between thirty-three and thirty-nine words in length:

I support a single-payer healthcare system because healthcare is a human right that everyone in the country should have equal access to. Private health insurance creates income inequalities, giving people with more money better access to healthcare.

This response is thirty-seven words; many were shorter than this (46%), and many were longer (53%). We report additional examples in the Appendix.

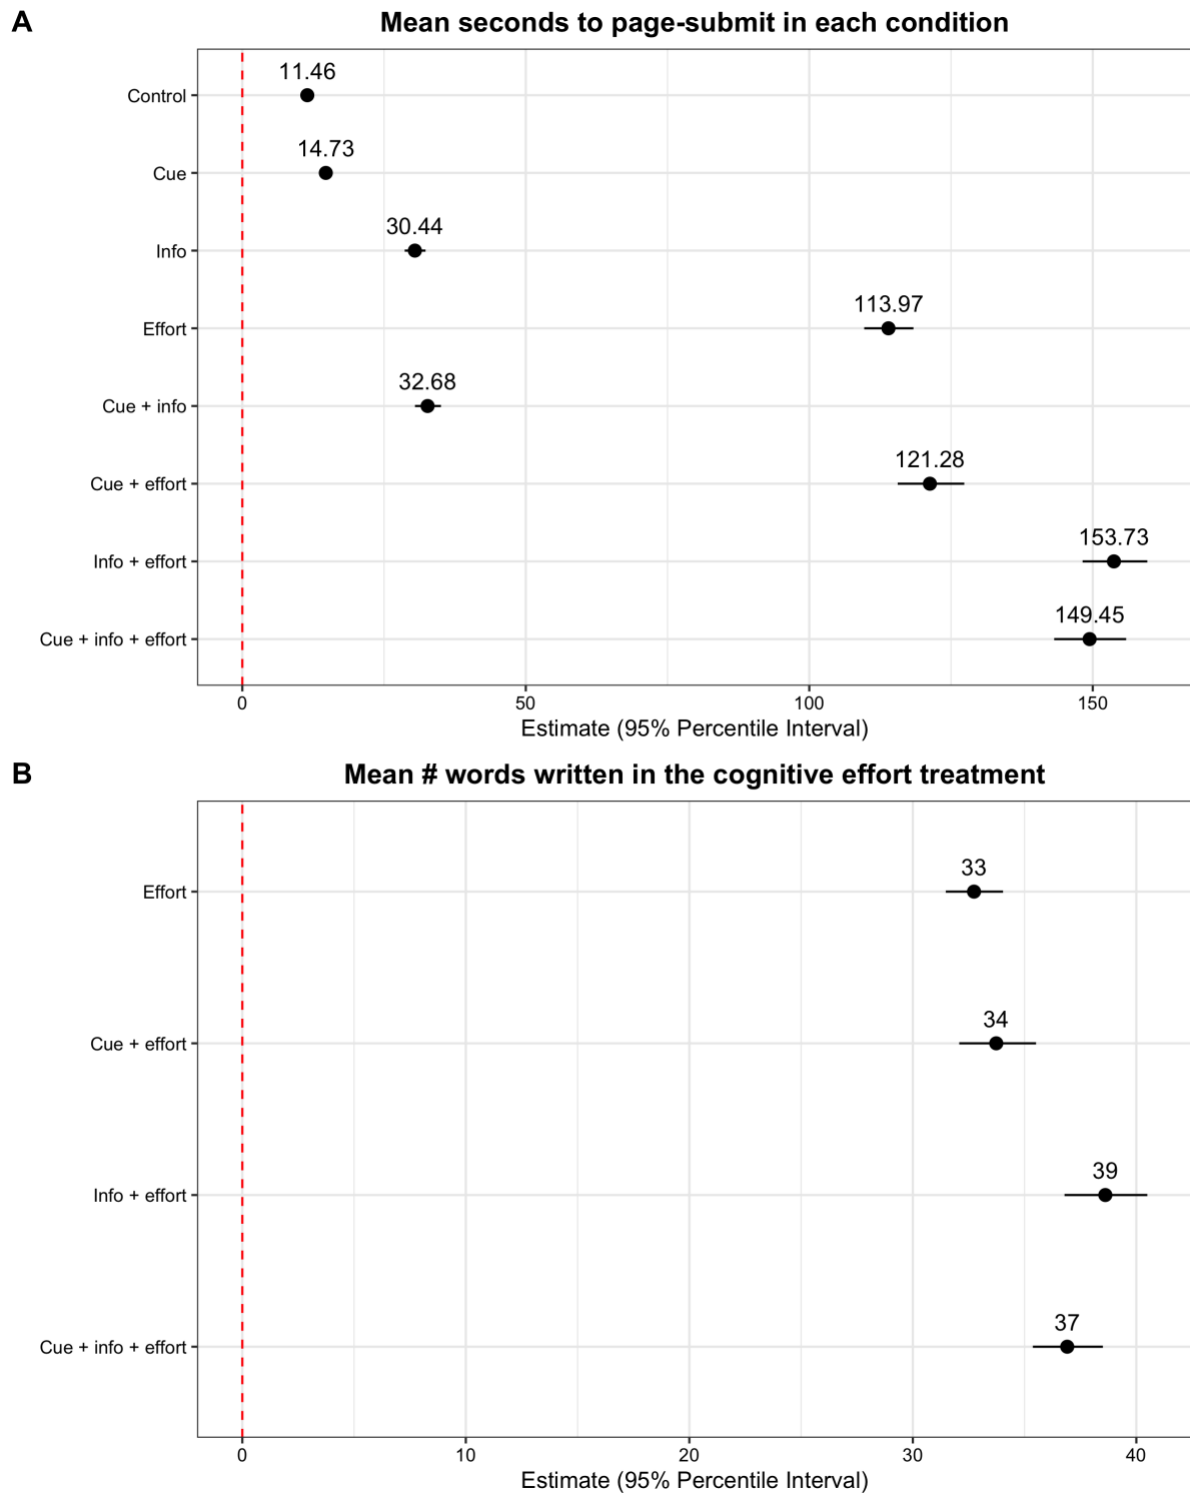


Figure 2. Results from the analysis of treatment compliance.

Aggregate results

Here we test the key prediction that the influence of party cues is diminished when people possess more information about policy, a greater propensity to engage in effortful thinking about policy, or both. The outcome variable is opinion toward the policy in question, given on a seven-point Likert scale, and coded such that higher scores indicate agreement with the in-party cue. For example, a Democrat respondent who strongly supported the *tax_breaks* policy (see Table 1)—originally a score of 7 on the scale—would receive a score of 1, since the in-party cue (Obama’s position) is to oppose this policy. In contrast, a score of 7 from a Republican respondent would retain its original coding, since the in-party cue (Trump’s position) is to support the policy. (And vice versa for issues where the cue direction is reversed.) Accordingly, the expected sign of the party cue treatment effect is positive for all partisan respondents and issues (political Independents are not included in the recoding scheme, nor in the analysis). We thus ignore respondent’s party identification in our analysis.

Figure 3A shows the distribution of policy opinions in each condition, with the raw mean (purple) and model-estimated mean (red) overlaid. The model-estimated means are from a Bayesian multilevel model fitted to the observation-level data from both experiments. The model includes dummy variables for each treatment (party cue, policy information, cognitive effort), and parameters for all interaction terms. The model allows all eight parameters to vary across policy issues as well as respondents.¹⁰ Thus, the estimates properly account for variation across policy issues and respondents (Yarkoni 2019). This is evident in Figure 3A: the 0.95 intervals on the raw means are so small that they cannot even be seen, whereas the intervals on the model-estimated means are wider because they account for the issue-level and respondent-

¹⁰ Unlike the treatment compliance models, this model was successfully identified whilst allowing all eight parameters to vary across respondents (as well as policy issues).

level clustering in the data. The prior distributions on the model parameters are vague and weakly-informative, allowing the data to speak for itself. Formal model specification and diagnostics are reported in the Appendix. The model is fitted on a total of 29,099 observations from 5,905 respondents who identified with either the Democratic or Republican Party.

Figure 3B shows the average treatment effect (ATE) of the party cue at baseline; that is, in the *absence* of the policy information and cognitive effort treatments (top row), as well as under the various combinations of information and effort treatments (rows 2-4). The baseline ATE is about half a scale point, and is bounded precisely and clearly to the right of zero; as expected, party cues reliably influenced people's policy opinions on average (Bullock 2020). Furthermore, qualitatively consistent with the key prediction of the heuristic theory, the party cue ATE is smaller under every combination of policy information and cognitive effort treatment (rows 2-4) compared with baseline (top row). In other words, the influence of the party cue was attenuated when people were exposed to substantive information about the policy, induced to engage in effortful thinking about the policy, or both. In terms of magnitude, when substantive policy information is available, the party cue ATE attenuates by ~35% ($1 - 0.32/0.49$) on average; when cognitive effort is induced, it attenuates by ~8% ($1 - 0.45/0.49$); and under both policy information and cognitive effort it attenuates by ~39% ($1 - 0.30/0.49$).

Figure 3C shows the key quantitative tests of the paper: whether these attenuation magnitudes are convincingly larger than zero. As indicated by the estimate in the top row, we can conclude with >95% probability that the attenuation of party cue influence is larger than zero when substantive policy information is available, consistent with the prediction of the heuristic theory. We can draw a similar conclusion when substantive policy information is available *and* effortful thinking about the policy is induced (third row). In contrast, however, there is little evidence that cognitive effort by itself attenuates the influence of the party cue: the data and model are quite compatible with the true decrease in party cue influence under

cognitive effort being zero (middle row). Furthermore, the results also imply that inducing cognitive effort in the presence of policy information does almost nothing to attenuate the influence of the party cue beyond exposure to the policy information per se; the point-estimates in the first and third rows are almost identical.

The aforementioned results are based on a model that assumes the outcome variable is metric, when in fact it is ordinal—specifically, a seven-point Likert scale. Furthermore, the raw data are right-skewed (Figure 3A), implying there may be a mild ceiling effect on the outcome scale in some conditions. Either or both of these factors can generate spurious interactions that fail to inform theory (Bürkner and Vuorre 2019; Rohrer and Arslan 2020). We thus additionally fitted both ordered-logistic and right-censored versions of the model described in this section, whose diagnostics and results are reported in the Appendix. The results of these models are substantively identical to those we reported in this section.

Taken together, the aggregate results offer qualified support for the causal interactions predicted by the heuristic theory of party cue influence: exposure to substantive information about policy causally attenuated the influence of party cues. The two qualifiers are, first, that we find little evidence for the causal role of cognitive effort in reducing the influence of party cues—either by itself or when paired with exposure to policy information. Second, the estimated attenuation in party cue influence under policy information is not large: between one-third and two-fifths of the total influence. In other words, despite the confidence we can have that the attenuation is greater than zero, the majority of the party cue influence remains intact, and is still bound convincingly to the right of zero (Figure 3B). In the next section, we explore heterogeneity in the causal interaction estimates across the 20 different policy issues.

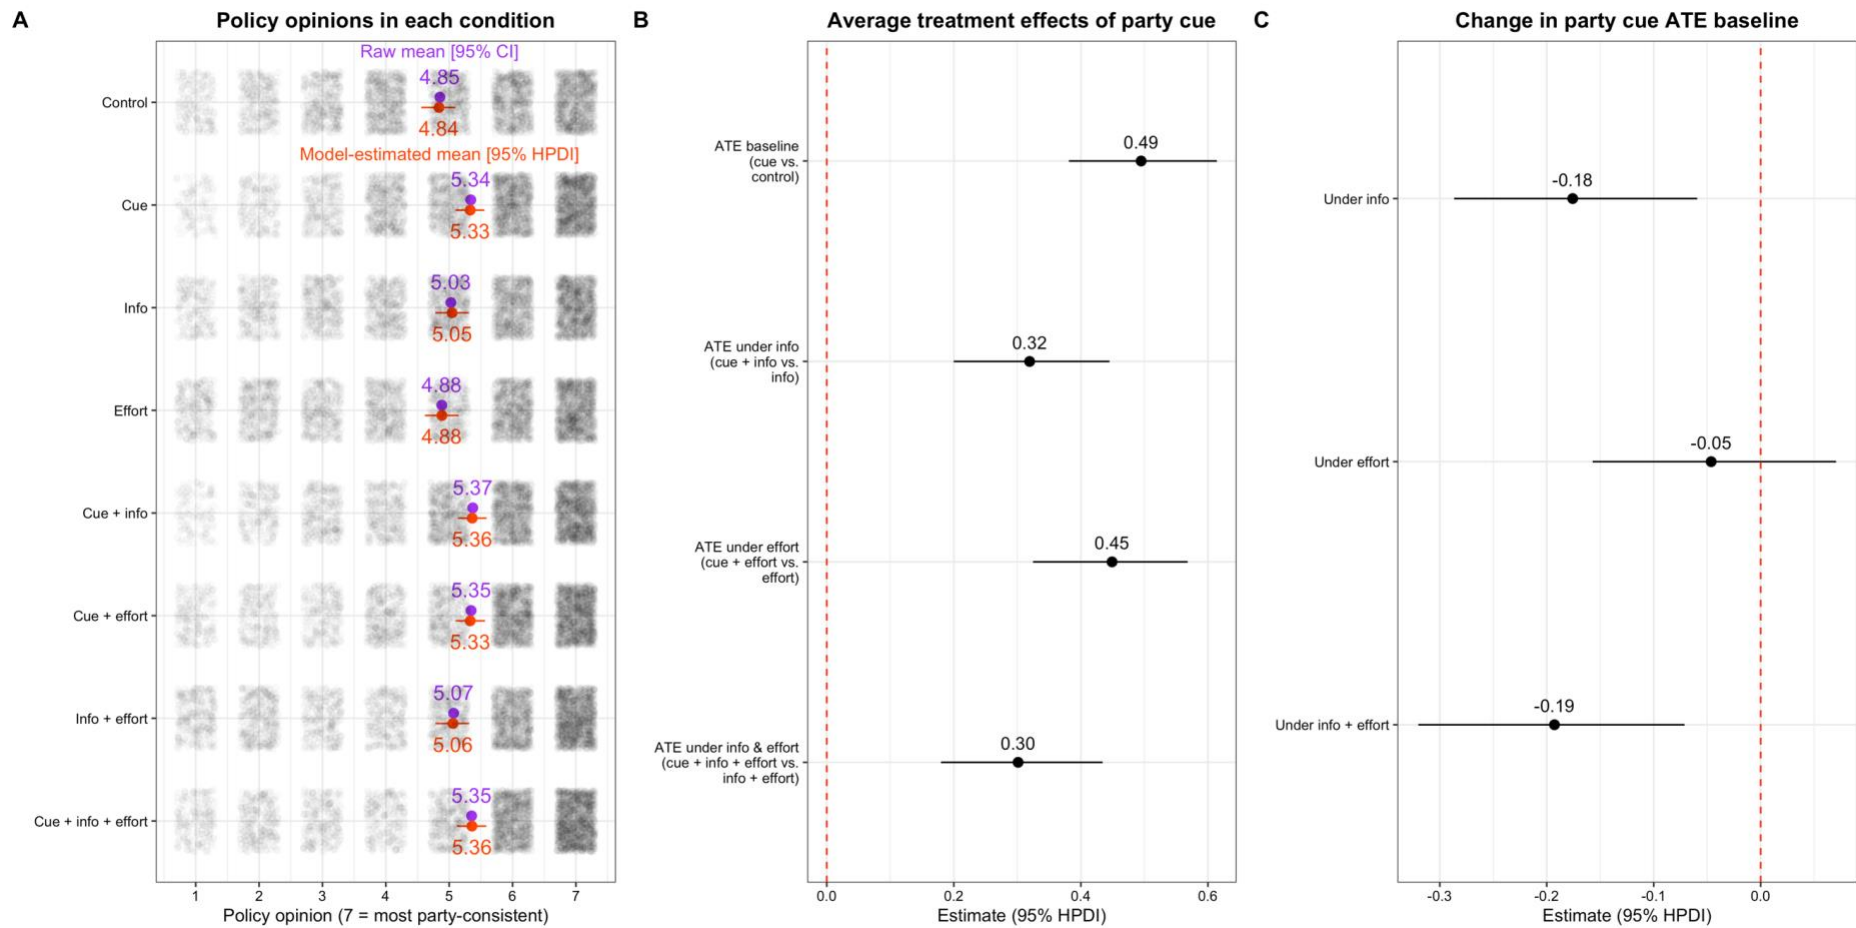


Figure 3. Aggregate results from the model testing the key prediction. HPDI = highest posterior density interval.

Heterogeneity by policy issue

Figure 4 shows the party cue ATEs for each policy issue in the sample. Recall that the issue questions are listed in Table 1. The estimates are computed from the full posterior distribution of the multilevel model fitted to obtain the aggregate results (described in the previous section), not the result of fitting separate models for every issue. This has the benefit of leveraging partial pooling to minimize the risk of overfitting the data for individual issues, thereby improving the out-of-sample predictive accuracy of the estimates on average (McElreath 2020). (In the Appendix, we report the distribution of policy opinions in each condition for each policy issue.)

Upon initial inspection, the estimates in Figure 4 suggest rather limited heterogeneity across policy issues in the predictions of the heuristic theory. Consistent with the aggregate results reported in the previous section, for all policy issues in our sample the influence of the party cue is qualitatively diminished when policy information is available, cognitive effort induced, or both (top row versus rows 2-4 in each panel). Furthermore, tracking across the panels, most of the attenuation magnitudes seem to fall in a similar range: approximately ~ 0.20 of a scale point decrease under policy information, and ~ 0.05 of a scale point decrease under cognitive effort. However, one can also consider these decreases with respect to the size of the baseline party cue ATE (top row), rather than as absolute movement on the opinion scale. In that case, there is somewhat more evidence of heterogeneity.

To illustrate this, Figure 5 shows the change in the baseline party cue ATE (y axis) under the different combinations of information and effort treatments (panels) as a function of the baseline ATE (x axis). Thus, estimates falling along the horizontal dashed black line indicate that the baseline party cue ATE did not attenuate under information or effort (or both) on the labelled issue; estimates along the diagonal sky blue line indicate 50% attenuation; and estimates along the diagonal deep blue line indicate 100% attenuation. The estimates are colored red if the 0.95 interval for the estimated change excludes zero, and black otherwise.

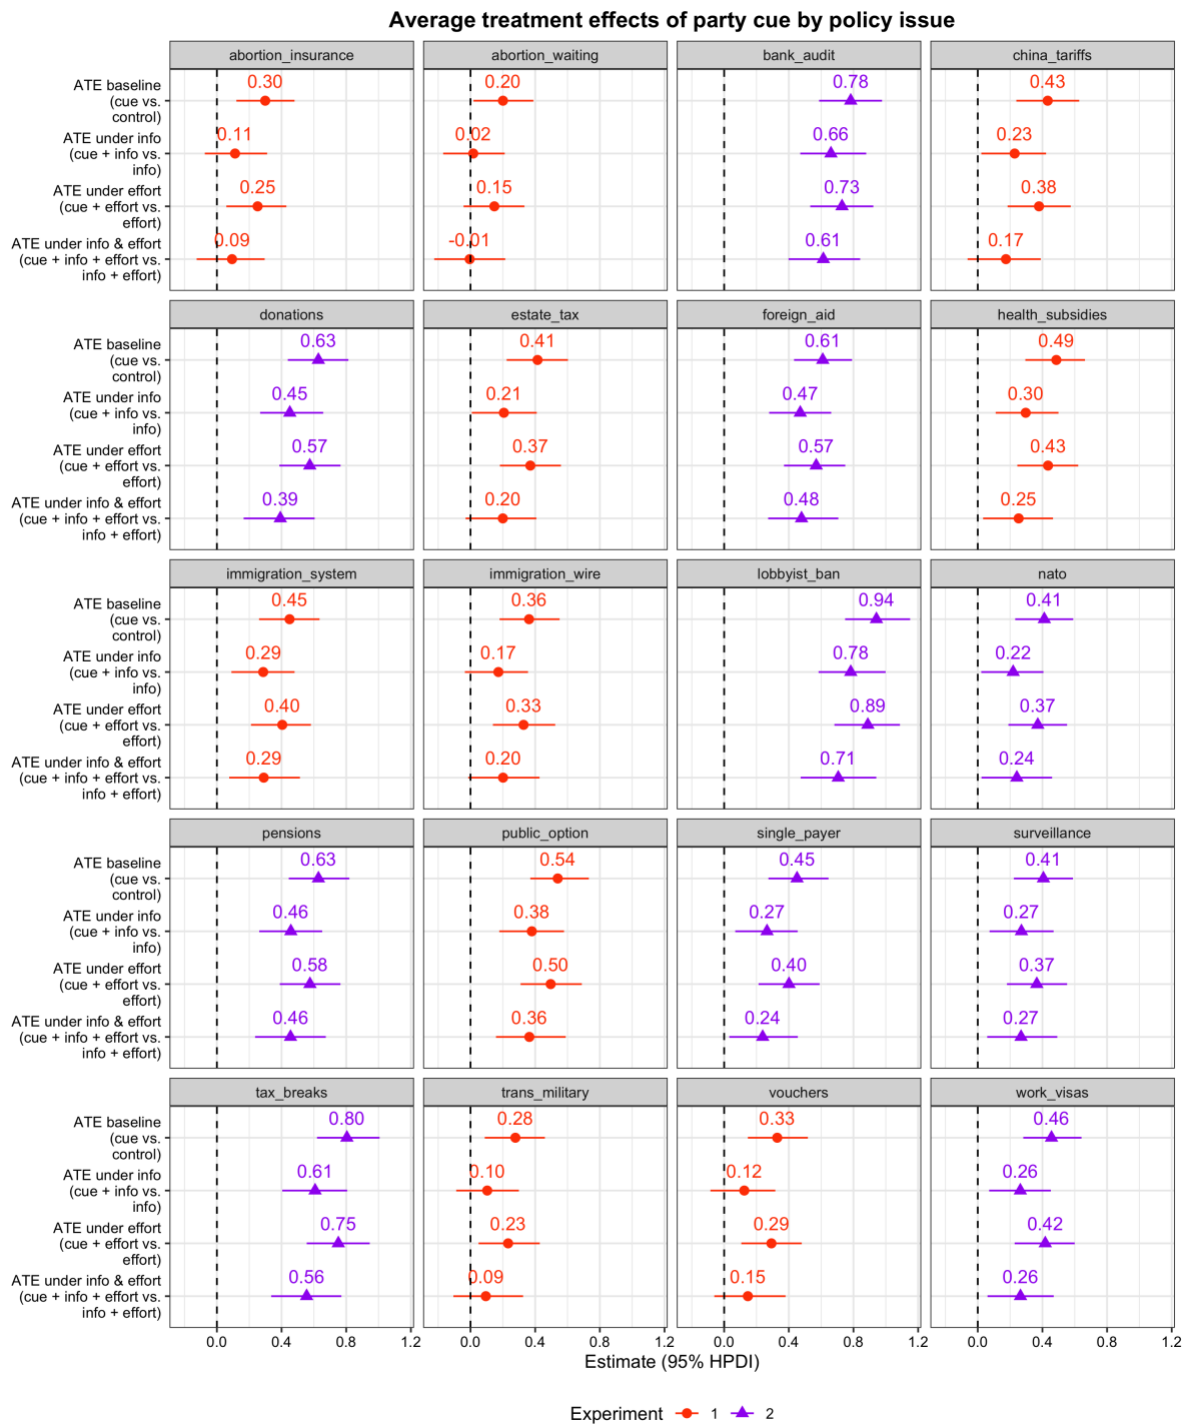


Figure 4. Average treatment effects of the party cue by policy issue.

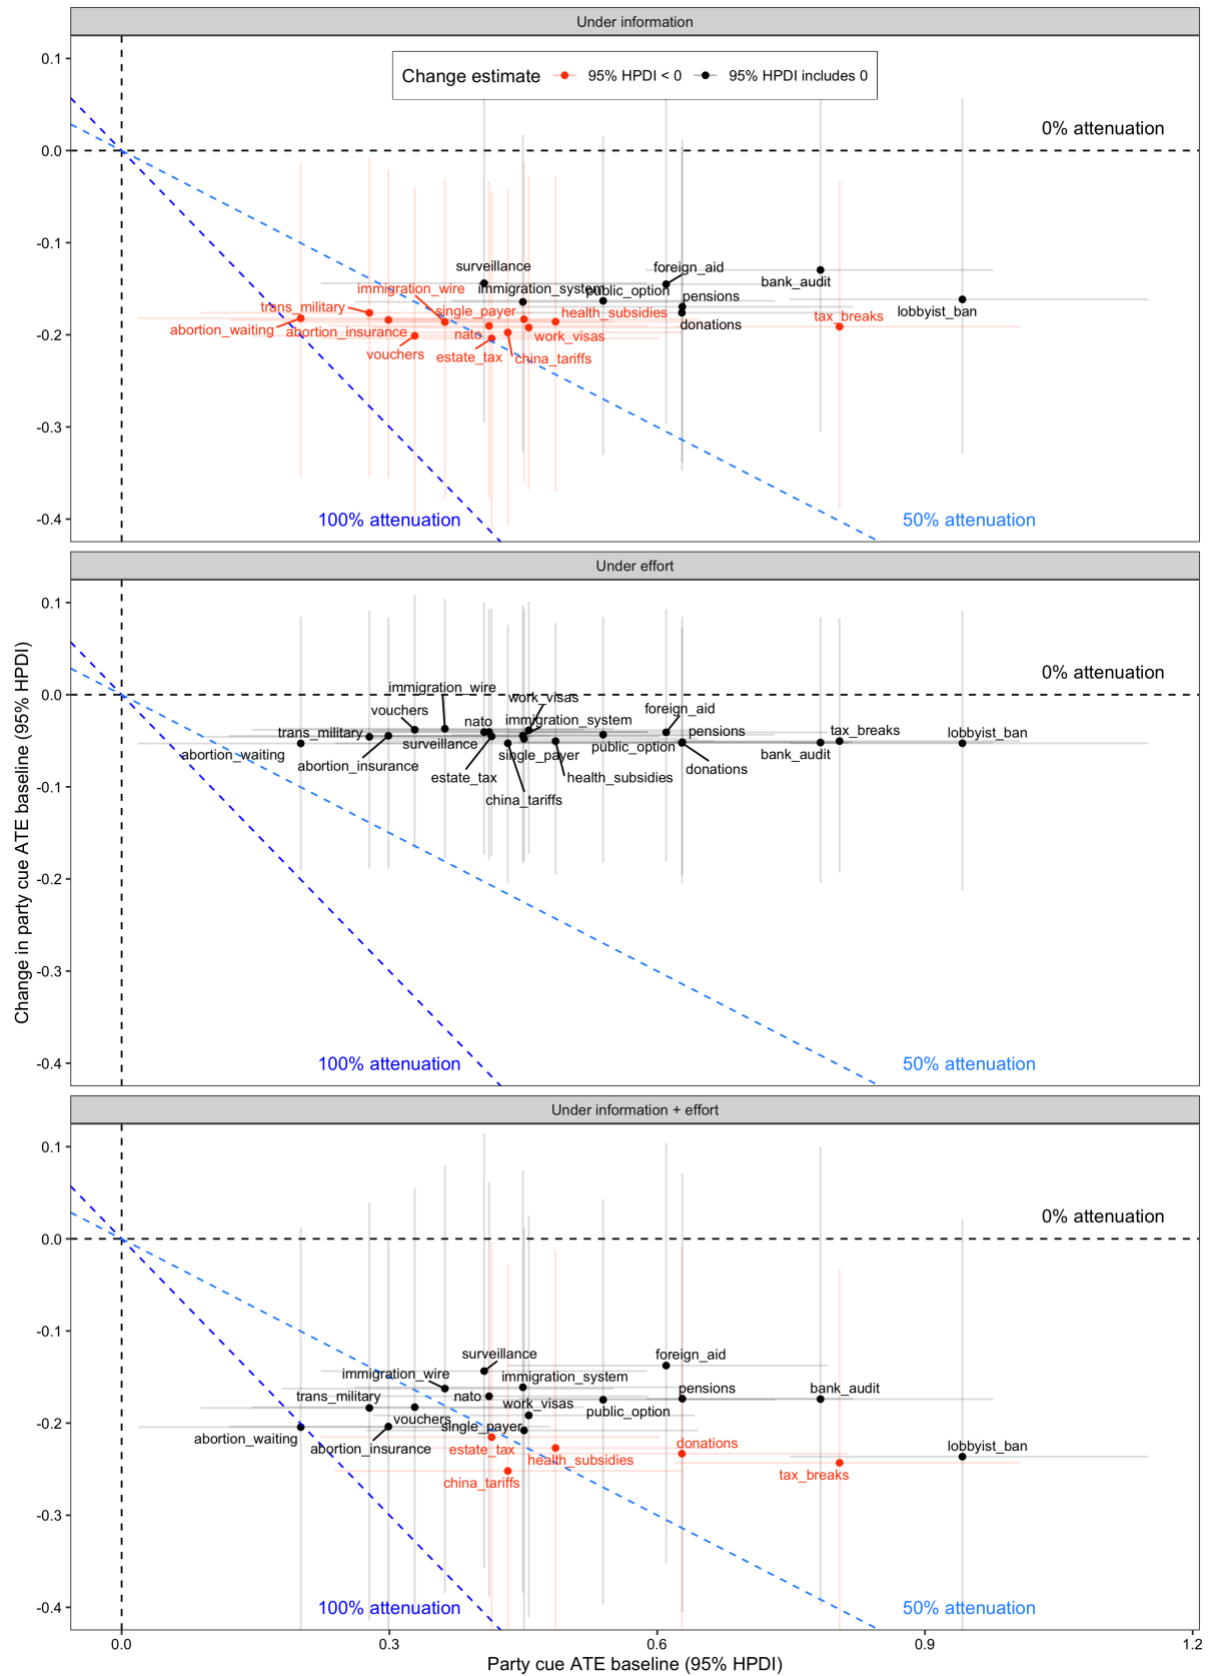


Figure 5. Change in the baseline party cue ATE as a function of the baseline ATE.

We draw attention to three patterns of interest revealed by Figure 5. First, the between-issue heterogeneity in the attenuation of the baseline ATE is largest under policy information and cognitive effort combined (bottom panel) than under either alone (top and middle); shown by the variation of the estimates with respect to the y axis. Second, across all three panels, the variation of the estimates with respect to the x axis is much greater than with respect to the y axis. Thus, between-issue heterogeneity in the baseline party cue ATE is much greater than the heterogeneity in the *attenuation* of the baseline ATE under information, effort, or both. At least, that is, when heterogeneity is assessed via absolute movement on the opinion scale.

However, as mentioned above, heterogeneity in the attenuations of the baseline party cue ATE can also be assessed with respect to size of the baseline ATE itself—for example, as a proportion of the baseline. Assessed this way, there is somewhat more compelling evidence of heterogeneity across policy issues. To illustrate, consider the three issues with the largest and smallest baseline ATEs: *bank_audit*, *tax_breaks* and *lobbyist_ban*, and *abortion_waiting*, *trans_military* and *abortion_insurance*, respectively. For those issues with the largest baseline ATEs, exposure to substantive policy information (top and bottom panels) attenuates the ATE by approximately 25-30%, a clear minority of the total. In contrast, for those issues with the smallest baseline ATEs, the estimated attenuations are in the region of 70-100% of the baseline ATE; a clear majority of the total. Observed in isolation, these different relative magnitudes of attenuation could paint a very different picture of support for the heuristic theory.

Such between-issue heterogeneity thus illustrates the broad importance of studying more than a handful of policy issues when developing and testing theories of party cue influence (Guisinger and Saunders 2017; Tappin 2020). The number of policy issues examined in the typical study is small—between one and four issues, and rarely more than that. It is easy to imagine different studies arriving at quite different conclusions regarding the extent to which, say, exposure to substantive policy information attenuates the influence of party cues;

all because they studied a different handful of the issues in Figure 5. By constraining empirical generalizability, studying a small and idiosyncratic set of policy issues can therefore foment inconsistencies in the published literature and stall theoretical progress (Yarkoni 2019).

Nevertheless, we close this section by highlighting that, under every combination of information and effort treatments (all panels), the majority of issues fall above the 50% attenuation line. Thus, reinforcing the results of the aggregate analysis, for most policy issues in our sample the majority of party cue influence remains intact when substantive policy information is available, effortful thinking about the issue is induced, or both. In the next section, we explore heterogeneity in the causal interaction estimates across respondents.

Heterogeneity by respondent characteristics

We focus on two demographic characteristics of our respondents: their educational attainment, and self-reported interest in politics. We focus on these characteristics due to their potential relevance for the causal interactions predicted by the heuristic theory.

In particular, we might expect that individuals with higher education or who are more interested in politics are more likely to be “treated” by the policy information and cognitive effort treatments. This could be the case if they possess a greater willingness and/or ability to consume political information and think about policy issues compared with their less educated and interested counterparts. Thus, the party cue influence could attenuate more strongly among the highly educated and interested. On the other hand, higher education and interest are likely to be associated with greater prior exposure to relevant information and deliberation over the policy issues in our sample. Thus, the marginal impact of the information and effort treatments delivered in our study may be lower for these individuals; meaning the influence of the party cue could attenuate *less* strongly among the highly educated and interested.

To explore these distinct possibilities, we fitted two additional versions of the model fitted to obtain the aggregate results, introducing a new dummy variable to indicate whether the respondent had (i) graduated college (with a Bachelor’s degree) or (ii) \geq median political interest ($\geq 5/7$ on a political interest Likert scale collected pretreatment), respectively. In each model, these dummy variables were interacted with the previous three dummies (and their interactions) that indicated the party cue, information and effort treatments. Formal model specifications and diagnostics are reported in the Appendix. Figure 6 shows the results.

Looking first at the education panels (A and B), among individuals with a Bachelor’s degree there appears to be slightly stronger attenuation of the party cue influence under policy information, cognitive effort, and both combined—consistent with this subgroup “complying” more with the information and effort treatments. However, the estimates among individuals who do not hold a Bachelor’s degree are not dramatically different. And, indeed, looking at the pairs of estimates in panel B, the posterior distribution of the difference between each of these pairs of estimates overlaps comfortably with zero (for brevity not shown here, see Appendix 5.4.3). This indicates that there is plausibly zero difference in attenuations between the subgroups. Interestingly, the picture is much the same among subgroups of individuals with higher versus lower interest in politics (panels C and D).

Thus, there appears to be limited evidence of heterogeneity in the causal interactions across subgroups of educational attainment and political interest, despite the fact we might expect these subgroups to be meaningful sources of heterogeneity here. The lack of evidence of heterogeneity along these dimensions has a noteworthy implication: it helps alleviate the concern that our aggregate results would not generalize well from our sample of respondents, given that it skews more educated and politically interested than the general U.S. population. The limited heterogeneity lends confidence to the person-level generalizability of our results.

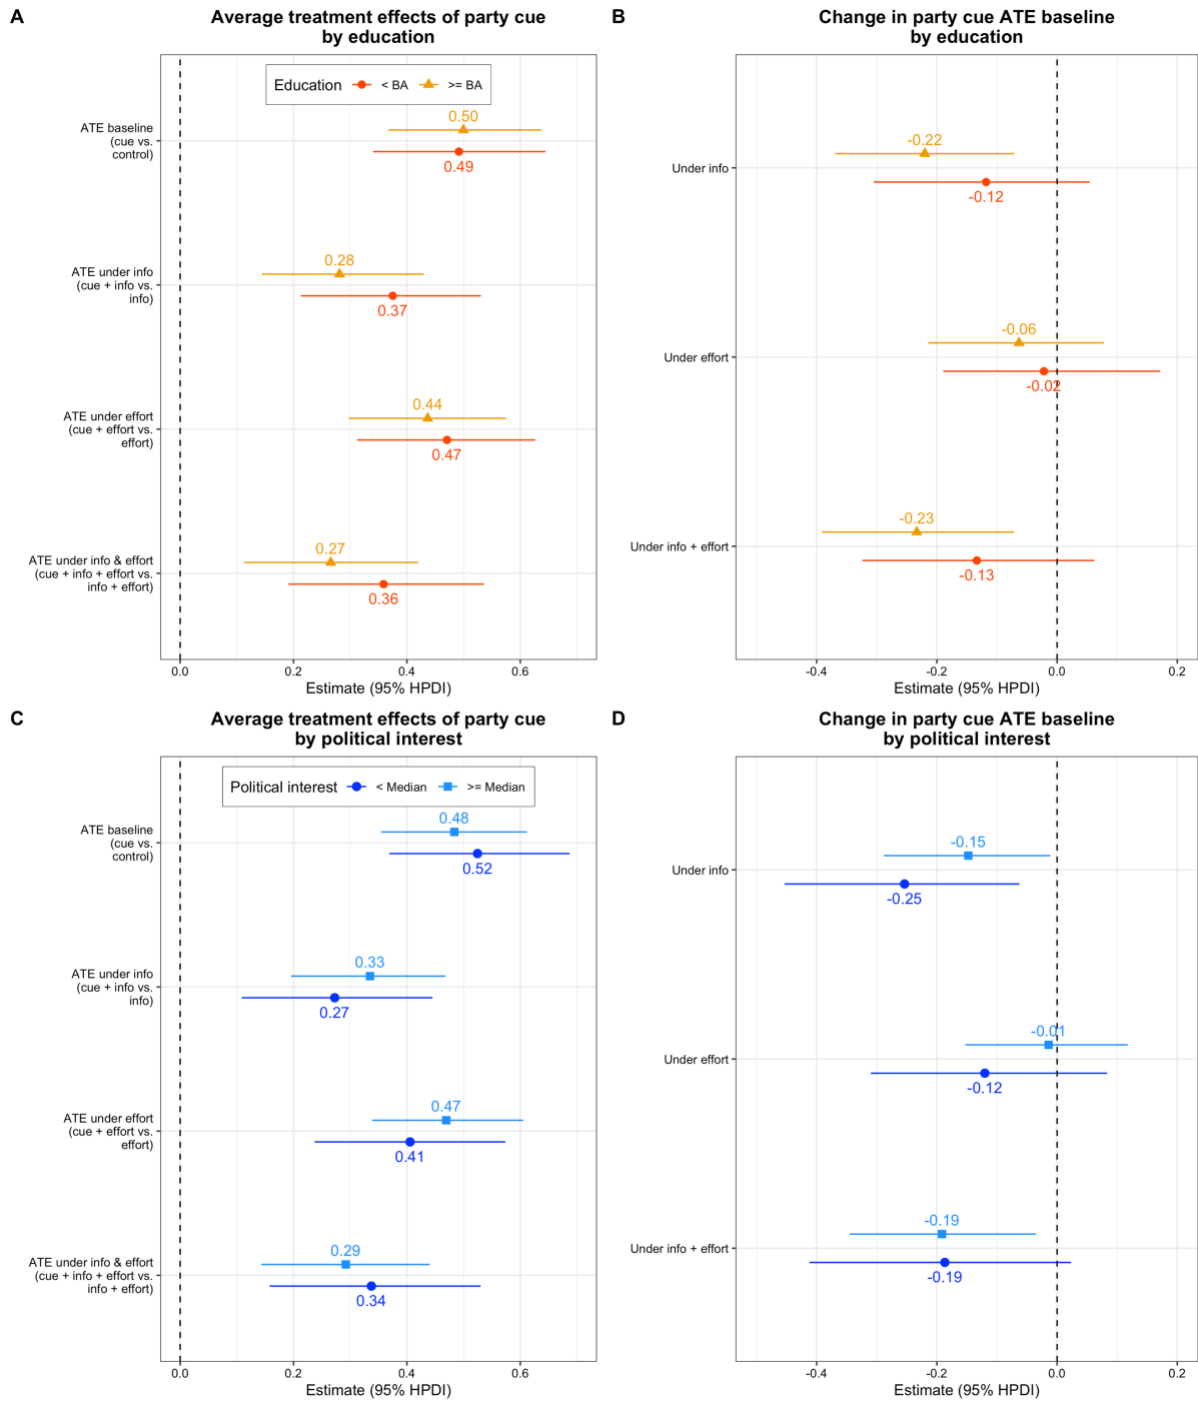


Figure 6. Heterogeneity in the aggregate results by respondent education and political interest.

Party cues cause people to behave as if better informed

We find that the influence of party cues is causally attenuated by exposure to substantive policy information. This is consistent with the heuristic theory, in particular with the notion that the party cue is used as an information “shortcut”. However, a key additional assumption of the information shortcut notion is that cue-following allows voters to form policy opinions that approximate those they would have formed were they more knowledgeable of the policy details (Hobolt 2006; Lupia 1994; Lupia and McCubbins 1998; Mondak 1993). In other words, the assumption is that the party cue is a *valid* information shortcut.

Our data offer a unique and powerful opportunity to interrogate this key assumption. With respect to our data, the assumption predicts that exposure to *either* the party cue *or* substantive policy information has a similar influence on people’s policy opinions, on average. That is, exposure to the party cue should cause people to form opinions *as-if* they were made better informed about the policy details. This isn’t entailed by our earlier results: exposure to policy information could attenuate the influence of party cues without moving opinions in the same direction, or even while moving them in the opposite direction. Previous influential studies testing the as-if assumption have drawn on observational data only, and have focused on a very small sample of policy issues (e.g., Hobolt 2006; Lupia 1994). Threats to causal inference and issue-level generalizability thus loom large. In contrast, our design is fully experimental and incorporates a large sample of distinct policy issues, obviating these threats.

To interrogate this assumption, we compare the average treatment effects of our party cue and policy information treatments. Before doing so, we distinguish between two patterns that would represent strong and weak corroborations of the as-if assumption, respectively. Strong corroboration is inferred if both cue and information ATEs have the same sign and their magnitudes are *not* distinguishable from each other with 95% probability. Weak corroboration is inferred if the ATEs have the same sign and are both distinguishable from zero with 95%

probability. Evidence of strong corroboration would be very surprising, given that the party cue treatments have a clear directional element, whereas the policy information treatments were balanced in content and brief—containing only descriptive information and a single argument for and against the policy. Evidence *inconsistent* with the as-if assumption would be represented by a policy information ATE that is either null or negatively-signed (with 95% probability) *and* that is distinguishable from the party cue ATE with 95% probability. Observing oppositely-signed party cue and policy information ATEs would imply that the cue has the opposite effect to the information; strongly inconsistent with the as-if assumption.

The relevant estimates are displayed in Figure 7. As before, the estimates are computed from the full posterior distribution of the multilevel model fitted to obtain the aggregate results, not the result of fitting new or separate models for every issue. Thus, the issue-level estimates are mildly regularized, improving their out-of-sample accuracy on average (McElreath 2020).

We draw attention to three points of interest in Figure 7. First, like the party cue ATEs (x axis), the estimates of the policy information ATEs (y axis) all lie to the right of zero; thus, across the board exposure to party cues had the same qualitative influence on opinions as exposure to substantive policy information. However, the influence of the information was generally smaller—note the scale limits on the x and y axes are different. Second, on aggregate we can conclude with >95% probability that exposure to policy information caused respondents to update their policy opinions ~0.21 scale points toward the in-party position, despite that they didn't in fact learn the in-party position. Third, we observe a fairly equal distribution of strong, weak, and inconsistent evidence of corroboration of the as-if assumption across policy issues. Accordingly, the aggregate estimate is classified as weak corroboration of the assumption overall: the aggregate party cue and policy information ATEs had the same sign and both were distinguishable from zero, but the party cue ATE was quantitatively larger. These results imply that party cues do cause (qualitatively) people to behave as-if better informed.




Figure 7. Party cue and policy information ATEs colored by corroboration of the as-if assumption.

Implications

The causal effect of party cues on people’s policy opinions is among the most well-established phenomena in political science. Yet, the reason *why* people are influenced by party cues remains conspicuously unclear (Bullock 2020). A long-standing and influential theory holds that party cues function as heuristics; people rely on them because they lack substantive information about policy, the motivation to think for themselves in order to form an opinion, or both. A key prediction of this theory is that, in the counterfactual world where people possessed more information about policy, a greater propensity for effortful thinking, or both, the influence of party cues would be diminished.

Prior work has not yielded a clear answer to, nor empirical test of, this key prediction. This is because the dominant study design operationalizes awareness of policy information and propensity for effortful thinking via observed covariates. As a recent review of research on party cues put it, “we [scholars] do not let our variables vary” (Bullock 2020, 16). This opens the door to a variety of confounding biases, described in the introduction of this paper, which have not been (and maybe *cannot* be) adequately addressed in studies using the covariate design. Among the small minority of studies that obviate these challenges via randomization, the results appear inconsistent, and their designs place considerable constraints on what is possible to learn from the results. Specifically, the focus on only a small handful of policy issues runs the risk of painting a highly idiosyncratic, and thus poorly generalizable, picture of the evidence for or against the key predictions of the theory. Furthermore, to our knowledge only policy information—not cognitive effort per se—has been the target of randomization.

We aimed to address all these limitations with our design. In doing so, we found that exposure to substantive policy information causally attenuated the influence of party cues, consistent with the prediction of the heuristic theory. In contrast, there was little evidence that cognitive effort per se causally attenuated the influence of party cues, inconsistent with the

theory. Our results provide new evidence, and have diverse implications, for the long-standing and influential theory that party cues function as heuristics.

Judging by the results of our treatment compliance analysis (Figure 2), cognitive effort was successfully induced (on average) by its namesake treatment in our design. Thus, our null result strongly implies that effortful thinking about policy *per se*—in the absence *or* presence of policy information—does not diminish reliance on party cues. This null result is strictly inconsistent with the long-standing and influential notion that party cues are heuristics for avoiding effortful thinking, and validates mounting circumstantial evidence that undercuts this notion. Bullock (2020, 13-14) recently summarized such evidence—for example, evidence that exposure to party cues may increase (rather than decrease) people’s response latency when considering policy issues (Petersen et al. 2013). Our null result offers more direct evidence to this point: randomly inducing cognitive effort did not reduce the influence of the party cue, as would be expected if the cue was used purely to avoid effortful thinking.

Accordingly, this null result also cuts against a long line of theorizing that party cues will be less influential when people think more deeply (Bakker, Lelkes, and Malka 2020; Bullock 2020; Kam 2005; Mondak 1993). Our result coheres with studies that find self-reported “need for cognition” does not moderate reliance on party cues (Bullock 2011; Kam 2005; Tappin 2020); but, crucially, enables superior confidence in this conclusion, because our design avoids the unobserved confounding that may have driven the previous results. Thus, current evidence suggests that, if propensity for cognitive effort *does* attenuate the influence of party cues, it must do so via a route that is alternative to deeper thinking *per se*, or deeper thinking about the policy information put in front of one’s eyes. One possibility is that cognitive effort causes greater *exposure* to substantive policy information—via information search behavior—which, in turn, attenuates the influence of party cues. This possibility also highlights how the notion that party cues are effort-saving heuristics may be rescued: exposure to party

cues may attenuate people’s search for further policy information. Good evidence speaking to this possibility is currently lacking (Bullock 2020).

In contrast to heuristics for thinking effort per se, our results do imply that party cues are heuristics in the sense of stand-ins for *information*. On this basis, our results offer a new perspective on the classic notion that the politically “aware” and “sophisticated” are less influenced by party cues (Bullock 2020; Kam 2005). In the covariate design, sophistication is often defined by a high score on a series of factual questions about politics, with a common assumption being that high-scoring individuals possess greater awareness of policy-relevant information; thereby, diminishing their reliance on the party cue. But, as described in the introduction of this paper, a high score is likely to co-vary with additional variables, like prior exposure to the party cue (before the study), which would similarly diminish high-scorers’ reliance on the party cue (during the study). This makes it difficult to interpret the simple finding that political sophistication correlates with diminished reliance on party cues, as reported by numerous studies (Anduiza, Gallego, and Muñoz 2013; Barber and Pope 2019; Boudreau and MacKenzie 2014; Kam 2005; Mondak 1993; Pannico 2020). By randomizing exposure to substantive policy information, and observing a corresponding attenuation of party cue influence, our results clarify that the aforementioned political sophistication findings may indeed be driven by greater awareness of policy-relevant information—and not simply by the higher likelihood of prior exposure to the party cue.

We provide new evidence consistent with the notion that party cues constitute a valid informational basis for people’s opinion formation: on average, party cues caused people to update their opinions in the same direction as substantive policy information. This needn’t have been the case. Exposure to policy information could have attenuated the influence of party cues without moving opinions in the same direction, or even while moving them in the opposite direction. Either of these outcomes would have undermined the notion that party cues are a

valid informational shortcut, which is predicated on the assumption that the cue allows voters to form policy opinions that approximate those they would have formed were they more knowledgeable of the policy details (Hobolt 2006; Lupia 1994; Lupia and McCubbins 1998; Mondak 1993). Thus, our data are broadly in line with the normatively-appealing assumption that environmental cues, such as those from political parties, offer a part-solution to the democratic dilemma of a politically ignorant electorate (Downs 1957; Lupia and McCubbins 1998; Mondak 1993; Popkin 1995; Sniderman, Brody, and Tetlock 1993).

Caution is warranted before taking our data as unbridled support for this normative interpretation, however. The idea that it is normative for democracy for politically-ignorant voters to rely on party cues in lieu of knowledge of policy information hinges not only on them both moving opinions in the same direction, but, also, on voters' motivations for forming policy opinions in the first place. If voters are primarily motivated to support policies whose *outcomes* benefit them—for example, whether they or their partners could get easier access to abortion if policy X were legislated—then our data seem to cohere with the normative interpretation. This is because a core assumption of normative theories of democracy is that voters form preferences on the basis of potential or observed policy outcomes (Achen and Bartels 2017).

But there are plausible reasons to think this may *not* be voters' primary motivation for forming their policy preferences. Like others (Downs 1957), advocates of this view note that, for most policy issues, the opinions and behavior of any individual voter have negligible effect on the realization of policy outcomes (Kahan 2016; Williams 2019, 2020). Thus, they argue that individual voters have negligible motivation to be sensitive to policy outcomes when forming their policy opinions. Instead, voters' primary motivation may typically be to conform to the positions of, and thus obtain rewards or avoid sanctions from, their communities; a motivation that seems rather *incoherent* with normative theories of democracy, and one that may be satiated by heuristics or otherwise. On this view, exposure to party cues and policy

information could both cause updating in the same direction because both provide information about the acceptable position in one's community. Thus, the extent to which our data support an interpretation of party-cues-as-heuristics that is normative for democracy depends upon the assumptions one is willing to make regarding voters' unobserved motivations.

Though we found that exposure to substantive policy information causally attenuated the influence of party cues by ~35% on average, we estimated notable variation in the magnitude of attenuation across policy issues; from as small as ~25% to as large as ~100% of the total influence (Figure 5). As argued earlier, variation such as this illustrates the critical importance of studying more than a handful of policy issues when testing theories of party cue influence. Historically, scholars have focused on their ability to generalize from the sample of survey respondents to the wider, unobserved population; but a sample of policy issues likewise comes from a wider, unobserved population—one to which we would (often) also like to generalize. Expanding the samples of policy issues in our studies is thus essential to building cumulative and generalizable knowledge, and to avoiding a published literature that consists of many studies canvassing a small number of policy issues that arrive at different and potentially inconsistent conclusions due to an idiosyncratic sample of issues (Guisinger and Saunders 2017; Tappin 2020; Yarkoni 2019).

Given that the average attenuation of party cue influence we observed under exposure to substantive policy information was just one-third of the total influence, a natural question to ask is: what explains the remaining party cue influence? One obvious explanation is that people still do not feel they possess sufficient policy information to fully ignore the party cue, despite the presence of the information treatment. This raises yet another question: if our information treatments were considerably longer, more detailed or otherwise different, would the average attenuation be closer to (or perhaps even further from) 100%? It is difficult to say. However,

our study provides a new empirical benchmark for future investigations of these and important related questions, and our design and analytic strategy offers a method for such investigations.

References

- Acharya, Avidit, Matthew Blackwell, and Maya Sen. 2018. "Analyzing Causal Mechanisms in Survey Experiments." *Political Analysis* 26(4): 357–78.
- Achen, Christopher H., and Larry M. Bartels. 2017. *Democracy for Realists: Why Elections Do Not Produce Responsive Government*. Princeton University Press.
- Agadjanian, Alexander. 2020. "When Do Partisans Stop Following the Leader?" *Political Communication*: 1–19.
- Anduiza, Eva, Aina Gallego, and Jordi Muñoz. 2013. "Turning a Blind Eye: Experimental Evidence of Partisan Bias in Attitudes Toward Corruption." *Comparative Political Studies* 46(12): 1664–92.
- Bakker, Bert N., and Yphtach Lelkes. 2018. "Selling Ourselves Short? How Abbreviated Measures of Personality Change the Way We Think about Personality and Politics." *The Journal of Politics* 80(4): 1311–25.
- Bakker, Bert N., Yphtach Lelkes, and Ariel Malka. 2020. "Understanding Partisan Cue Receptivity: Tests of Predictions from the Bounded Rationality and Expressive Utility Perspectives." *The Journal of Politics* 82(3): 1061–77.
- Barber, Michael, and Jeremy C. Pope. 2019. "Does Party Trump Ideology? Disentangling Party and Ideology in America." *American Political Science Review* 113(1): 38–54.
- Bisbee, James, and Diana Lee. 2020. *Mobility and Elite Cues: Partisan Responses to Covid-19*. Politics and International Relations. preprint.
<https://preprints.apsanet.org/engage/apsa/article-details/5f4818bb40d6e400121eda81> (November 5, 2020).
- Boudreau, Cheryl, Christopher S. Elmendorf, and Scott A. MacKenzie. 2019. "Roadmaps to Representation: An Experimental Study of How Voter Education Tools Affect Citizen Decision Making." *Political Behavior* 41(4): 1001–24.
- Boudreau, Cheryl, and Scott A. MacKenzie. 2014. "Informing the Electorate? How Party Cues and Policy Information Affect Public Opinion about Initiatives." *American Journal of Political Science* 58(1): 48–62.
- Bullock, John G. 2011. "Elite Influence on Public Opinion in an Informed Electorate." *American Political Science Review* 105(3): 496–515.
- . 2020. "Party Cues." *The Oxford Handbook of Electoral Persuasion*.
<https://www.oxfordhandbooks.com/view/10.1093/oxfordhb/9780190860806.001.0001/oxfordhb-9780190860806-e-2> (February 18, 2020).

- Bürkner, Paul-Christian. 2017. "Brms: An R Package for Bayesian Multilevel Models Using Stan." *Journal of Statistical Software* 80(1): 1–28.
- Bürkner, Paul-Christian, and Matti Vuorre. 2019. "Ordinal Regression Models in Psychology: A Tutorial." *Advances in Methods and Practices in Psychological Science* 2(1): 77–101.
- Cacioppo, John T., and Richard E. Petty. 1982. "The Need for Cognition." *Journal of Personality and Social Psychology* 42(1): 116–31.
- Ciuk, David J., and Berwood A. Yost. 2016. "The Effects of Issue Salience, Elite Influence, and Policy Content on Public Opinion." *Political Communication* 33(2): 328–45.
- Cohen, Geoffrey L. 2003. "Party Over Policy: The Dominating Impact of Group Influence on Political Beliefs." *Journal of Personality and Social Psychology* 85(5): 808–22.
- Converse, Philip E. 2006. "The Nature of Belief Systems in Mass Publics (1964)." *Critical Review* 18(1–3): 1–74.
- Coppock, Alexander. 2019. "Generalizing from Survey Experiments Conducted on Mechanical Turk: A Replication Approach." *Political Science Research and Methods* 7(3): 613–28.
- Coppock, Alexander, Alan S. Gerber, Donald P. Green, and Holger L. Kern. 2017. "Combining Double Sampling and Bounds to Address Nonignorable Missing Outcomes in Randomized Experiments." *Political Analysis* 25(2): 188–206.
- Coppock, and McClellan. 2019. "Validating the Demographic, Political, Psychological, and Experimental Results Obtained from a New Source of Online Survey Respondents." *Research & Politics* 6(1): 2053168018822174.
- Delli Carpini, Michael X., and Scott Keeter. 1996. *What Americans Know about Politics and Why It Matters*. Yale University Press.
- Democracy Fund, and UCLA. 2020. "Nationscape Data Set Release 1." *Democracy Fund Voter Study Group*. <https://www.voterstudygroup.org/publication/nationscape-data-set-release-1> (May 22, 2020).
- Downs, Anthony. 1957. "An Economic Theory of Political Action in a Democracy." *Journal of Political Economy* 65(2): 135–50.
- Ehret, Phillip J., Leaf Van Boven, and David K. Sherman. 2018. "Partisan Barriers to Bipartisanship: Understanding Climate Policy Polarization." *Social Psychological and Personality Science* 9(3): 308–18.
- Gaines, Brian J., James H. Kuklinski, and Paul J. Quirk. 2007. "The Logic of the Survey Experiment Reexamined." *Political Analysis* 15(1): 1–20.
- Gilens, Martin, and Naomi Murakawa. 2002. "Elite Cues and Political Decision Making." *Research in Micropolitics*: 15–49.

- Gomila, Robin, and Chelsey S. Clark. 2020. "Missing Data in Experiments: Challenges and Solutions." *Forthcoming in Psychological Methods*. <https://psyarxiv.com/mxenv/> (October 17, 2020).
- Gottfried, Adele Eskeles et al. 2017. "Trajectories from Academic Intrinsic Motivation to Need for Cognition and Educational Attainment." *The Journal of Educational Research* 110(6): 642–52.
- Guisinger, Alexandra, and Elizabeth N. Saunders. 2017. "Mapping the Boundaries of Elite Cues: How Elites Shape Mass Opinion across International Issues." *International Studies Quarterly* 61(2): 425–41.
- Hobolt, Sara Binzer. 2006. "How Parties Affect Vote Choice in European Integration Referendums." *Party Politics* 12(5): 623–47.
- Imai, Kosuke, Dustin Tingley, and Teppei Yamamoto. 2013. "Experimental Designs for Identifying Causal Mechanisms." *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 176(1): 5–51.
- Kahan, Dan M. 2016. "The Politically Motivated Reasoning Paradigm, Part 1: What Politically Motivated Reasoning Is and How to Measure It." In *Emerging Trends in the Social and Behavioral Sciences*, , 1–16. <https://onlinelibrary.wiley.com/doi/abs/10.1002/9781118900772.etrds0417> (June 4, 2019).
- Kam, Cindy D. 2005. "Who Toes the Party Line? Cues, Values, and Individual Differences." *Political Behavior* 27(2): 163–82.
- Kinder, Donald R., and Nathan P. Kalmoe. 2017. *Neither Liberal nor Conservative: Ideological Innocence in the American Public*. University of Chicago Press.
- Lau, Richard R., and David P. Redlawsk. 2001. "Advantages and Disadvantages of Cognitive Heuristics in Political Decision Making." *American Journal of Political Science* 45(4): 951–71.
- Lerner, Jennifer S., and Philip E. Tetlock. 1999. "Accounting for the Effects of Accountability." *Psychological Bulletin* 125(2): 255–75.
- Lupia, Arthur. 1994. "Shortcuts Versus Encyclopedias: Information and Voting Behavior in California Insurance Reform Elections." *American Political Science Review* 88(1): 63–76.
- Lupia, Arthur, and M. D. McCubbins. 1998. *The Democratic Dilemma: Can Citizens Learn What They Need to Know?* Cambridge University Press.
- Marshall, John. 2019. "Signaling Sophistication: How Social Expectations Can Increase Political Information Acquisition." *The Journal of Politics* 81(1): 167–86.
- McElreath, Richard. 2020. *Statistical Rethinking: A Bayesian Course with Examples in R and STAN*. CRC Press.

- Mérola, Vittorio, and Matthew P. Hitt. 2016. "Numeracy and the Persuasive Effect of Policy Information and Party Cues." *Public Opinion Quarterly* 80(2): 554–62.
- Mondak, Jeffery J. 1993. "Public Opinion and Heuristic Processing of Source Cues." *Political Behavior* 15(2): 167–92.
- Mullinix, Kevin J., Thomas J. Leeper, James N. Druckman, and Jeremy Freese. 2015. "The Generalizability of Survey Experiments*." *Journal of Experimental Political Science* 2(2): 109–38.
- Nicholson, Stephen P. 2011. "Dominating Cues and the Limits of Elite Influence." *The Journal of Politics* 73(4): 1165–77.
- Pannico, Roberto. 2020. "Parties Are Always Right: The Effects of Party Cues and Policy Information on Attitudes towards EU Issues." *West European Politics* 43(4): 869–93.
- Petersen, Michael Bang, Martin Skov, Søren Serritzlew, and Thomas Ramsøy. 2013. "Motivated Reasoning and Political Parties: Evidence for Increased Processing in the Face of Party Cues." *Political Behavior* 35(4): 831–54.
- Peterson, Erik. 2017. "The Role of the Information Environment in Partisan Voting." *The Journal of Politics* 79(4): 1191–1204.
- . 2019. "The Scope of Partisan Influence on Policy Opinion." *Political Psychology* 40(2): 335–53.
- Popkin, Samuel L. 1995. "Information Shortcuts and the Reasoning Voter." In *Information, Participation, and Choice: An Economic Theory of Democracy in Perspective*, University of Michigan Press.
- Prior, Markus. 2018. *Hooked: How Politics Captures People's Interest*. Cambridge University Press.
- Rahn, Wendy M. 1993. "The Role of Partisan Stereotypes in Information Processing about Political Candidates." *American Journal of Political Science* 37(2): 472–96.
- Riggle, Ellen D. et al. 1992. "Bases of Political Judgments: The Role of Stereotypic and Nonstereotypic Information." *Political Behavior* 14(1): 67–87.
- Rohrer, Julia M., and Ruben C. Arslan. 2020. "Precise Answers to Vague Questions: Issues with Interactions." <https://psyarxiv.com/7fm2j/> (December 24, 2020).
- Slothuus, Rune. 2010. "When Can Political Parties Lead Public Opinion? Evidence from a Natural Experiment." *Political Communication* 27(2): 158–77.
- . 2016. "Assessing the Influence of Political Parties on Public Opinion: The Challenge from Pretreatment Effects." *Political Communication* 33(2): 302–27.
- Slothuus, Rune, and Claes H. de Vreese. 2010. "Political Parties, Motivated Reasoning, and Issue Framing Effects." *The Journal of Politics* 72(3): 630–45.

- Sniderman, Paul M., Richard A. Brody, and Phillip E. Tetlock. 1993. *Reasoning and Choice: Explorations in Political Psychology*. Cambridge University Press.
- Somin, Ilya. 2016. *Democracy and Political Ignorance: Why Smaller Government Is Smarter, Second Edition*. Stanford University Press.
- Tappin, Ben M. 2020. *Estimating the Between-Issue Variation in Party Elite Cue Effects*. PsyArXiv. preprint. <https://osf.io/p48zb> (October 8, 2020).
- VanderWeele, Tyler J., and Eric J. Tchetgen Tchetgen. 2014. “Attributing Effects to Interactions.” *Epidemiology* 25(5): 711–22.
- Vössing, Konstantin. 2020. “Shaping Public Opinion about Regional Integration: The Rhetoric of Justification and Party Cues.” *Political Studies*: 0032321720905130.
- Westfall, Jacob, and Tal Yarkoni. 2016. “Statistically Controlling for Confounding Constructs Is Harder than You Think.” *PLOS ONE* 11(3): e0152719.
- Williams, Daniel. 2019. “Socially Adaptive Belief.” *Forthcoming in Mind & Language*.
- . 2020. “Motivated Ignorance, Rationality, and Democratic Politics.” *Synthese*. <http://link.springer.com/10.1007/s11229-020-02549-8> (December 26, 2020).
- Yarkoni, Tal. 2019. *The Generalizability Crisis*. PsyArXiv. preprint. <https://osf.io/jqw35> (December 9, 2019).