Affect or Ideology?: The Heterogeneous Effects of Political Cues on Policy Support*

Sam Fuller†[1], Nicolás de la Cerda‡[2], and Jack Rametta§[1]

[1]University of California, Davis
[2]University of North Carolina at Chapel Hill

April 1, 2023

Abstract

Affective polarization has grown into a burgeoning subfield of American politics and political behavior research. Focusing on an individual's attachment toward their own party and animosity toward the other party, mounting evidence points to an unprecedented rise in partisan animosity in the United States. However, little research has been done to probe and understand the latent dimensions that inform partisan affect *and* influence an individual's policy support. Using a nationally representative sample of 3,300 US adults, we aim to better understand the relationship between affect and policy support. Specifically, we contribute to the literature by first identifying multiple dimensions of an individual's level of partisan affect using various scaling methods, while also using Bayesian Aldrich-McKelvey scaling to recover a robust estimate of individuals' ideology. Second, we look at how these levels of affect and ideology influence and moderate the effect of in- and out-party cues. Furthermore, we employ a new method, causal forest, to explore heterogeneous treatment effects based on observable characteristics, with an emphasis on individuals' levels of affect and ideological position. Overall, we find that out-party cues, but not in-party ones, have heterogeneous effects on individuals' policy support. Specifically, we find that affect (both thermometer-level and social distance), not ideology, conditions the effects of out-party policy cues. Those who have larger differences of in- minus out-party affect having larger negative reactions to these cues.

*As an early working paper, mentions of analysese in the Appendix are placeholders. A more complete draft of this paper will include all referenced Appendix analyses.

†PhD Candidate, Department of Political Science. E-mail: sjfuller@ucdavis.edu.

‡PhD Candidate, Department of Political Science, University of North Carolina at Chapel Hill. E-mail: ndelacerda@unc.edu.

§PhD Candidate, Department of Political Science. E-mail: jtrametta@ucdavis.edu.

1 Introduction

Research in psychology and political science has consistently found that citizens process information in biased ways. The kind of information we pay attention to, the way in which we incorporate it, and how we evaluate it, is driven by motivation. Reasoning is always motivated in the sense that it is directed toward a goal (Kunda 1990). Prior stored information plays a key role in guiding the way in which we evaluate new information in at least two ways. First, when faced with new evidence, individuals are more likely to accept information that is consistent with their previous beliefs and more critical with information that challenges their priors (Ditto and Lopez 1992). Second, cognition is often affectively charged, concepts stored in long-term memory are associated with positive or negative affect. In turn, these affectively charged concepts shape the way in which we process new information (Lodge and Taber 2013).

In politics, these two mechanisms mostly work in tandem. Partisans' reasoning, this is, reasoning driven by the need of arriving to a conclusion that is consistent with the one of the party a citizen identifies with (Lodge and Hamill 1986), can be explained by either of these two causal channels. Republicans may accept and incorporate new information based on the consistency of the new information with their prior beliefs (motivated bias) or because the information is provided by a Republican official triggering positive affects (hot cognition).

However, because these two mechanisms work together, at an empirical level, it is very difficult to disentangle their unique contribution to attitude bias. In this article, we evaluate the extent to which political affect and ideology moderate information processing.

We evaluate the extent to which affect and ideology shape attitudinal bias, and thus policy support, among Democrats and Republicans using a survey experiment. Participants in our study were exposed to three hypothetical legislative bill proposals on a proposed $12/hour minimum wage increase, a border protection budget increase, and an increase in American farm subsidies. We chose the first two of those bills because they can easily be mapped to the current partisan divide, but are relatively moderate and lower salience than, for example, abortion or gun control. While Democrats support and Republicans disapprove increasing minimum wages, the opposite is true for attitudes regarding border protection. The farm subsidies bill was chosen as a relatively non-partisan policy. Overall, these three bills were selected to 1) generally cover a left-right continuum, and 2) to be moderate enough that the partisan cue is plausible. For each individual and each one of these bills, respondents were randomly exposed to one of three conditions: a Republican cue, Democratic cue, and non-partisan group cue.

By allowing combinations that are ideologically inconsistent (e.g., Republicans supporting minimum wage increases or Democrats supporting border protection budget increases) we seek to identify the distinct effect of affect and ideology on political cueing. While we expect the effect of affect and ideology to be in the same direction when participants are exposed to ideologically consistent cues, we expect the opposite for ideologically inconsistent ones. Accordingly, the experimental design provides us with a unique opportunity to untangle the effects of affect and ideology.

Our contributions to this literature are two-fold: First, we estimate affect and ideology utilizing scaling methodologies that both more granularly capture these concepts and that adjust for differential item-functioning (DIF). Previous studies on cueing effects have tended to use rough measures, particularly self-identified ideology. Second, whereas previous literature relies on standard conditional average treatment effect (CATE) calculations (mostly by using interactive ordinary least squares regressions; e.g., Kam 2005, Barber and Pope 2019), we exploit a methodology designed explicitly for detecting treatment effect heterogeneity and calculating CATEs: causal forests (CF; Wager and Athey 2018). This agnostic, and relatively conservative, approach to detecting heterogeneity, which incorporates all measured covariates, gives us confidence that our CATEs are both not spurious and accurately calculated. Furthermore, we also utilize CF to evaluate the relative importance of affect and ideology as cueing moderators, finding that both are integral for determining out-party cueing treatment effects.

In line with previous research, we find that partisanship has a strong effect on information processing. Both Democratic and Republican cues strongly shape policy support. However, contrary to both recent and seminal work on cuing effects, we find that while in-party cues do have positive effects, they are not conditioned by other covariates (with the exception of Republican respondents with the border funding policy). We believe that these contrary findings are the result of using both better measures of affect and ideology (and our use of spatial knowledge, rather than political knowledge) as well as our use of causal forests. Furthermore, our results from our causal forest show that our multiple measures of affect are by far the most important moderators of treatment effects. The relative importance of affect outweighs the capacity of any other predictor at accurately predicting heterogeneity in the effects, including ideology. Contrary to what previous research suggests, we find little evidence of political interest or spatial knowledge as a source of heterogeneity.

2 Affect, Ideology, and Cognition

The literature on political cognition has systematically shown that partisanship plays a fundamental role in shaping information processing. Both Democrats and Republicans are more likely to accept information that is congruent with their partisan loyalties (Ditto and Lopez 1992), leading to sharp gaps in perceptions of political events (Bartels 2002).

However, it is not clear what are the key mechanisms that drive these perceptual differences. According to Lodge and Taber's dual-process model of information processing (2013), perceptual differences arise from unconscious affective processes that then shape conscious and cognitive reactions to political events. Accordingly, the way in which we evaluate political events is deeply related to affectively charged, difficult to control processes. Lodge and Taber (2013) highlight two channels through which information processing becomes affectively charged: (1) hot cognition and motivated bias, and (2) affect priming and affect contagion. While both of these channels operate at an unconscious level and are affectively driven, their theoretical underpinning is different. While the first channel is related to feelings that are intrinsic to the stimulus, the second channel is related to affect that is only incidentally related. In simple terms, on one side, hot cognition and motivated bias arise from the content of the message: the nature of a given policy or the stand regarding a particular issue or debate. On the other, affective priming and contagion are more related to the way in which a message is conveyed and who the messenger is.

This means that partisan perceptual gaps can arise through two commonly related but distinct channels, the content of a particular message, and the characteristics of the messenger. Accordingly, we expect these two channels to have distinct effects over the way in which citizens process information.

Prior attitudes play a fundamental role for the first of these two channels. The way in which citizens feel about an issue can positively or negatively charge new information. In particular, individuals are more likely to accept information that is consistent with their prior views and to reject inconsistent information. While consistent information is quickly incorporated, individuals tend to use more cognitive resources counterarguing inconsistent information. That is, when people feel strongly about an issue, they tend to engage in both confirmation and disconfirmation biases (Taber and Lodge 2016). Consistent with this theoretical model, previous literature on partisan cueing has found that when individuals receive information about the policy positions of either in- or out-party elites, they adjust their positions (Bullock 2011, 2020).

The second channel is driven by elements that are incidental to a particular stimulus, particularly the messenger's identity and characteristics, and how the message is delivered. Individuals who have high negative affect towards the Republican party can reject a stimuli, for example a policy proposal, not based on the nature of the proposal, but because the proposal is supported by someone who the individual does not like. This means that individuals can negatively charge a message, even in cases when they agree with the policy content. In fact, previous research has shown that individuals' partisan affect can increase the probability of accepting in-party congruent misinformation and decrease the probability of accepting out-party congruent misinformation (Jenke 2023).

Because these two channels have different cognitive underpinnings, we expect each one of them to operate independently and to have distinct effects over the ways in which individuals process political information.

- Hypothesis 1: An (out-)in-party cue will (decrease) increase support for a given policy regardless of the ideological bent of the policy.
- Hypothesis 2: Liberal and conservative ideological positions increase the support for ideologically congruent policies, regardless of what partisan group supports the policy.
- Hypothesis 3a: Positive in-party affect increases support for policies that are supported by a co-partisan group, regardless of the particular content of the policy.
- Hypothesis 3b: Negative out-party affect decreases support for policies that are supported by an out-party group, regardless of the particular content of the policy.

Finally, it is worth highlighting that the two channels through which information becomes affectively charged commonly work in tandem. Because the content of a message is usually related to the characteristics of the messenger, when citizens are exposed to information that is congruent with their priors, it usually comes from a source that is associated with positive affect. In practical terms, Democrats tend to support liberal policies and Republican tend to support conservative ones. Accordingly, separating the effects of these two channels is particularly challenging at an empirical level. In the next section, we present an experimental design that is uniquely tailored to confront these challenges.

2.1 Experiment Details

Our experiment follows a relatively standard cueing experiment, where individuals are presented with three bills (randomly ordered):

{Cue} has proposed a bill to increase the federal minimum wage to \$12 an hour from its current level of \$ 7.25 an hour. This change is intended to increase incomes for low-wage workers while limiting related costs for affected businesses.

{Cue} has proposed to increase the US Customs and Border Protection budget by 20%. The stated purpose is to increase efficiency and safety at US border crossings.

{Cue} has proposed a bill to increase subsidies for American farmers due to rising competition from markets abroad. This is intended to both protect American farmers and to keep the cost of food low for Americans.

Where "{Cue}" is block-randomized within the party identification blocks of our respondents, specifically Democrats (plus leaners), Republicans (plus leaners), and Independents. Respondents were block randomized into three different groups for each one of the bills: control, Republican treatment, Democratic treatment.

We selected these proposals based on their general ideological bent, and made sure to have them be relatively "moderate" proposals to ensure that the proposals and the groups that supported them seemed believable. Specifically, the minimum wage bill is generally expected to be seen as liberal, the Border Protection budget as conservative, and the farming subsidy as nonpartisan/non-ideological. Because the cues are randomized for each one of the policy proposals, some respondents were presented conservative proposals with a Democratic cues, and liberal proposals with a Republican cue.

Figure SX in the appendix illustrates the baseline policy support of our three policies in the control (nonpartisan group) condition by PID. In general, . Our range of support, and the granularity of our support variable, across all three of our policies and within each PID, gives us reason to believe that our experiment is not deeply subject to ceiling or floor effects.

2.1.1 Block Randomization

Before the experiment, respondents were first asked for their party identification. This was done to allow for block-randomization in our sample, where individuals within each PID (Democrat, Republican, Independent) were then randomly assigned to treatments for each cue. This was done to ensure that we had adequate power to estimate the effects of each cue, by each policy, by each PID. If standard randomization were used, it could lead to imbalanced PID by cues and policies. For example, Respondents assigned to the Republican Elite cue for the minimum wage policy, leading us to be unable to estimate the effect of that cue on that policy for Democratic or Independent respondents. To ensure the success of this randomization approach, we first simulated data before fielding and, after the survey was completed, we checked the distribution of respondents to each cue by each policy by PID. The largest difference between any PID, cue, and policy combination is 4% with the vast majority of combinations maintaining a less than 1% difference. We also conduct standard balance tests for each treatment condition by our set of analyzed covariates. We find no difference between any of the conditions across all of our covariates, on average. The distribution of our respondents and the results of our balance tests are reported in Table SX and Table SY in our appendix.

3 Data

Our data comes from a nationally-representative, novel dataset of American adults collected online by the survey firm Lucid in September and October of 2022. The dataset contains 3,359 observations, and was collected with quotas based on Census data for gender, race, region, and age; details regarding the representativeness of this sample can be found in Table A1 of the appendix. Along with this data, we also collect data on other covariates: demographics, religion, political interest and identification, political attitudes, affect and ideology.

To ensure quality responses from our respondents, we include a single attention check and consent to provide honest and accurate answers. Respondents who failed the attention check or refused to consent to providing good answers, were immediately prevented from completing the survey and thus entering our dataset. Furthermore, we also check for "straightliners" (individuals who answer quickly and select the same response for all stimuli in a grid-question) among those who did complete the entire survey (Kim et al. 2019). We identify and remove 261 individuals who straightlined in various parts of our survey resulting in a complete dataset of 3,098 respondents. [1]

3.1 Scaling Ideology & Affect

We collect standard measures of both affect and ideology. We use standard measures from the American National Election Study (ANES) for self-reported ideology (on a 7-point Likert scale from "Extremely liberal" to "Extremely conservative") and feeling thermometers for both the Democratic and Republican Parties (0–100 with 0 being coldest and 100 being warmest feelings). Traditionally, a standard partisan affect measure is then calculated by taking the difference between an individual's feeling thermometer towards their in-party and out-party (Iyengar et al. 2019), resulting in a -100–100 point scale.

While the previous literature has tended to use these measures in their raw format (e.g., Barber and Pope 2019), simply taking an individual's self-reported ideology and, in the rare cases partisan affect is measured, raw affect measures, we instead opt to leverage three different scaling methodologies to recover more robust measures of ideology, thermometer-level affect, social distance, and trust, respectively.[2] First, we both include multiple measures of affect and scale them given previous findings suggesting that simple thermometer-level affect does not capture the entire phenomenon (Bankert 2020; Druckman and Levendusky 2019;

[1] This sample size is still more than sufficient to satisfy 90% power for the average treatment effects reported in this study.

[2] Importantly, we do test self-reported ideology and the standard measure of affect for thermometers, in-party thermometer minus out-party thermometer. However, social distance and trust scores require some sort of scaling method (Druckman and Levendusky 2019), so we do not test another formulation. Using both self-reported ideology and the standard measure of thermometer-affect we recover similar directional results, albeit with larger confidence bands, in all cases. Consequently, we opt to report the results using our scaled measures and raw/standard measures in the main text and SECTION S3 of the appendix, respectively.

Kingzette 2020). In general, this gives us more leverage to identify the general effects, or at least the importance, of various, robust measures of affect. Second, we opt to scale ideology using Bayesian Aldrich-McKelvey Scaling (BAM), given the distinct problem of differential-item functioning that occurs with ideological self-placements (Hare et al. 2015). Overall, and especially compared to the previous literature, these robust and various measures of affect and ideology give us the ability to more confidently calculate conditional treatment effects in our sample. For interested readers, the details of these specific procedures, including information on the raw measures collected in the survey, are included in section S3 of the appendix.

4 Results

This section proceeds by first testing how the effect of political cues is moderated by partisanship. Consistent with previous research, we find that the way in which citizens process information is shaped by partisan loyalties. With those baseline results, we then test for heterogeneous treatment effects of affect and ideology using Causal Forest. This methodology better enables us to first test for any treatment effect heterogeneity and then, if there is, to identify which variables are causing the heterogeneity.

4.1 Baseline Results

Using a simple OLS regression, with unadjusted standard errors, we calculate conditional average treatment effects (CATEs) by respondent PID, reporting our results in Figure 1.[3] Figure 1 is broken down into two panels, with the left representing the pooled Democratic cues and the right representing the pooled Republican cues. Effects for these cues (Y-axis) are presented with 95 and 80% confidence-intervals—thin and thick lines, respectively—and are broken down within each panel by a respondent's self-reported PID (X-axis; leaners are

[3]For a table of these same results, please see section SZ in the Appendix. Furthermore, we also calculate these effects using a Causal Forest and bootstrapped standard errors, finding similar effect sizes and confidence intervals. Results are also included in section SZ of the Appendix.

included in weak PIDs). Overall, we find results that are highly consistent with previous research. In general, partisans are influenced by both in- and out-party cues across the pooled policies. Furthermore, the effect sizes and significance increase by the strength of PID, for both Democrats and Republicans. Independents, as we may expect when not controlling for ideology, consistently have null effects across both Democratic and Republican treatments.

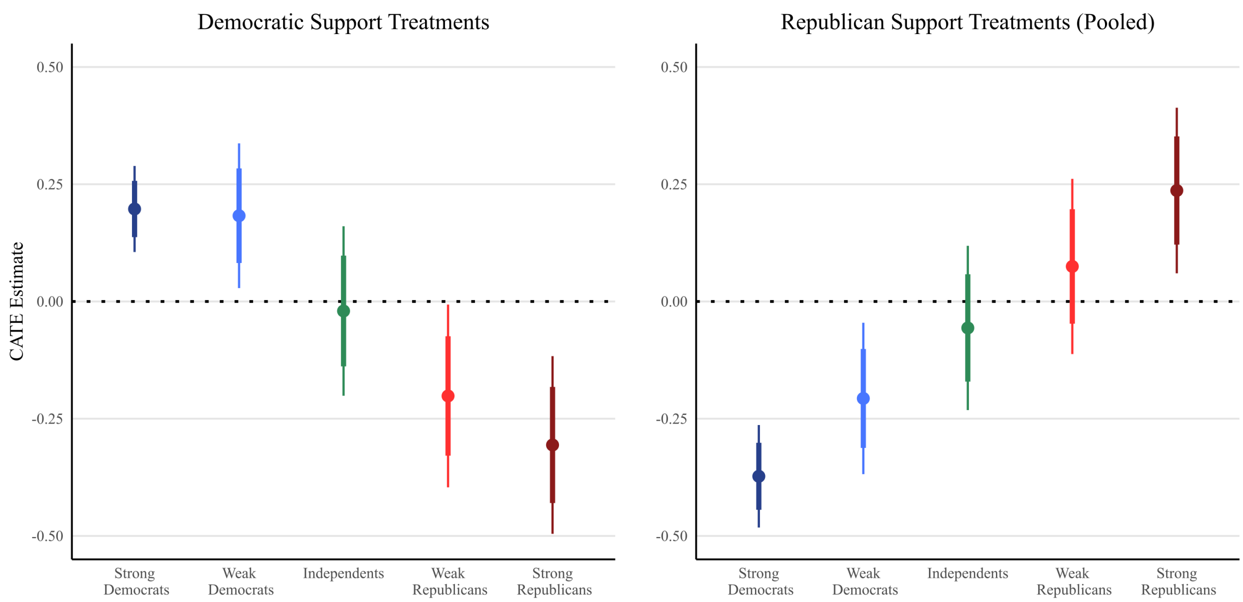

Figure 1: Pooled Cue CATE Estimates by Respondent and Cue PID

One notable trend among these results is the general size of the effects of out-partisan cues (e.g., Republican treatments among Democratic respondents). In general, the effect sizes tend to be larger and the confidence bands further away from zero. These results clearly suggest that PID or, as is more likely, PID as some rough substitute for some dimension of underlying ideological and/or affective beliefs are influencing the effects of these cues. Consequently, we move to further investigate these potential influences.

While these results are promising, in that they are consistent with previous findings surrounding simple partisan policy cues, there are limitations to using standard OLS for calculating CATEs. Specifically, there is no clear test for detecting real treatment effect heterogeneity above and beyond checking the statistical significance of the conditioning variable.

Furthermore, the specification of the regression used to calculate any CATEs has significant influence on whether or not these variables do appear statistically significant. Specifically, this manifests itself in two ways: 1) the choice of a simple interaction between the treatment assignment and the conditioning variable of interest or in calculating treatment effects by subsetting each regression to the conditioned variable group; and 2) in the decisions of both what variables to condition on and what covariates should be included in the calculation of any given CATE. Finally, as with any estimation technique, OLS necessarily assumes additive, linear effects, which significantly constrains the estimation of any CATEs.

Given these limitations, and our distinct interest in accurately estimating CATEs for respondents' ideology, affect, and spatial knowledge, we opt to use a causal forest to not only calculate our CATEs, but to first even determine if treatment effect heterogeneity exists in our experiment. The following section proceeds by first detailing the causal forest methodology; second, reporting the results of our causal forest heterogeneity tests and subsequent variable importance score calculations; and finally, reporting our CATE estimations, including interaction effects, for important conditioning variables.

4.2 Causal Forests

In general, machine learning methods have been employed for the task of prediction, and not necessarily for inference. However, recently a variety of machine learning algorithms have been adapted specifically for different causal inference problems (Athey and Imbens 2016; Athey and Wager 2019; Chernozhukov et al. 2022). [4] In this paper, we employ one of these new methods, causal forests, a method designed to detect and estimate treatment effect heterogeneity.

Causal forests adapt the random forest estimator (Breiman 2001). Random forests are an ensemble of individual decision trees optimized to minimize prediction error (i.e., MSE, AUC, etc.). Instead of optimizing for prediction error, causal trees optimize for maximum

[4] See Grimmer, Roberts, and Stewart 2021 for more detail.

treatment effect heterogeneity. Additionally, the causal forest employs subsample splitting: half of the data is using to choose splits in the causal trees, and the other half is used to estimate the leafs. Subsample splitting ensures the causal trees are not overfitting on the training data. Once the causal forest is estimated, to test for treatment effect heterogeneity we employ a heteroskedasticity robust, omnibus test developed by Chernozhukov et al. 2022. [5] In brief, the test takes an input causal forest model and generates a best linear fit of the CATEs on held-out data. It then uses the results of this linear fit to evaluate the quality of the CATE estimates. This test acts as both an omnibus heterogeneity test and also, returns calibration information such that the fit of the causal forest model can be evaluated.

While we believe this estimation procedure is superior to subgroup analysis, in that it's more robust to model specification and more difficult to game, it does come at a cost. Because the causal forest employs sub-sample splitting, it doesn't use all of the training data to estimate the CATEs. As such, this approach does add some noise relative to subgroup analysis, and can thus be viewed as conservative in some contexts. However in our view, this is a price worth paying, as subgroup analysis is only superior to causal forest when the model is very well specified, and moreover, that the researcher knows this perfect specification ex ante. This is hardly ever the case.

For the purpose of this article, we execute 12 separate causal forest models across the three different policy cues and subsetting by party ID.

4.3 Treatment Effect Heterogeneity and Variable Importance

Based on the test detailed above, we check to see if treatment effect heterogeneity exists in our experiment based on our full set of covariates. We run this test for each policy separately, subsetting also by PID, reporting the results below in Table 1. Across all cues we find distinct evidence for treatment effect heterogeneity, but only for *out-party cues*, with the exception of the in-party cue for increased border funding for Republican respondents. This suggests

[5] Implemented by the *grf* R package.

that the effect of an in-party cue on support for any of these policies, is generally consistent across all of our measured covariates. However, for these out-party cues, for both Democrat and Republican respondents, there appears to be significant heterogeneity.

Table 1: Treatment Effect Heterogeneity Checks by Policy Cue and Pooled

Treatment	Party ID	Min. Wage	Border Funding	Farm Subsidy	Pooled
Democratic Cue	Democrats				
	Republicans		*X*		***X***
	Independents				
Republican Cue	Democrats	*X*	*X*	*X*	***X***
	Republicans		*X*		
	Independents				***X***

Consequently, we further employ causal forests to determine *which* variables are leading to this heterogeneity, reporting the bootstrapped (100 trial) calculation of variable importance for predicting heterogeneous treatment effects below in Figure 2.[6]

Importance is calculated for each variable as a weighted sum of how often, and at what depth, a variable appears in all of the trees of a given causal forest model. The logic of this is that splits earlier on in any given tree are necessarily more important for prediction than those later, and since each tree is a random selection of variables its important to sum across all trees in the forest. Finally, the scores themselves are relative, in that the sum of all scores is equal to 1. The higher a variable's score, the higher its relative importance is in accurately predicting heterogeneity in our experiment. For a detailed explanation of the calculation of these variable importance scores, see SZ of the appendix.

As seen in Figure 2, consistent across all three policies (and pooled policies), for out-party cues and for Democratic and Republican respondents, our measures of affect, specifically thermometer and social distance affect, are the most important variables for capturing treatment effect heterogeneity (with the exception of the border funding Republican cue and respondents). Interestingly, ideology is by far the most important predictor for independent

[6]We employ bootstrapping here as changes in the random number generator employed can lead to slightly different variable importance estimates.

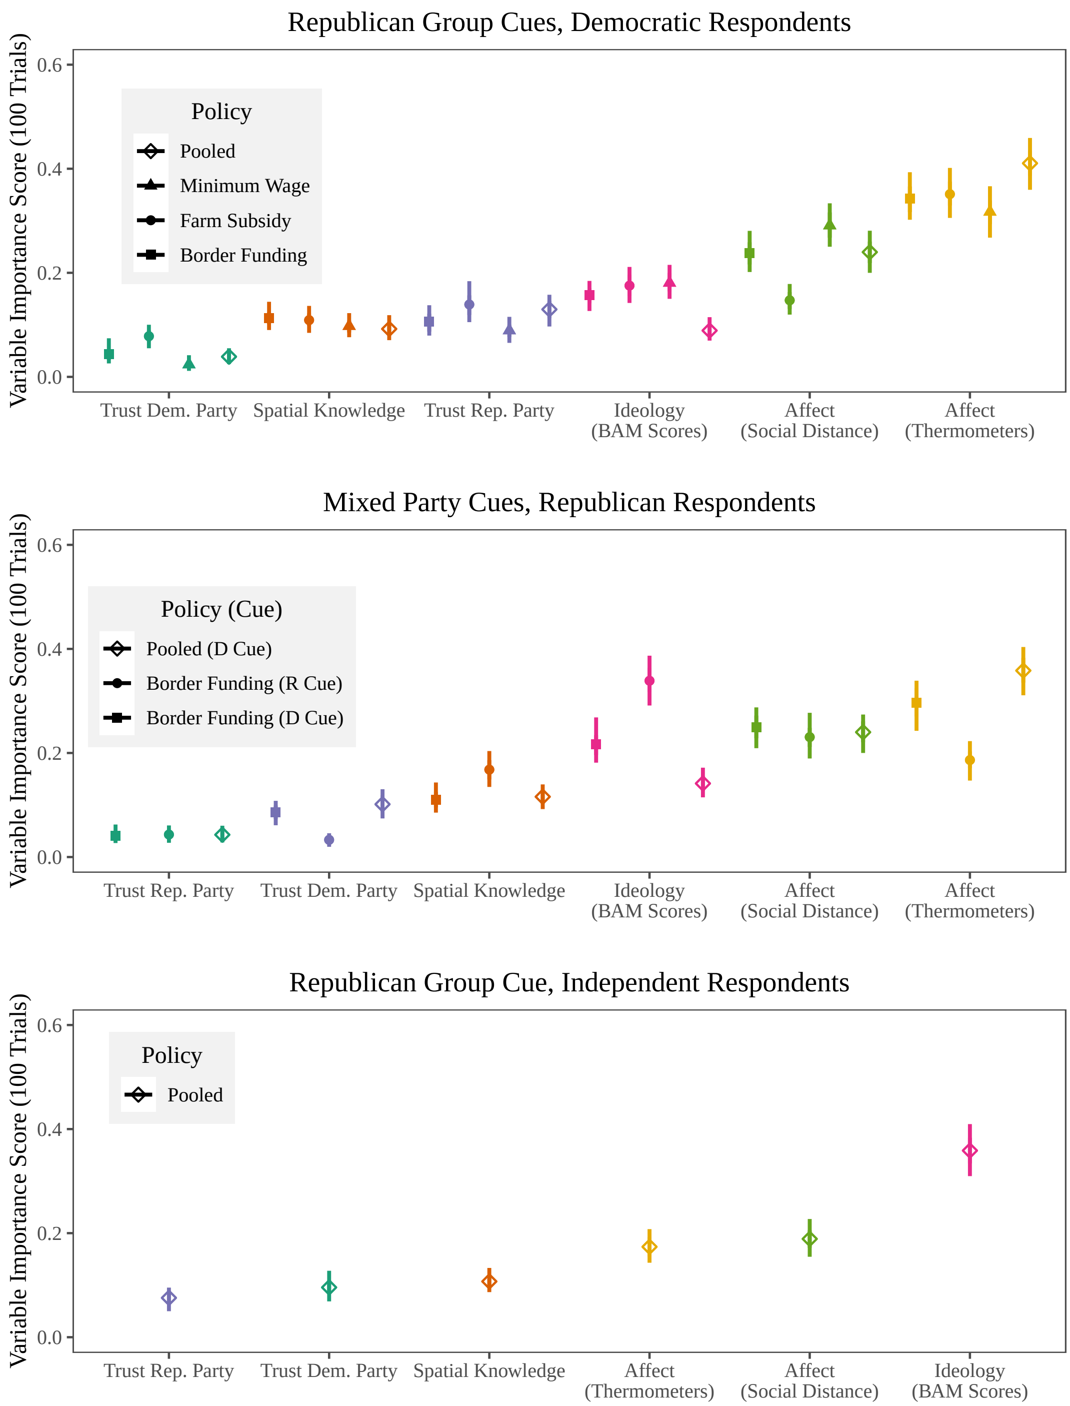

Figure 2: Variable Importance for Estimating Heterogeneous Treatment Effects

respondents. Those who are on the right (conservative) part of the scale have positive reactions to Republican cues, whereas those on the left (liberal) part of the scale have negative reactions. This makes sense, in that it appears that ideology for independent respondents acts as almost a proxy for party identification.

4.4 CATE Estimation for Affect and Ideology

We proceed in this section by calculating CATEs for affect (thermometer-level affect and social distance) and ideology by both Democratic and Republican respondents for out-party cues.[7] To do this, we map the causal forest CATE predictions over different pre-treatment covariates. These CATE estimates are derived from separate causal forest models trained on the pooled policy cues and subsetting respondents by party ID.

One benefit of using this method, is that we are able to calculate CATEs while including all three of our moderating variables of interest. Consequently, the significant results below are not the effects of these moderating variables in a vacuum, rather they are the effects of these variables when including all other covariates. For all of our CATE estimations, we include ideology, thermometer-level affect, social distance, partisan trust, and spatial knowledge. Figure 3 shows individual CATEs for each variable (rows) by PID (columns) for pooled, out-party cues. Each panel also shows the ATE (with 95% confidence bands) for the given PID.

[7]Results broken down by policy for Democratic respondents reveal the same underlying treatment effect heterogeneity. As well, plots for our other covariates reveal no significant treatment effect heterogeneity. As a result, we opt to report plots for all policies and covariates in section SX in the appendix.

Figure 3: CATEs for Democratic Respondents by Variable and PID

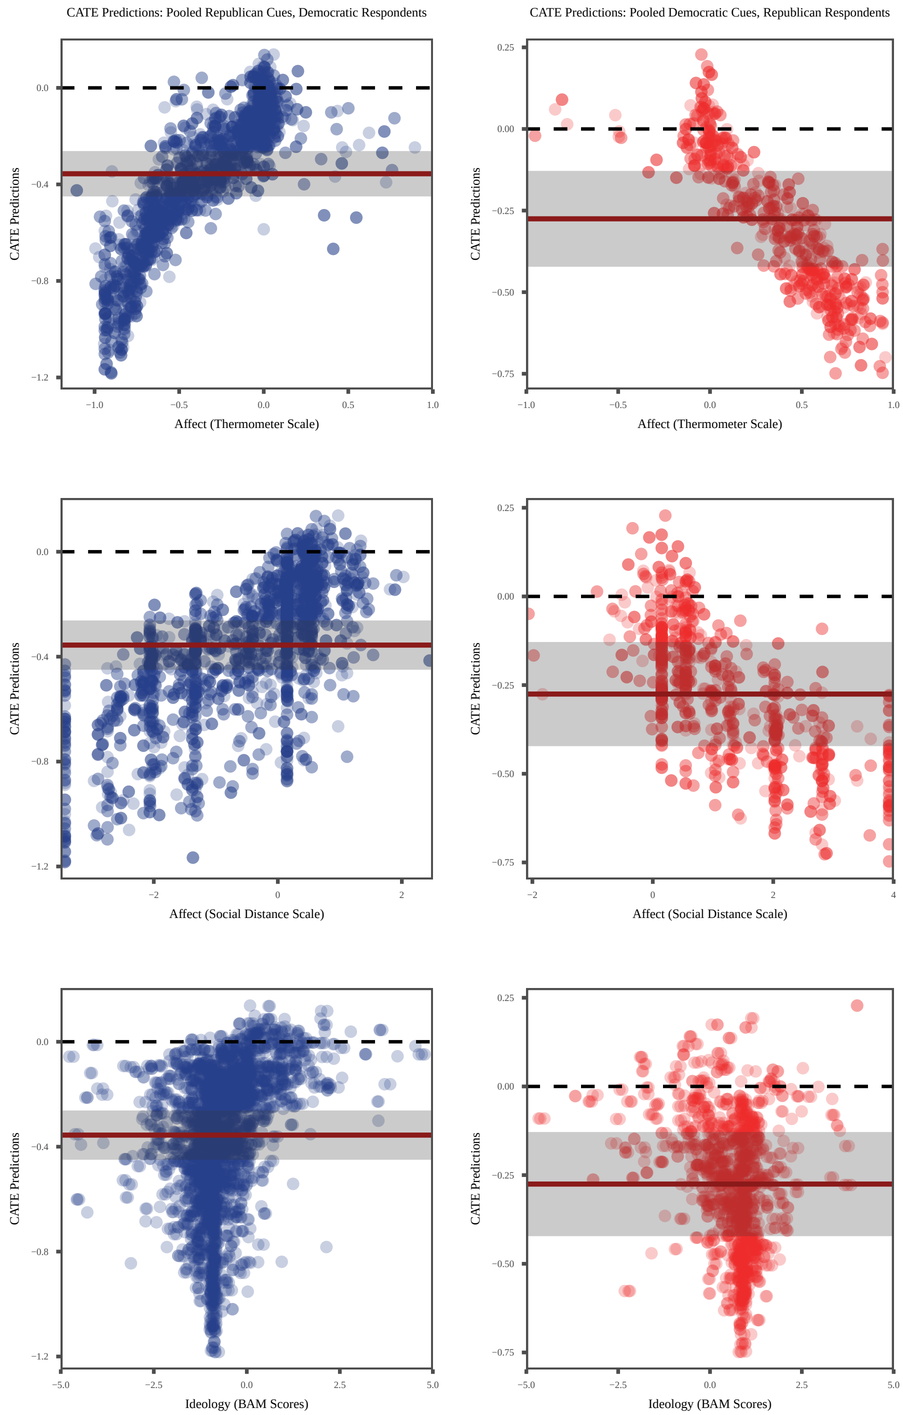

Clearly, there is significant treatment effect heterogeneity across both thermometer-level affect and social distance. For Democratic (Republican) respondents, those on the left (right) side of both scales have very large negative treatment effects whereas those on the right (left) side of the scale have essentially null effects. In contrast, ideology has no clear conditioning role for treatment effects for either Democratic or Republican respondents.

Based on the results above, we opt to explore in-depth the potential interaction CATEs for both thermometer-level affect and social distance. We do this by calculating the partial dependence between affect and social distance for out-party cues with pooled policies for Democratic and Republican respondents. This is done simply by calculating the CATE by *both* a given level of thermometer and social distance affect from both pooled, out-party cue causal forests. Figure 4 reports the partial dependence plots for thermometer (Y-axis) and social distance (X-axis) affect for out-party cues by PID (columns) for the pooled policies.

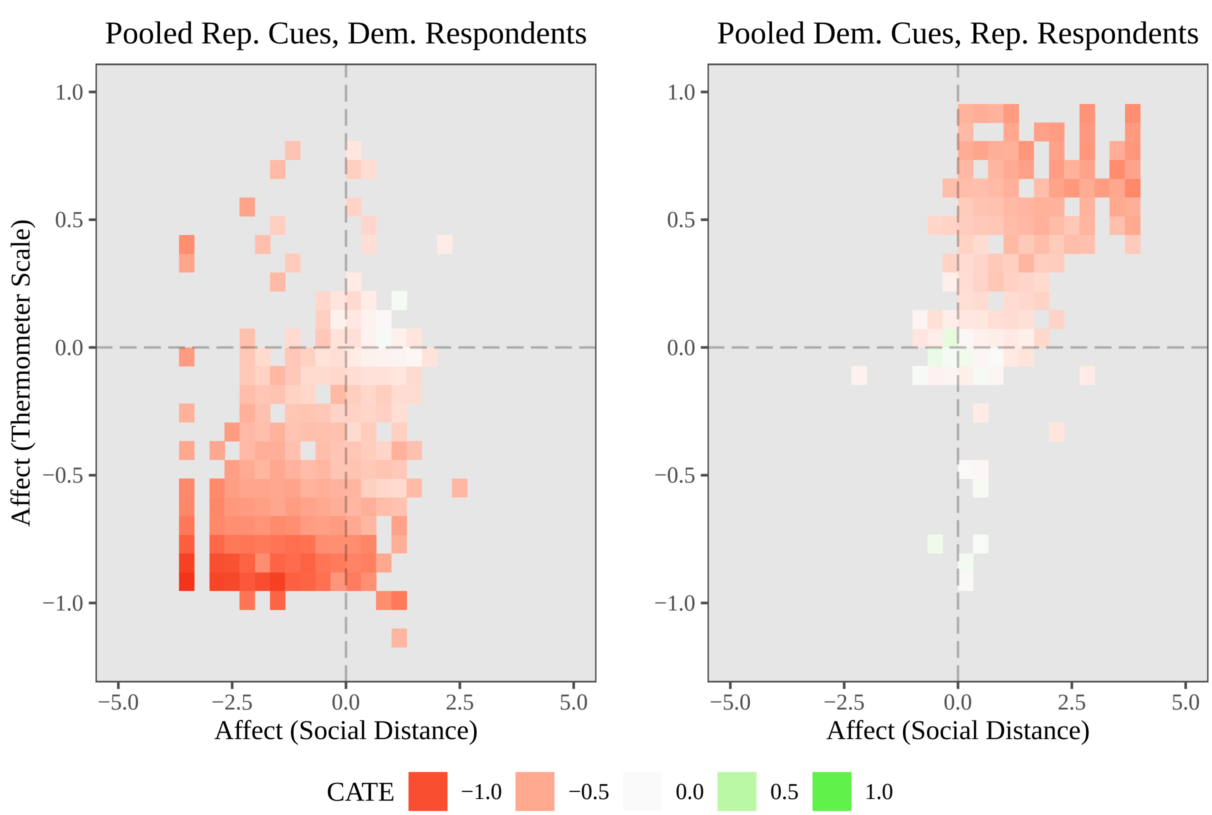

Figure 4: Partial Dependence Plots for Thermometer & Social Distance Affect

Clearly, those who are at the extreme end of both affect scales have greater negative treatment effects than those who are less extreme on one or both of the scales. Simply put, those who have higher partisan affect as a sum of both of our measures of affect, have greater negative treatment effects than their counterparts.

5 Discussion

In this paper we explore the role of affect and ideology as moderators of information processing. Based on Lodge and Taber's (2013) dual-process model of information processing, we argue that affect and ideology play a distinct role shaping the way in which political information is processed. While the model assumes that both processes are unconscious and lead to affectively charged processing, they have different theoretical underpinnings. We theorize that perceptual gaps can arise through intrinsic and extrinsic elements of a particular message. In concrete, this means that affect and ideology can shape information processing in different ways.

We test these expectations using an experimental design that allows us to separate the effects of affect and ideology. By randomizing Democratic and Republican cues to conservative and liberal proposal, our experimental design allowed us to evaluate if affect and ideology in fact play a distinct role shaping information processing. Given that we are interested in heterogeneous treatment effects, we use Causal Forest to evaluate these expectations.

Somewhat surprisingly, we find that affect, *not* ideology, is the most important moderator on the effect of partisan policy cues. However, we argue that this may be the case based on the relatively moderate bent of our policies. Essentially, we believe that individuals do not have highly concrete preferences for these policy proposals that are directly determined by their underlying ideology. In a different experiment with more salient, ideological policies, we would expect ideology to play more of a role (at least with predicting baseline policy support). Importantly, we also find that the importance of affect is far greater than political knowledge, a variable that previous work has identified as a key moderator (e.g., Barber and Pope 2019). Furthermore, we only find significant heterogeneous effects for out-partisan cues, whereas for in-party cues, it appears that cueing effects are relatively constant.

Finally, we find that both thermometer-level affect *and* social distance affect have highly significant, interactive effects on predicting treatment effects. Those who have higher differences in their in- minus out-party affect react much more negatively to out-party cues than their other, less extreme copartisans. At a theoretical level, our findings suggest that affect priming and affect contagion play a much more important role than hot cognition and motivated bias, affectively charging political information, at least for the American public.

References

Athey, Susan, and Guido Imbens. 2016. "Recursive partitioning for heterogeneous causal effects." *Proceedings of the National Academy of Sciences* 113 (27): 7353–7360.

Athey, Susan, and Stefan Wager. 2019. "Estimating treatment effects with causal forests: An application." *Observational Studies* 5 (2): 37–51.

Bankert, Alexa. 2020. "Negative and Positive Partisanship in the 2016 US Presidential Elections." *Political Behavior,* 1–19.

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.

Bartels, Larry M. 2002. "Beyond the running tally: Partisan bias in political perceptions." *Political behavior* 24:117–150.

Breiman, Leo. 2001. "Random forests." *Machine learning* 45:5–32.

Bullock, John G. 2011. "Elite influence on public opinion in an informed electorate." *American Political Science Review* 105 (3): 496–515.

———. 2020. "The Oxford Handbook of Electoral Persuasion." Chap. Party Cues. Oxford University Press, June. https://doi.org/10.1093/oxfordhb/9780190860806.001.0001. https://doi.org/10.1093/oxfordhb/9780190860806.001.0001.

Chernozhukov, Victor, Mert Demirer, Esther Duflo, and Iván Fernández-Val. 2022. *Generic Machine Learning Inference on Heterogenous Treatment Effects in Randomized Experiments.* arXiv: 1712.04802 [stat.ML].

Ditto, Peter H, and David F Lopez. 1992. "Motivated skepticism: Use of differential decision criteria for preferred and nonpreferred conclusions." *Journal of personality and social psychology* 63 (4): 568.

Druckman, James N, and Matthew S Levendusky. 2019. "What Do We Measure When We Measure Affective Polarization?" *Public Opinion Quarterly* 83 (1): 114–122.

Grimmer, Justin, Margaret E Roberts, and Brandon M Stewart. 2021. "Machine learning for social science: An agnostic approach." *Annual Review of Political Science* 24:395–419.

Hare, Christopher, David A Armstrong, Ryan Bakker, Royce Carroll, and Keith T Poole. 2015. "Using Bayesian Aldrich-McKelvey Scaling to Study Citizens' Ideological Preferences and Perceptions." *American Journal of Political Science* 59 (3): 759–774.

Iyengar, Shanto, Yphtach Lelkes, Matthew Levendusky, Neil Malhotra, and Sean J Westwood. 2019. "The Origins and Consequences of Affective Polarization in the United States." *Annual Review of Political Science* 22:129–146.

Jenke, Libby. 2023. "Affective Polarization and Misinformation Belief." *Political Behavior,* 1–60.

Kam, Cindy D. 2005. "Who toes the party line? Cues, values, and individual differences." *Political Behavior* 27:163–182.

Kim, Yujin, Jennifer Dykema, John Stevenson, Penny Black, and D Paul Moberg. 2019. "Straightlining: Overview of Measurement, Comparison of Indicators, and Effects in Mail–Web Mixed-Mode Surveys." *Social Science Computer Review* 37 (2): 214–233.

Kingzette, Jon. 2020. "Who Do You Loathe? Feelings toward Politicians vs. Ordinary People in the Opposing Party." *Journal of Experimental Political Science,* 1–10.

Kunda, Ziva. 1990. "The case for motivated reasoning." *Psychological bulletin* 108 (3): 480.

Lodge, Milton, and Ruth Hamill. 1986. "A partisan schema for political information processing." *American Political Science Review* 80 (2): 505–519.

Lodge, Milton, and Charles S Taber. 2013. *The rationalizing voter.* Cambridge University Press.

Taber, Charles S, and Milton Lodge. 2016. "The illusion of choice in democratic politics: The unconscious impact of motivated political reasoning." *Political Psychology* 37:61–85.

Wager, Stefan, and Susan Athey. 2018. "Estimation and inference of heterogeneous treatment effects using random forests." *Journal of the American Statistical Association* 113 (523): 1228–1242.