Appendix: Estimating the persistence of party cue influence in a survey experiment

Contents

1 Policy issues and party cue treatment

The policy issues are listed in Table A1, with the corresponding party cue information used for the party cue treatment. The party cue treatment is illustrated in Figure A1.

Table A1: Policy issues and party cue treatment information.

Policy question	Obama position	Trump position
Should an in-state sales tax apply to online purchases of in-state buyers from out-of-state sellers?	Yes	No
Should pension payments be increased for retired government workers?	Yes	No
Should the Federal Reserve Bank be audited by Congress?	No	Yes
Should the U.S. increase or decrease foreign aid spending?	Increase	Decrease
Do you support a single-payer healthcare system?	Yes	No

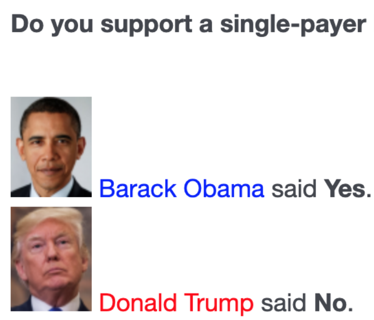

Figure A1: Illustration of party cue treatment.

2 Attrition

As per the main text, it is important to analyze rates of attrition in online survey experiments in general - and ours in particular - in order to diagnose the potential risk of bias in the treatment effects caused by attrition that is differential between treatment and control groups.

Table A2 shows a summary of the results of binomial logistic regression models predicting respondent post-treatment dropout (1 = dropped out, 0 = did not drop out) during the course of the panel survey, as a function of assignment to party cue treatment, for each policy issue in our sample. The results suggest that respondents assigned to the control (no cue) condition on the foreign aid issue were more likely to dropout than those assigned to the party cue on this issue ($p < .001$). The raw numbers are shown in Table A3. While we struggle to imagine a coherent mechanism by which learning the party cue on this particular issue (and not others) should cause respondents to remain in the survey at a higher rate - and, thus, it could be statistical noise - we must nevertheless consider whether this dropout could bias our estimates of party cue persistence. Specifically, this dropout could bias our estimates if policy opinions over foreign aid are correlated with the 'type' of person who drops out if-untreated. If this is the case, then the policy opinions of the control (no cue) and treatment (party cue) groups are no longer balanced in expectation, undermining random assignment and biasing the estimated party cue treatment effect (Gomila & Clark, 2020).

Table A2: Results from models predicting post-treatment dropout from party cue treatment.

item_index	predictor	estimate	std.error	statistic	p.value
Bank audit	cue	0.02	0.22	0.09	.931
Foreign aid	cue	-0.85	0.24	-3.61	< .001
Pension payments	cue	-0.29	0.22	-1.28	.200
Sales tax	cue	0.05	0.22	0.24	.813
Single payer	cue	-0.11	0.22	-0.50	.617

Note:
Estimates are log-odds.

Table A3: Raw numbers of respondents as a function of post-treatment dropout and party cue on foreign aid issue.

Post-treatment dropout	cue	n
0	0	372
0	1	401
1	0	63
1	1	29

We adopt two strategies to mitigate this potential risk of bias. First, in addition to our primary models fit on the initial and follow-up survey data (reported in the main text, and described in the Results section further below), we fit an additional model in each case in which we adjust for a range of pre-treatment covariates that are correlated with policy opinions on the foreign aid issue. Tables A4 and A5 show estimated associations between pre-treatment covariates and policy opinions on foreign aid at initial and follow-up surveys, respectively. Note that policy opinions are recoded such that higher values equal greater alignment with the in-party cue, as per the main text and our primary analyses. The covariates are: gender (1 = female, 0 = not female), education (1 = college graduate, 0 = not college graduate), party ID (1 = Democrat, 0 = not Democrat), strength of party ID (1 = strongest, 0 = weakest), ideology ID (1 = liberal, 0 = not liberal), ideological extremity (1 = most extreme, 0 = least extreme) and self-reported vote choice in the 2020 US presidential election (1 = Joe Biden, 0 = not Joe Biden).

If these covariates sufficiently capture the 'type' of person who drops out if-untreated, we can say that dropout is ignorable *conditional on* these covariates; thus, avoiding a biased treatment effect estimate (Gomila & Clark, 2020). However, this is a strong assumption because we do not have a large number of covariates, and they are measured with some nonzero amount of measurement error. Therefore, our second strategy is to fit yet another model on the initial and follow-up survey data (described further below), this time excluding all policy opinions provided on the foreign aid issue - in either the treatment (party cue) or control (no cue) group. This reduces our statistical power, but more convincingly obviates any bias introduced to our estimated treatment effect by differential attrition on this issue. Reassuringly, both of these strategies yield results that are substantively similar to those from our primary analyses, negating serious concerns about bias.

Table A4: Policy opinions over foreign aid at t0 predicting pre-treatment covariates.

predictor	estimate	std.error	statistic	p.value	covariate_outcome
t0 policy opinions (recoded)	-0.02	0.01	-1.31	.191	female
t0 policy opinions (recoded)	0.00	0.01	0.16	.875	college
t0 policy opinions (recoded)	-0.10	0.01	-8.24	< .001	democrat
t0 policy opinions (recoded)	0.03	0.01	3.35	.001	strength of party ID
t0 policy opinions (recoded)	-0.05	0.01	-3.77	< .001	liberal
t0 policy opinions (recoded)	0.04	0.01	4.13	< .001	ideological extremity
t0 policy opinions (recoded)	-0.04	0.01	-2.93	.004	biden2020

Note:
Policy opinions recoded such that higher values equal greater alignment with in-party cue.

Table A5: Policy opinions over foreign aid at t1 predicting pre-treatment covariates.

predictor	estimate	std.error	statistic	p.value	covariate_outcome
t1 policy opinions (recoded)	-0.02	0.01	-1.74	.083	female
t1 policy opinions (recoded)	0.00	0.01	0.03	.975	college
t1 policy opinions (recoded)	-0.10	0.01	-10.97	< .001	democrat
t1 policy opinions (recoded)	0.02	0.01	2.37	.018	strength of party ID
t1 policy opinions (recoded)	-0.06	0.01	-5.88	< .001	liberal
t1 policy opinions (recoded)	0.02	0.01	3.10	.002	ideological extremity
t1 policy opinions (recoded)	-0.05	0.01	-5.26	< .001	biden2020

Note:
Policy opinions recoded such that higher values equal greater alignment with in-party cue.

Finally, before moving onto the results, Table A6 shows a summary of the results of a binomial logistic regression model predicting respondent post-treatment dropout (1 = dropped out, 0 = did not drop out) during the course of the panel survey, as a function of assignment to opinion-measurement condition (t0-t1 vs. t1-only). There is little evidence consistent with differential dropout by measurement condition ($p > .05$).

Table A6: Results from model predicting post-treatment dropout from opinion-measurement condition.

predictor	estimate	std.error	statistic	p.value
t1_only	-0.22	0.22	-0.98	.326

Note:
Estimates are log-odds.

3 Results

3.1 Estimates at t0

As per the main text, we fitted a Bayesian multilevel model to estimate the average treatment effect (ATE) of the party cue at t0. The outcome variable is designated as Y_i, which corresponds to the observed policy opinion data at t0, re-coded such that larger numbers indicate a more party-consistent response (as described in the main text). We standardize the outcome variable by the mean and SD in the control (no cue) group for modeling. There is a single dummy variable, `cue`, which takes a value of 1 if the ith observation was assigned to the party cue treatment, and 0 otherwise. The model allows the intercept and dummy parameters to vary across *policy issues*, as well as across *respondents*. Thus, we allow the ATE to vary both across policy issues and across respondents. The priors on all parameters are vague and weakly-informative, allowing the data to speak for themselves. The formal specification of the model is as follows, where i indexes observations, J indexes the vector of parameters for policy issues, K indexes the vector of parameters for respondents, and both $\mathbf{R}_J$ and $\mathbf{R}_K$ are 2×2 correlation matrices:

$$
\begin{aligned}
Y_i &\sim \text{Normal}(\mu_i, \sigma) \\
\mu_i &= \phi + \lambda \texttt{cue}_i \\
\phi &= \alpha + \alpha_{J[i]} + \alpha_{K[i]} \\
\lambda &= \beta + \beta_{J[i]} + \beta_{K[i]} \\
\begin{bmatrix} \alpha_J \\ \beta_J \end{bmatrix} &\sim \text{MVNormal}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \mathbf{S}_J \right) \\
\begin{bmatrix} \alpha_K \\ \beta_K \end{bmatrix} &\sim \text{MVNormal}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \mathbf{S}_K \right) \\
\mathbf{S}_J &= \begin{pmatrix} \sigma_{\alpha_J} & 0 \\ 0 & \sigma_{\beta_J} \end{pmatrix} \mathbf{R}_J \begin{pmatrix} \sigma_{\alpha_J} & 0 \\ 0 & \sigma_{\beta_J} \end{pmatrix} \\
\mathbf{S}_K &= \begin{pmatrix} \sigma_{\alpha_K} & 0 \\ 0 & \sigma_{\beta_K} \end{pmatrix} \mathbf{R}_K \begin{pmatrix} \sigma_{\alpha_K} & 0 \\ 0 & \sigma_{\beta_K} \end{pmatrix} \\
\alpha, \beta &\sim \text{Normal}(0, 0.5) \\
\sigma_{\alpha_J}, \sigma_{\beta_J} &\sim \text{Exponential}(5) \\
\sigma_{\alpha_K}, \sigma_{\beta_K} &\sim \text{Exponential}(5) \\
\sigma &\sim \text{Exponential}(1) \\
\mathbf{R}_J, \mathbf{R}_K &\sim \text{LKJcorr}(4).
\end{aligned}
$$

A summary of the model results is shown in Table A7. The corresponding traceplots are shown in Figure A2.

Table A7: Results from ATE-at-t0 model.

Group	Term	Estimate	Est.Error	l-95% CI	u-95% CI	Rhat	Bulk_ESS	Tail_ESS
fixed	Intercept	-0.02	0.16	-0.33	0.29	1.00	3908.00	5029.00
fixed	cue	0.27	0.06	0.15	0.38	1.00	7329.00	4988.00
residual	sigma	0.87	0.02	0.84	0.90	1.00	2973.00	4819.00
item	sd(Intercept)	0.35	0.12	0.19	0.65	1.00	5675.00	5596.00
item	sd(cue)	0.07	0.06	0.00	0.22	1.00	3878.00	4989.00
item	cor(Intercept,cue)	-0.10	0.32	-0.68	0.53	1.00	14875.00	5797.00
pid	sd(Intercept)	0.35	0.04	0.28	0.42	1.00	2443.00	3737.00
pid	sd(cue)	0.15	0.10	0.01	0.37	1.00	589.00	1461.00
pid	cor(Intercept,cue)	-0.09	0.29	-0.57	0.54	1.00	2475.00	4592.00

Note:
Group item = policy issue. Group pid = respondent. ESS = effective sample size.

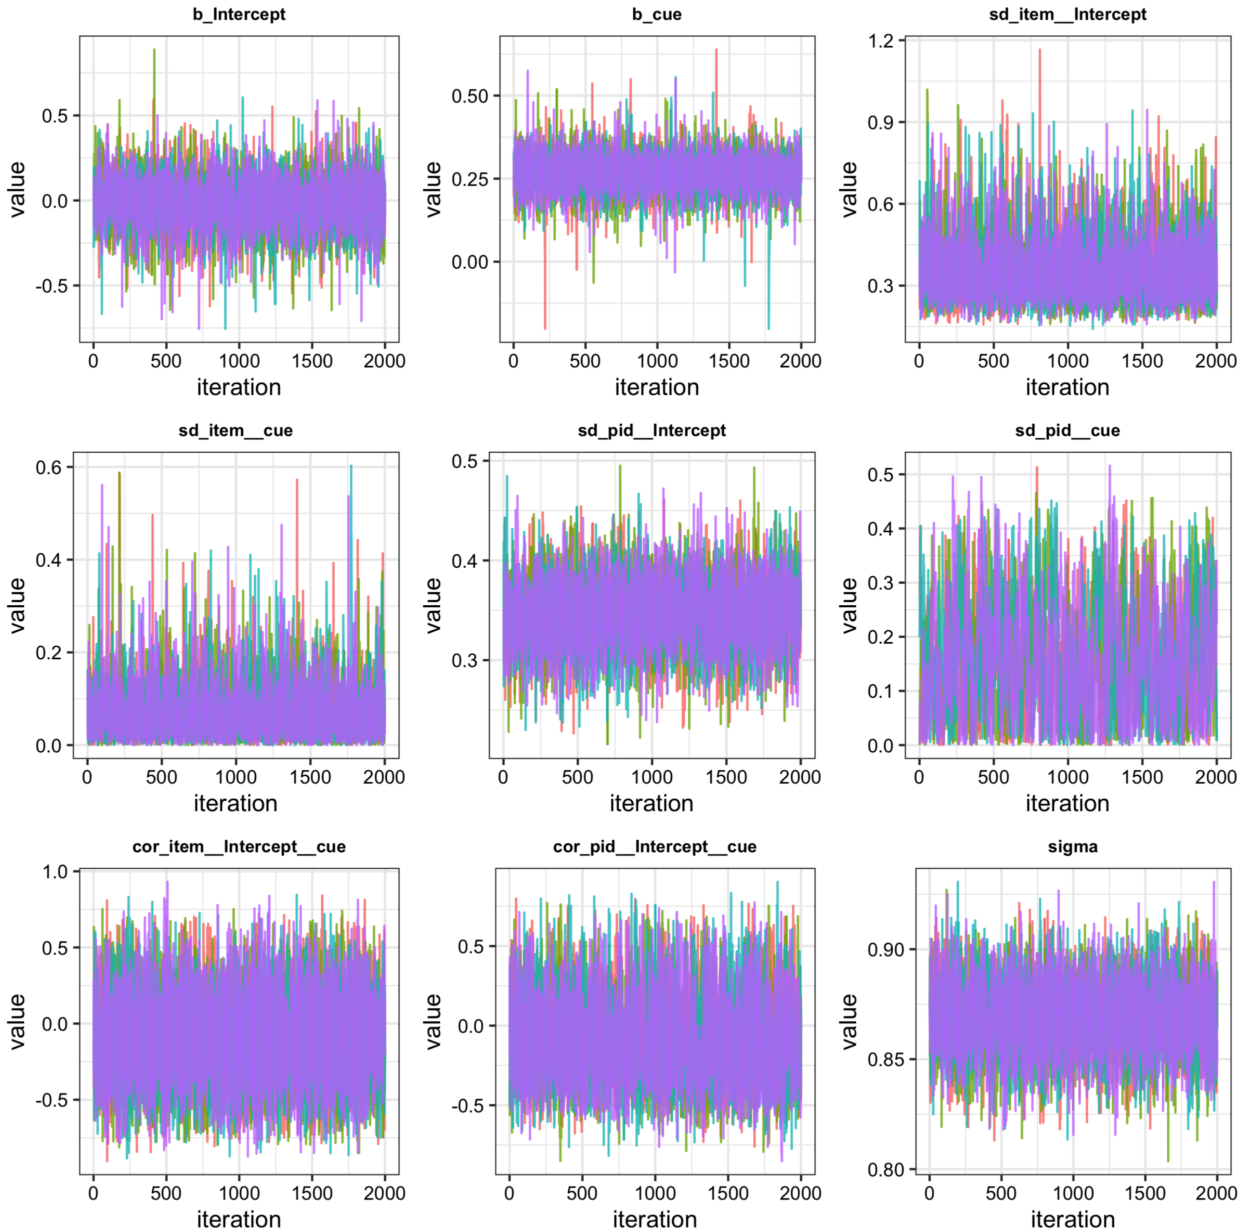

Figure A2: Traceplots for ATE-at-t0 model.

3.2 Estimates at t1

Here we fit a Bayesian multilevel model to estimate the average treatment effect (ATE) of the party cue at t1. The outcome variable is designated as Y_i, which here corresponds to the observed policy opinion data at t1, re-coded such that larger numbers indicate a more party-consistent response (as before). We standardize the outcome variable by the mean and SD in the control (no cue) group for modeling. There are three dummy variables: the first, `cue`, takes a value of 1 if the ith observation was assigned to the party cue treatment, and 0 otherwise. The second, `t1_only`, takes a value of 0.5 if the ith observation was assigned to the t1-only measurement condition, and -0.5 if assigned to the t0-t1 measurement condition (the standard persistence design). Because there are approximately equal numbers of observations in the t1-only and t0-t1 measurement conditions, centering the dummy this way allows us to interpret the party cue ATE averaged across measurement conditions. The third dummy variable is equal to $\texttt{cue} \times \texttt{t1\_only}$, for the interaction between the first two dummies. The parameter on this interaction term tells us whether the ATE is larger (positive values) or smaller (negative values) in the t1-only condition relative to the t0-t1 condition.

The model allows all parameters to vary across *policy issues*, and the intercept and party cue ATE parameters to vary across *respondents*. (Measurement condition is assigned at the level of respondents, and so we do not ask the model to fit separate measurement condition parameters for each respondent.) The priors on all parameters are vague and weakly-informative, allowing the data to speak for themselves. The formal specification of the model is as follows, where i indexes observations, J indexes the vector of parameters for policy issues, K indexes the vector of parameters for respondents, $\mathbf{R}_J$ is a 4×4 correlation matrix and $\mathbf{R}_K$ is a 2×2 correlation matrix:

$$
\begin{aligned}
Y_i &\sim \text{Normal}(\mu_i, \sigma) \\
\mu_i &= \phi + \lambda_1 \texttt{cue}_i + \lambda_2 \texttt{t1\_only}_i + \lambda_3 (\texttt{cue}\times\texttt{t1\_only})_i \\
\phi &= \alpha + \alpha_{J[i]} + \alpha_{K[i]} \\
\lambda_1 &= \beta_1 + \beta_{1J[i]} + \beta_{1K[i]} \\
\lambda_2 &= \beta_2 + \beta_{2J[i]} \\
\lambda_3 &= \beta_3 + \beta_{3J[i]} \\
\begin{bmatrix} \alpha_J \\ \beta_{1J} \\ \beta_{2J} \\ \beta_{3J} \end{bmatrix} &\sim \text{MVNormal}\left(\begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \mathbf{S}_J \right) \\
\begin{bmatrix} \alpha_K \\ \beta_{1K} \end{bmatrix} &\sim \text{MVNormal}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \mathbf{S}_K \right) \\
\mathbf{S}_J &= \begin{pmatrix} \sigma_{\alpha_J} & 0 & 0 & 0 \\ 0 & \sigma_{\beta_{1J}} & 0 & 0 \\ 0 & 0 & \sigma_{\beta_{2J}} & 0 \\ 0 & 0 & 0 & \sigma_{\beta_{3J}} \end{pmatrix} \mathbf{R}_J \begin{pmatrix} \sigma_{\alpha_J} & 0 & 0 & 0 \\ 0 & \sigma_{\beta_{1J}} & 0 & 0 \\ 0 & 0 & \sigma_{\beta_{2J}} & 0 \\ 0 & 0 & 0 & \sigma_{\beta_{3J}} \end{pmatrix} \\
\mathbf{S}_K &= \begin{pmatrix} \sigma_{\alpha_K} & 0 \\ 0 & \sigma_{\beta_{1K}} \end{pmatrix} \mathbf{R}_K \begin{pmatrix} \sigma_{\alpha_K} & 0 \\ 0 & \sigma_{\beta_{1K}} \end{pmatrix} \\
\alpha &\sim \text{Normal}(0, 0.5) \\
\beta_1, \beta_2, \beta_3 &\sim \text{Normal}(0, 0.5) \\
\sigma_{\alpha_J}, \sigma_{\beta_{1J}}, \sigma_{\beta_{2J}}, \sigma_{\beta_{3J}} &\sim \text{Exponential}(5) \\
\sigma_{\alpha_K}, \sigma_{\beta_{1K}} &\sim \text{Exponential}(5) \\
\sigma &\sim \text{Exponential}(1) \\
\mathbf{R}_J, \mathbf{R}_K &\sim \text{LKJcorr}(4).
\end{aligned}
$$

A summary of the model results is shown in Table A8. The corresponding traceplots are shown in Figure A3.

Table A8: Results from ATE-at-t1 model.

Group	Term	Estimate	Est.Error	l-95% CI	u-95% CI	Rhat	Bulk_ESS	Tail_ESS
fixed	Intercept	-0.05	0.16	-0.36	0.27	1.00	1862.00	3025.00
fixed	cue	0.13	0.04	0.05	0.20	1.00	6646.00	4583.00
fixed	t1_only	-0.09	0.06	-0.20	0.02	1.00	5470.00	5061.00
fixed	cue:t1_only	0.08	0.07	-0.07	0.21	1.00	6352.00	5126.00
residual	sigma	0.89	0.01	0.87	0.91	1.00	4730.00	5926.00
item	sd(Intercept)	0.36	0.12	0.19	0.65	1.00	3493.00	5032.00
item	sd(cue)	0.05	0.04	0.00	0.15	1.00	4129.00	4012.00
item	sd(t1_only)	0.05	0.05	0.00	0.18	1.00	3660.00	3931.00
item	sd(cue:t1_only)	0.07	0.07	0.00	0.24	1.00	3942.00	4508.00
item	cor(Intercept,cue)	-0.06	0.29	-0.60	0.52	1.00	8988.00	5603.00
item	cor(Intercept,t1_only)	-0.07	0.29	-0.61	0.51	1.00	9004.00	6285.00
item	cor(cue,t1_only)	0.00	0.30	-0.58	0.58	1.00	8183.00	6234.00
item	cor(Intercept,cue:t1_only)	-0.05	0.29	-0.61	0.52	1.00	10622.00	5537.00
item	cor(cue,cue:t1_only)	-0.00	0.30	-0.58	0.58	1.00	7543.00	5678.00
item	cor(t1_only,cue:t1_only)	-0.01	0.31	-0.60	0.58	1.00	7113.00	5913.00
pid	sd(Intercept)	0.32	0.03	0.26	0.37	1.00	2094.00	3703.00
pid	sd(cue)	0.12	0.07	0.01	0.26	1.01	524.00	1152.00
pid	cor(Intercept,cue)	0.18	0.26	-0.34	0.69	1.00	3550.00	4848.00

Note:

Group item = policy issue. Group pid = respondent. ESS = effective sample size.

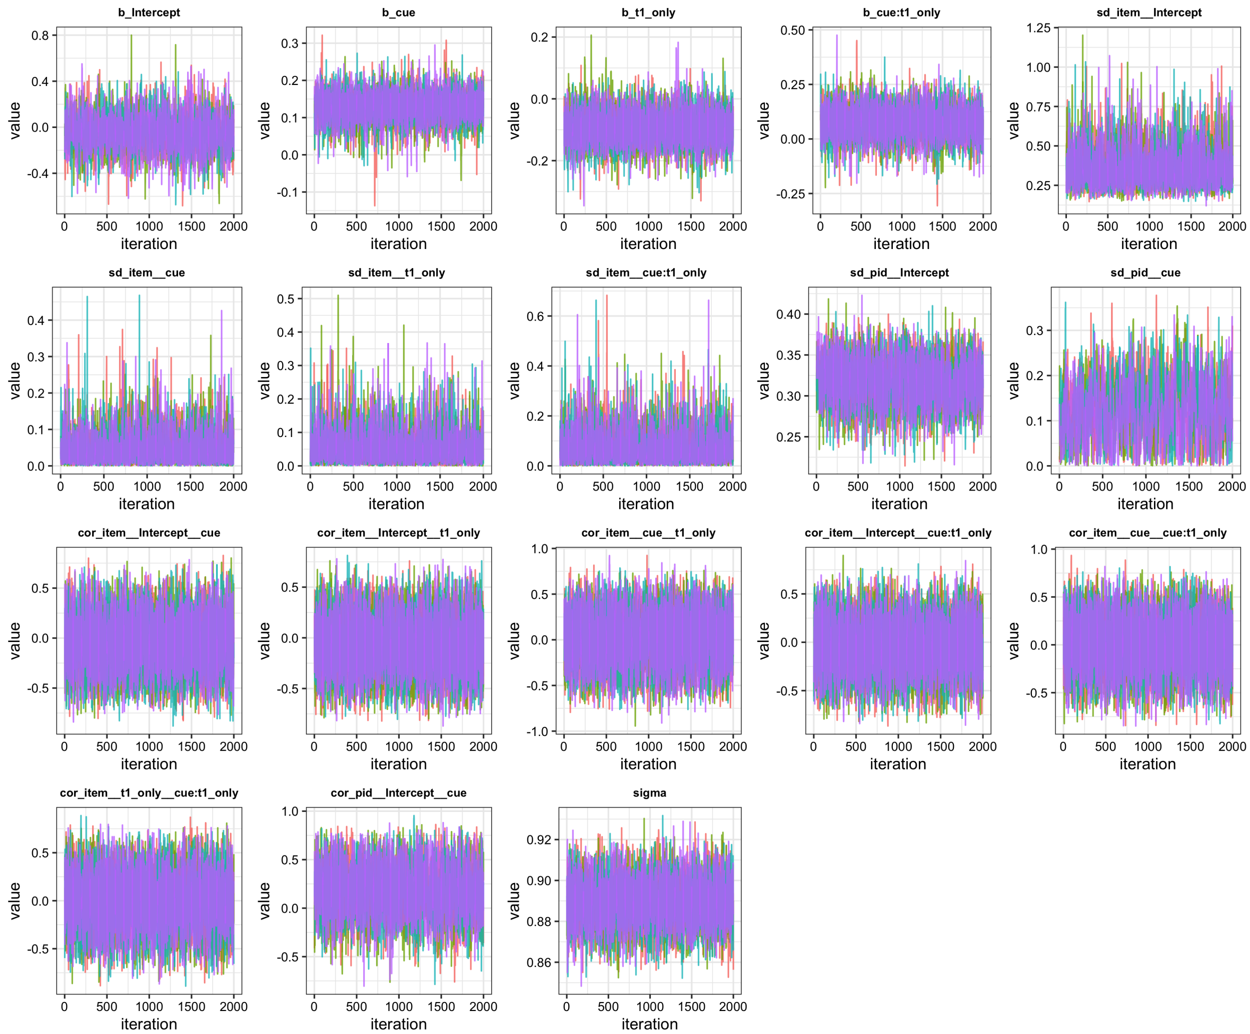

Figure A3: Traceplots for ATE-at-t1 model.

3.3 Robustness to attrition

As per the above section on attrition, we fitted two additional models on both the initial (t0) and follow-up survey (t1) data to mitigate the risk of any bias caused by potential differential dropout on the foreign aid issue. We describe these robustness checks on the estimates at t0 first, and then on the estimates at t1.

3.3.1 Estimates at t0

The **first** additional model adjusts for a range of pre-treatment covariates: gender (1 = female, 0 = not female), education (1 = college graduate, 0 = not college graduate), party ID (1 = Democrat, 0 = not Democrat), strength of party ID (1 = strongest, 0 = weakest), ideology ID (1 = liberal, 0 = not liberal), ideological extremity (1 = most extreme, 0 = least extreme) and self-reported vote choice in the 2020 US presidential election (1 = Joe Biden, 0 = not Joe Biden). These variables are all mean-centered and entered as fixed effects in the primary t0 model specification. The corresponding parameters are assigned a vague and weakly-informative prior distribution of `Normal(0, 0.5)`. A summary of the model results is shown in Table A9. The corresponding traceplots are shown in Figure A4. The results are substantively identical to those of the primary specification.

Table A9: Results from ATE-at-t0 model: with covariate adjustment.

Group	Term	Estimate	Est.Error	l-95% CI	u-95% CI	Rhat	Bulk_ESS	Tail_ESS
fixed	Intercept	-0.01	0.15	-0.32	0.30	1.00	3815.00	4948.00
fixed	cue	0.27	0.06	0.16	0.38	1.00	7975.00	4730.00
fixed	female_c	-0.07	0.05	-0.16	0.02	1.00	11781.00	6369.00
fixed	college_c	0.04	0.05	-0.05	0.14	1.00	12644.00	6326.00
fixed	democrat_c	-0.41	0.09	-0.59	-0.24	1.00	10381.00	6975.00
fixed	party_strength_c	0.11	0.09	-0.06	0.29	1.00	10368.00	6317.00
fixed	liberal_c	-0.15	0.08	-0.31	0.01	1.00	10685.00	6644.00
fixed	ideo_strength_c	0.48	0.10	0.29	0.68	1.00	8680.00	6764.00
fixed	biden2020_c	0.01	0.07	-0.14	0.15	1.00	11971.00	6306.00
residual	sigma	0.87	0.02	0.84	0.90	1.00	3621.00	4464.00
item	sd(Intercept)	0.35	0.12	0.19	0.65	1.00	6388.00	5784.00
item	sd(cue)	0.06	0.05	0.00	0.20	1.00	4354.00	5624.00
item	cor(Intercept,cue)	-0.09	0.33	-0.68	0.56	1.00	16170.00	5391.00
pid	sd(Intercept)	0.23	0.04	0.14	0.31	1.00	2348.00	2818.00
pid	sd(cue)	0.15	0.09	0.01	0.34	1.00	777.00	1839.00
pid	cor(Intercept,cue)	-0.05	0.31	-0.59	0.58	1.00	3449.00	5409.00

Note:
Group item = policy issue. Group pid = respondent. ESS = effective sample size.

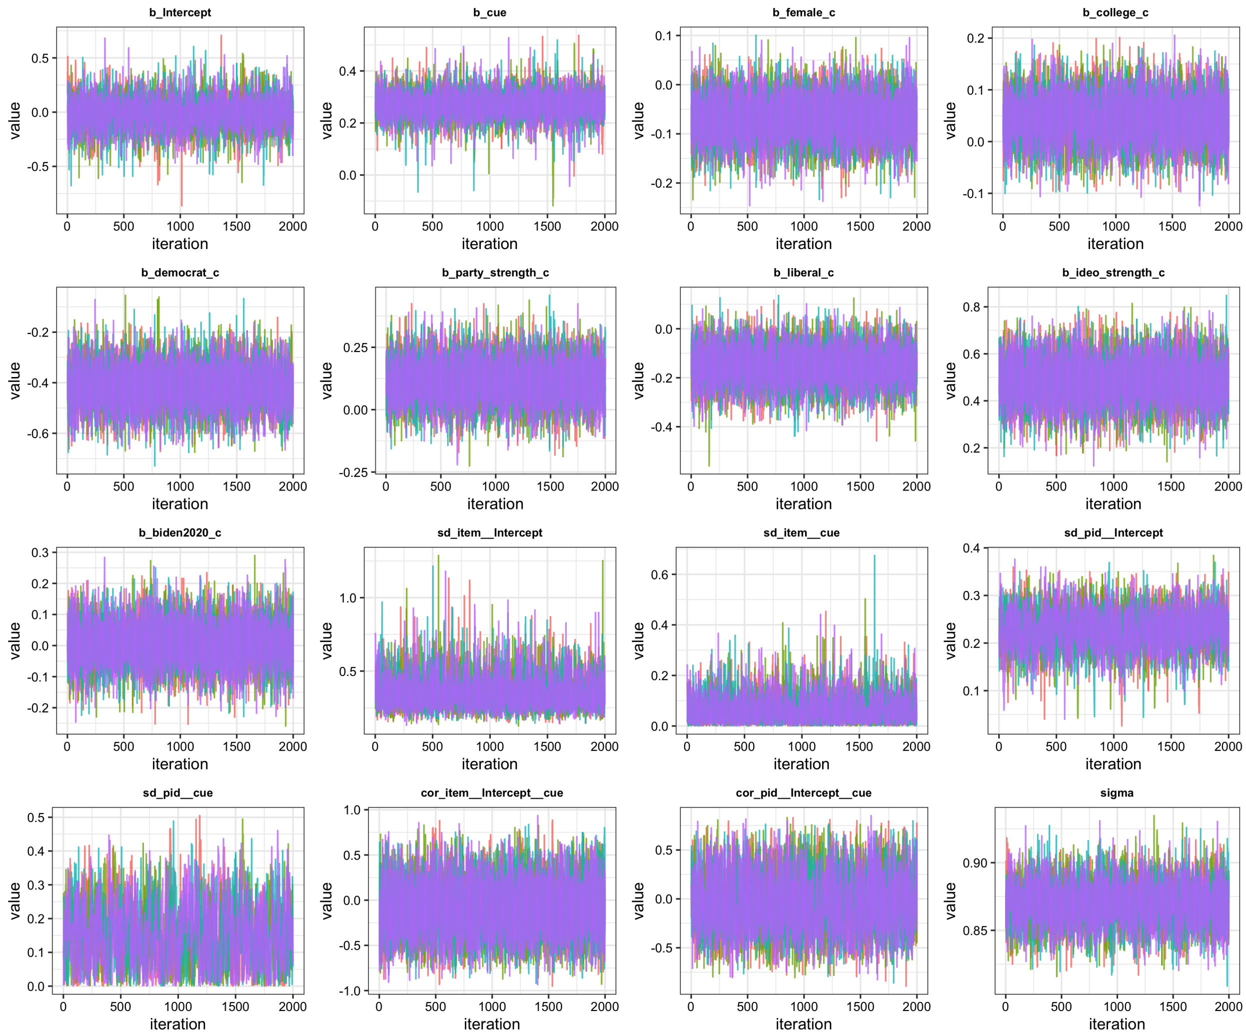

Figure A4: Traceplots for ATE-at-t0 model: with covariate adjustment.

The **second** additional model excludes *all* observations on the foreign aid issue, treatment and control group. (This equates to dropping approximately 400 observations from the model.) Thus, differential dropout on this particular issue can no longer bias the estimated ATE of the party cue at t0. The model specification is identical to the primary t0 model specification, save for dropping these observations. A summary of the model results is shown in Table A10. The corresponding traceplots are shown in Figure A5. The results are substantively identical to those of the primary specification.

Table A10: Results from ATE-at-t0 model: excluding foreign aid issue.

Group	Term	Estimate	Est.Error	l-95% CI	u-95% CI	Rhat	Bulk_ESS	Tail_ESS
fixed	Intercept	-0.01	0.19	-0.40	0.36	1.00	4713.00	5354.00
fixed	cue	0.28	0.07	0.13	0.42	1.00	8036.00	4705.00
residual	sigma	0.90	0.02	0.86	0.94	1.00	3286.00	4370.00
item	sd(Intercept)	0.39	0.14	0.21	0.74	1.00	6552.00	5659.00
item	sd(cue)	0.08	0.08	0.00	0.28	1.00	3583.00	4481.00
item	cor(Intercept,cue)	-0.09	0.32	-0.67	0.56	1.00	17265.00	6125.00
pid	sd(Intercept)	0.24	0.06	0.11	0.34	1.00	1381.00	1168.00
pid	sd(cue)	0.13	0.10	0.00	0.36	1.00	1268.00	2025.00
pid	cor(Intercept,cue)	-0.14	0.33	-0.69	0.52	1.00	4311.00	5788.00

Note:
Group item = policy issue. Group pid = respondent. ESS = effective sample size.

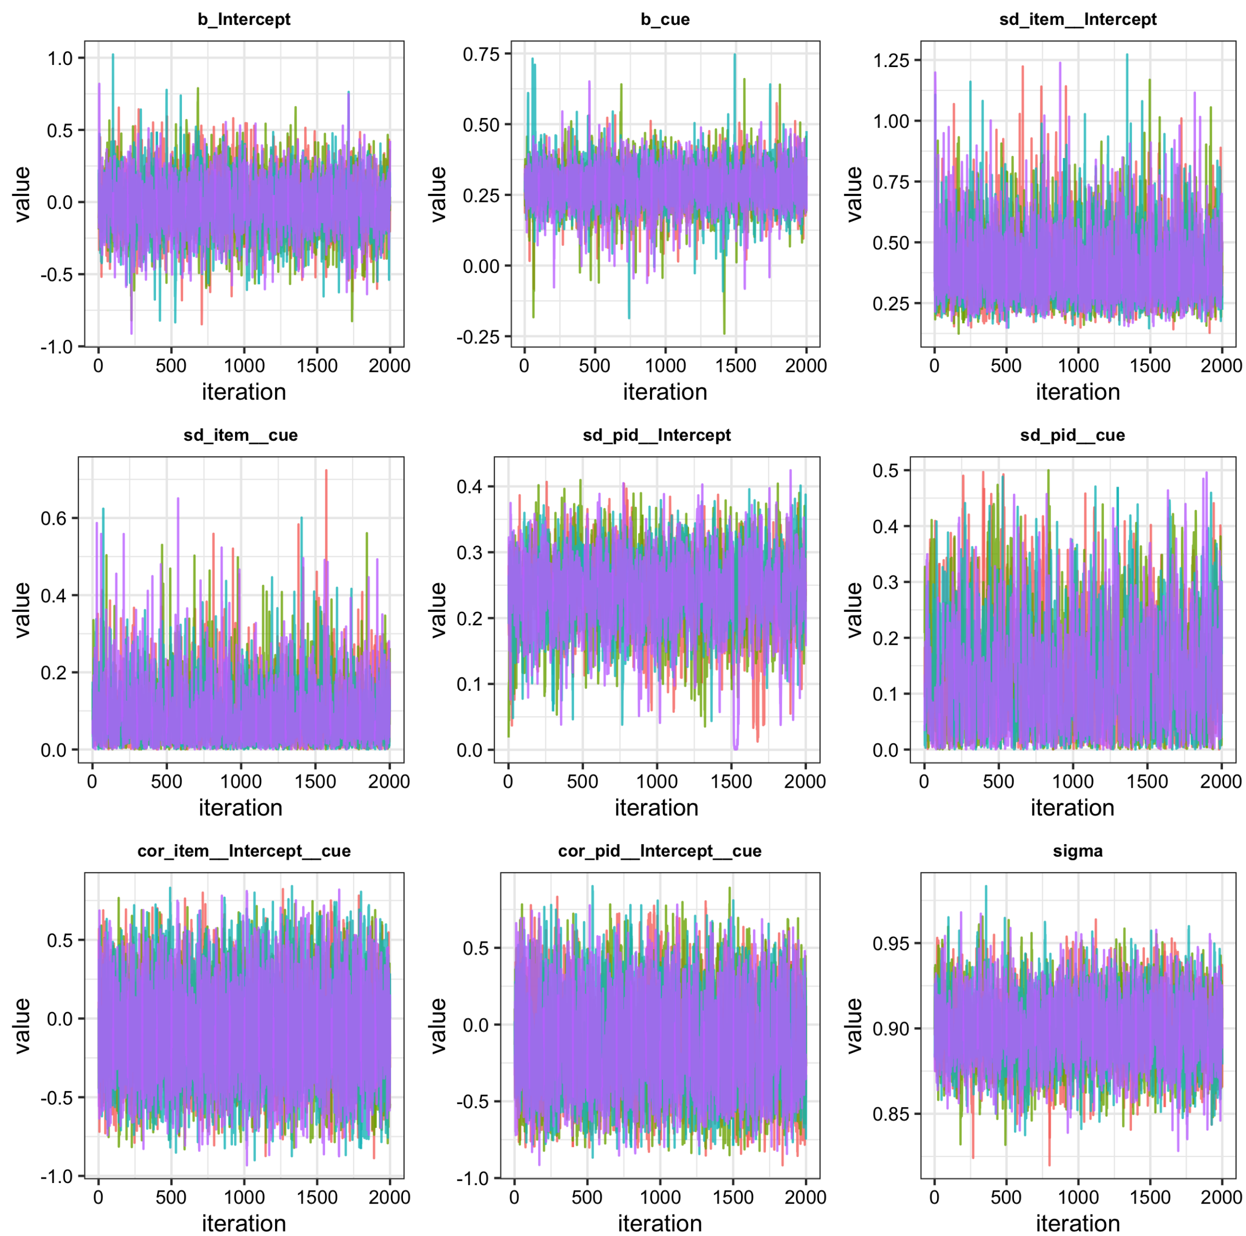

Figure A5: Traceplots for ATE-at-t0 model: excluding foreign aid issue.

3.3.2 Estimates at t1

Here we report the two foregoing analyses on our *t1* primary specification. Table A11 shows a summary of the results from the model with covariate adjustment. The corresponding traceplots are shown in Figure A6. Table 12 shows a summary of the results from the model with observations from the foreign aid issue omitted. The corresponding traceplots are shown in Figure A7. In both cases, the results are substantively identical to those of the primary specification.

Table A11: Results from ATE-at-t1 model: with covariate adjustment.

Group	Term	Estimate	Est.Error	l-95% CI	u-95% CI	Rhat	Bulk_ESS	Tail_ESS
fixed	Intercept	-0.05	0.16	-0.37	0.26	1.00	3048.00	4003.00
fixed	cue	0.13	0.04	0.06	0.21	1.00	8255.00	5373.00
fixed	t1_only	-0.08	0.05	-0.18	0.03	1.00	9006.00	5446.00
fixed	female_c	-0.06	0.03	-0.12	0.00	1.00	11173.00	6528.00
fixed	college_c	0.10	0.03	0.03	0.17	1.00	12081.00	6885.00
fixed	democrat_c	-0.46	0.06	-0.58	-0.34	1.00	8682.00	7013.00
fixed	party_strength_c	0.15	0.07	0.02	0.28	1.00	10034.00	6330.00
fixed	liberal_c	-0.11	0.06	-0.23	0.00	1.00	9636.00	6686.00
fixed	ideo_strength_c	0.37	0.07	0.22	0.51	1.00	9053.00	6557.00
fixed	biden2020_c	0.01	0.05	-0.09	0.11	1.00	8950.00	6062.00
fixed	cue:t1_only	0.06	0.07	-0.08	0.20	1.00	8927.00	5968.00
residual	sigma	0.89	0.01	0.87	0.91	1.00	5607.00	5366.00
item	sd(Intercept)	0.36	0.12	0.20	0.65	1.00	4624.00	5712.00
item	sd(cue)	0.04	0.04	0.00	0.15	1.00	5184.00	4847.00
item	sd(t1_only)	0.05	0.05	0.00	0.18	1.00	4287.00	4460.00
item	sd(cue:t1_only)	0.07	0.06	0.00	0.24	1.00	4333.00	4792.00
item	cor(Intercept,cue)	-0.05	0.30	-0.60	0.53	1.00	15374.00	5282.00
item	cor(Intercept,t1_only)	-0.08	0.29	-0.62	0.48	1.00	16616.00	6202.00
item	cor(cue,t1_only)	0.01	0.31	-0.58	0.59	1.00	12479.00	6473.00
item	cor(Intercept,cue:t1_only)	-0.04	0.29	-0.59	0.52	1.00	18063.00	6369.00
item	cor(cue,cue:t1_only)	-0.00	0.30	-0.58	0.56	1.00	11308.00	6146.00
item	cor(t1_only,cue:t1_only)	-0.01	0.30	-0.58	0.57	1.00	8861.00	6574.00
pid	sd(Intercept)	0.21	0.03	0.13	0.27	1.00	1905.00	2972.00
pid	sd(cue)	0.12	0.06	0.01	0.24	1.01	815.00	2087.00
pid	cor(Intercept,cue)	0.17	0.28	-0.41	0.69	1.00	5035.00	4640.00

Note:
Group item = policy issue. Group pid = respondent. ESS = effective sample size.

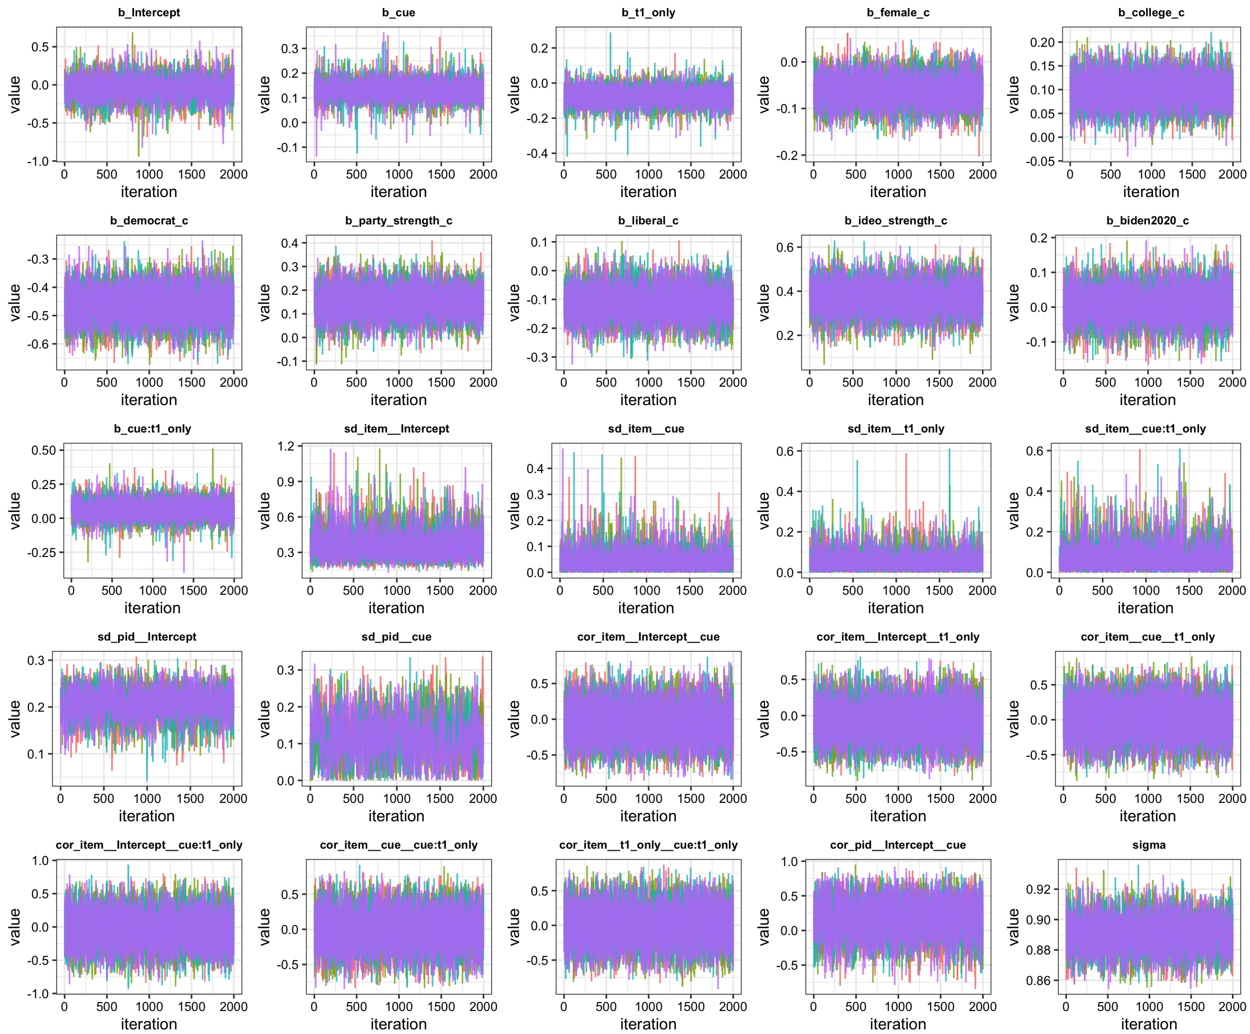

Figure A6: Traceplots for ATE-at-t1 model: with covariate adjustment.

Table A12: Results from ATE-at-t1 model: excluding foreign aid issue.

Group	Term	Estimate	Est.Error	l-95% CI	u-95% CI	Rhat	Bulk_ESS	Tail_ESS
fixed	Intercept	-0.05	0.18	-0.42	0.32	1.00	2613.00	4323.00
fixed	cue	0.14	0.06	0.03	0.25	1.00	5355.00	4106.00
fixed	t1_only	-0.10	0.07	-0.23	0.04	1.00	6875.00	5087.00
fixed	cue:t1_only	0.09	0.09	-0.08	0.26	1.00	7877.00	5541.00
residual	sigma	0.92	0.01	0.89	0.95	1.00	3688.00	4685.00
item	sd(Intercept)	0.39	0.13	0.21	0.74	1.00	4918.00	5362.00
item	sd(cue)	0.06	0.06	0.00	0.22	1.00	3612.00	4387.00
item	sd(t1_only)	0.07	0.06	0.00	0.24	1.00	4373.00	4994.00
item	sd(cue:t1_only)	0.09	0.08	0.00	0.30	1.00	4309.00	4516.00
item	cor(Intercept,cue)	-0.07	0.29	-0.61	0.50	1.00	13105.00	5621.00
item	cor(Intercept,t1_only)	-0.06	0.29	-0.62	0.50	1.00	13814.00	5692.00
item	cor(cue,t1_only)	0.01	0.30	-0.57	0.57	1.00	10730.00	5792.00
item	cor(Intercept,cue:t1_only)	-0.04	0.29	-0.58	0.52	1.00	14510.00	6036.00
item	cor(cue,cue:t1_only)	-0.00	0.30	-0.58	0.57	1.00	10906.00	6491.00
item	cor(t1_only,cue:t1_only)	0.00	0.30	-0.58	0.58	1.00	9737.00	5808.00
pid	sd(Intercept)	0.20	0.04	0.10	0.27	1.00	1278.00	1096.00
pid	sd(cue)	0.10	0.07	0.00	0.25	1.00	858.00	1580.00
pid	cor(Intercept,cue)	0.11	0.31	-0.50	0.68	1.00	5382.00	5024.00

Note:
Group item = policy issue. Group pid = respondent. ESS = effective sample size.

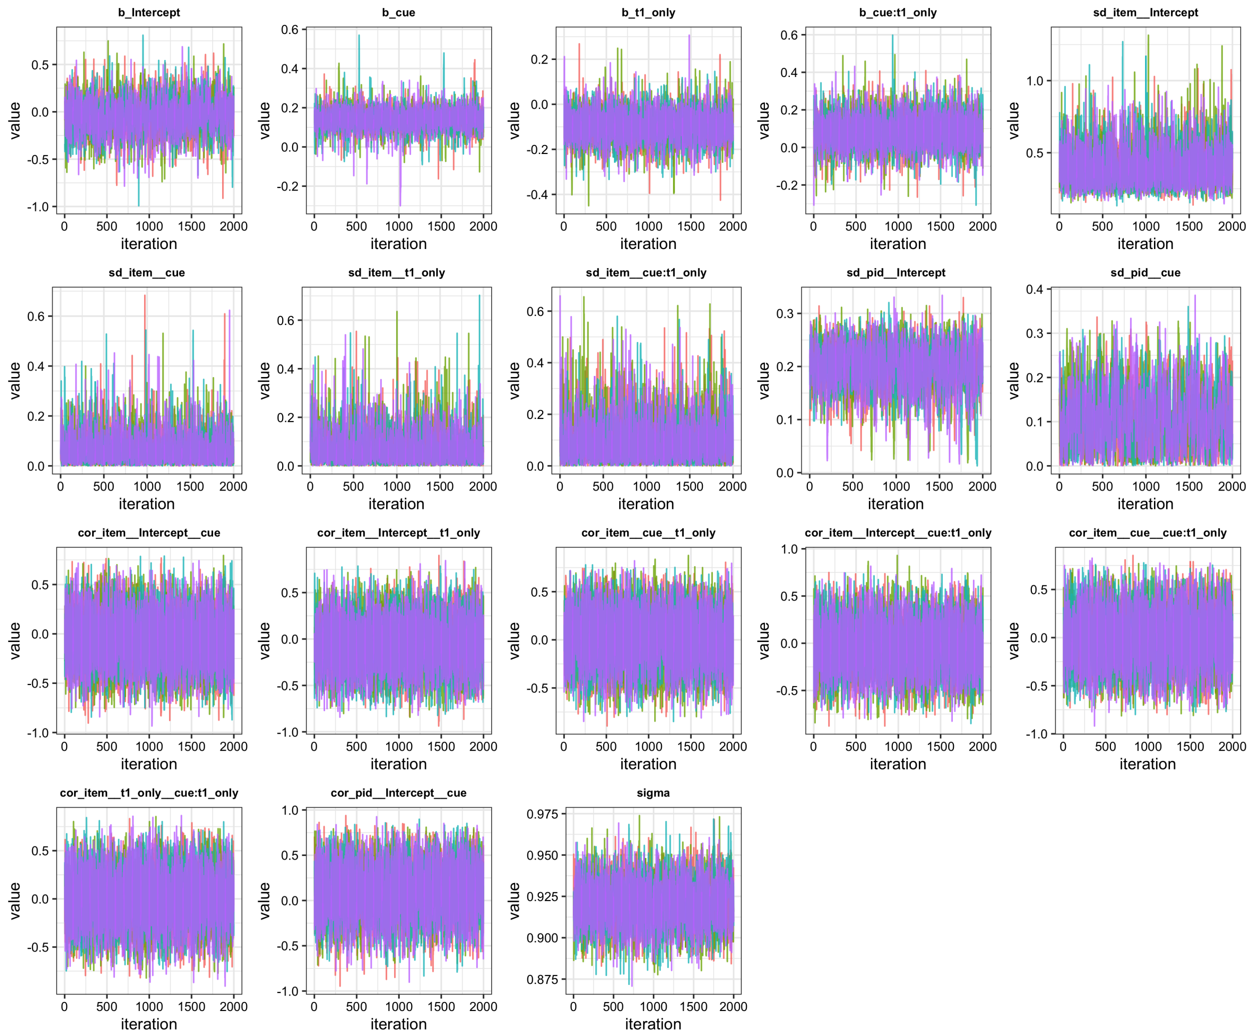

Figure A7: Traceplots for ATE-at-t1 model: excluding foreign aid issue.

3.4 Models excluding political Independents

As per the main text, our primary model specifications included political Independents. Thus, below we report the results of models excluding respondents who identified as Independents - those who selected 4 on the 1-7 party ID scale. As the tables below show, at both t0 and t1 the key estimates are substantively identical to those from our primary specifications.

3.4.1 Estimates at t0

A summary of the model results is shown in Table A13, and the corresponding traceplots are shown in Figure A8.

Table A13: Results from ATE-at-t0 model: excluding Independents.

Group	Term	Estimate	Est.Error	l-95% CI	u-95% CI	Rhat	Bulk_ESS	Tail_ESS
fixed	Intercept	-0.02	0.17	-0.39	0.32	1.00	2070.00	3798.00
fixed	cue	0.32	0.07	0.18	0.44	1.00	5317.00	4875.00
residual	sigma	0.84	0.02	0.81	0.88	1.00	2864.00	4828.00
item	sd(Intercept)	0.39	0.13	0.22	0.71	1.00	3772.00	4770.00
item	sd(cue)	0.09	0.07	0.00	0.27	1.00	3259.00	3319.00
item	cor(Intercept,cue)	-0.16	0.32	-0.73	0.49	1.00	7350.00	5522.00
pid	sd(Intercept)	0.32	0.04	0.24	0.40	1.00	2830.00	4601.00
pid	sd(cue)	0.18	0.11	0.01	0.39	1.01	405.00	1499.00
pid	cor(Intercept,cue)	-0.04	0.29	-0.55	0.56	1.00	2018.00	3913.00

Note:
Group item = policy issue. Group pid = respondent. ESS = effective sample size.

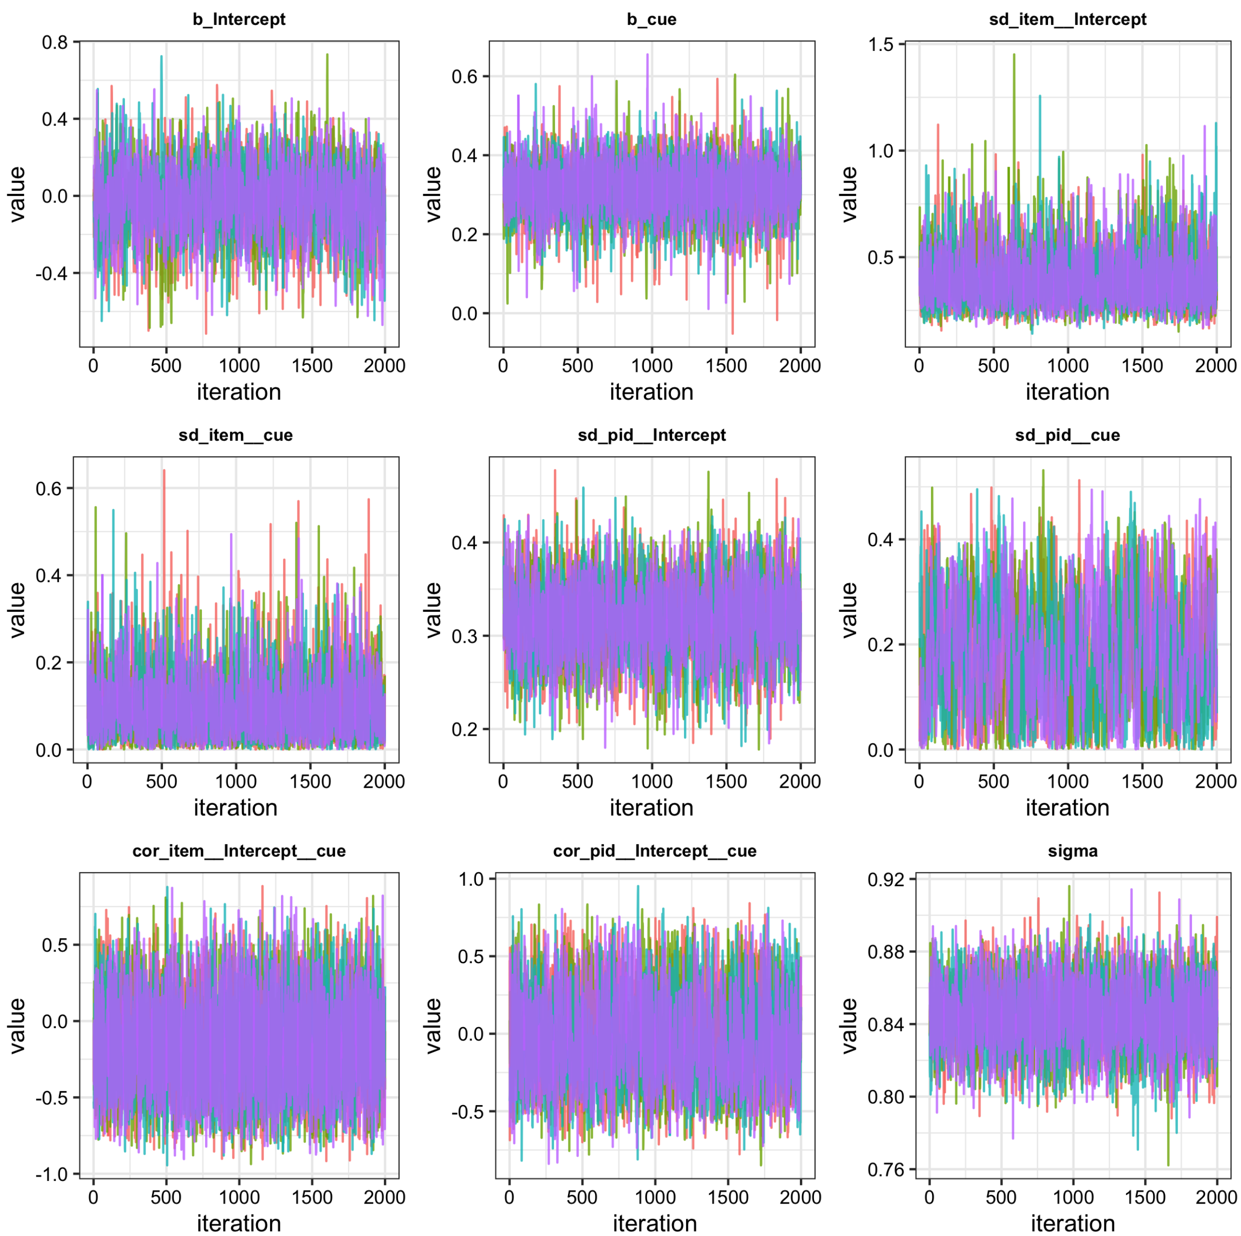

Figure A8: Traceplots for ATE-at-t0 model: excluding Independents.

3.4.2 Estimates at t1

A summary of the model results is shown in Table A14, and the corresponding traceplots are shown in Figure A9.

Table A14: Results from ATE-at-t1 model: excluding Independents.

Group	Term	Estimate	Est.Error	l-95% CI	u-95% CI	Rhat	Bulk_ESS	Tail_ESS
fixed	Intercept	-0.05	0.17	-0.39	0.29	1.00	2748.00	3922.00
fixed	cue	0.14	0.05	0.05	0.23	1.00	7246.00	4784.00
fixed	t1_only	-0.10	0.06	-0.22	0.02	1.00	7225.00	5791.00
fixed	cue:t1_only	0.06	0.08	-0.09	0.21	1.00	8627.00	5714.00
residual	sigma	0.88	0.01	0.85	0.90	1.00	5727.00	5854.00
item	sd(Intercept)	0.40	0.13	0.23	0.71	1.00	4599.00	5237.00
item	sd(cue)	0.06	0.05	0.00	0.18	1.00	3945.00	4458.00
item	sd(t1_only)	0.05	0.05	0.00	0.17	1.00	5420.00	5215.00
item	sd(cue:t1_only)	0.06	0.06	0.00	0.21	1.00	5091.00	4450.00
item	cor(Intercept,cue)	-0.10	0.29	-0.64	0.48	1.00	14911.00	6000.00
item	cor(Intercept,t1_only)	-0.08	0.29	-0.63	0.51	1.00	13906.00	5427.00
item	cor(cue,t1_only)	0.01	0.30	-0.57	0.58	1.00	11692.00	5743.00
item	cor(Intercept,cue:t1_only)	-0.01	0.29	-0.57	0.55	1.00	16487.00	5960.00
item	cor(cue,cue:t1_only)	-0.00	0.30	-0.58	0.57	1.00	11791.00	6226.00
item	cor(t1_only,cue:t1_only)	-0.02	0.30	-0.59	0.56	1.00	9593.00	6007.00
pid	sd(Intercept)	0.30	0.03	0.23	0.35	1.00	2647.00	3739.00
pid	sd(cue)	0.12	0.07	0.01	0.26	1.00	701.00	1787.00
pid	cor(Intercept,cue)	0.17	0.28	-0.40	0.69	1.00	4864.00	5188.00

Note:
Group item = policy issue. Group pid = respondent. ESS = effective sample size.

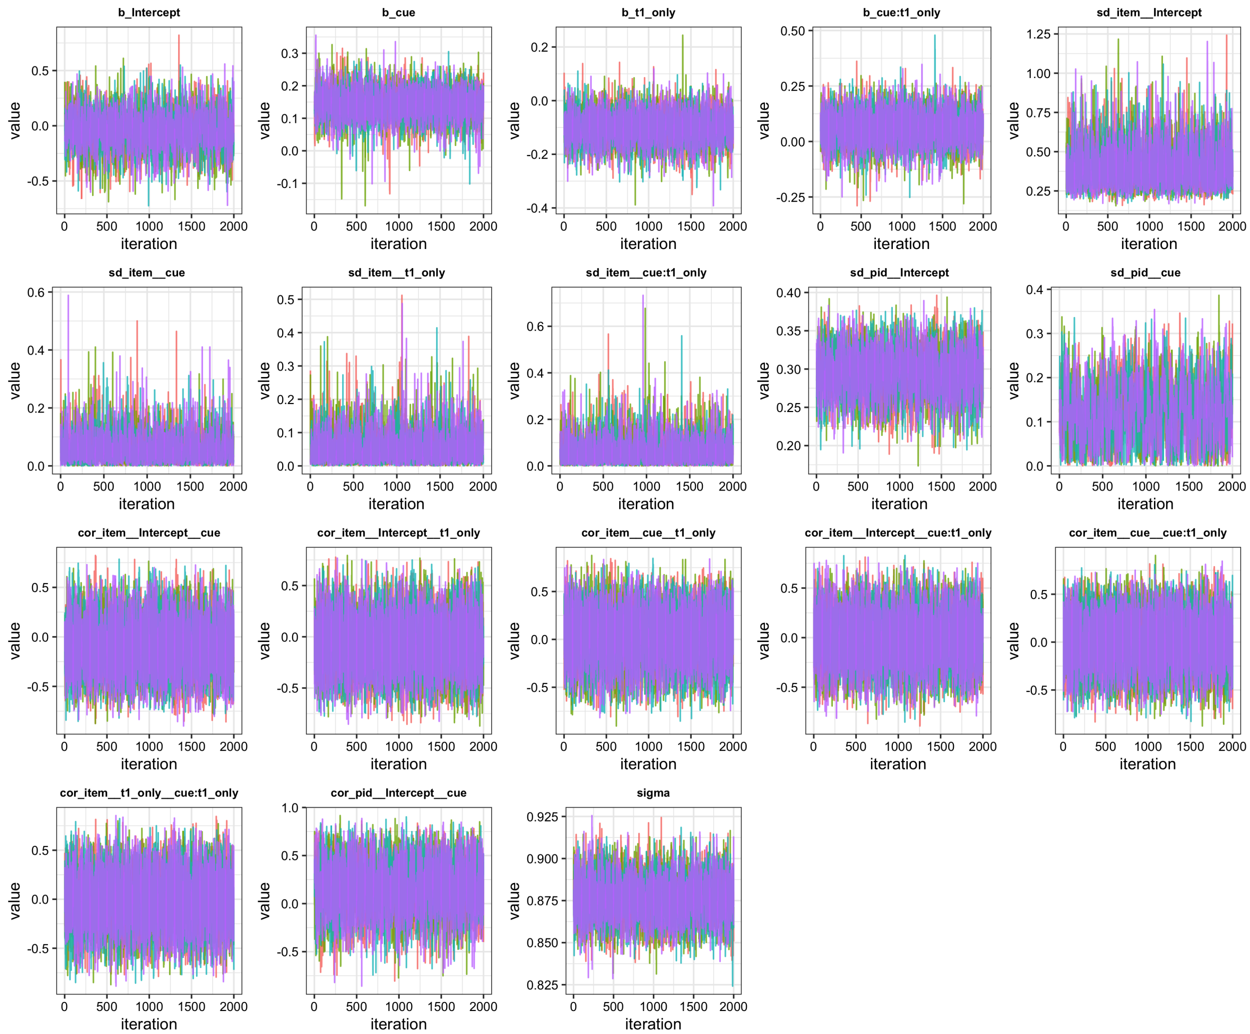

Figure A9: Traceplots for ATE-at-t1 model: excluding Independents.

3.5 Models separated by partisanship

Figure A10 shows the standardized estimates from models fitted separately to respondents who identified with the Democratic and Republican Party. The Republican estimates are much noisier because there are significantly fewer Republicans in the sample. Despite this, the pattern is roughly symmetric: the party cue ATE is attenuated by approximately 50% by follow-up among both Democrats and Republicans.

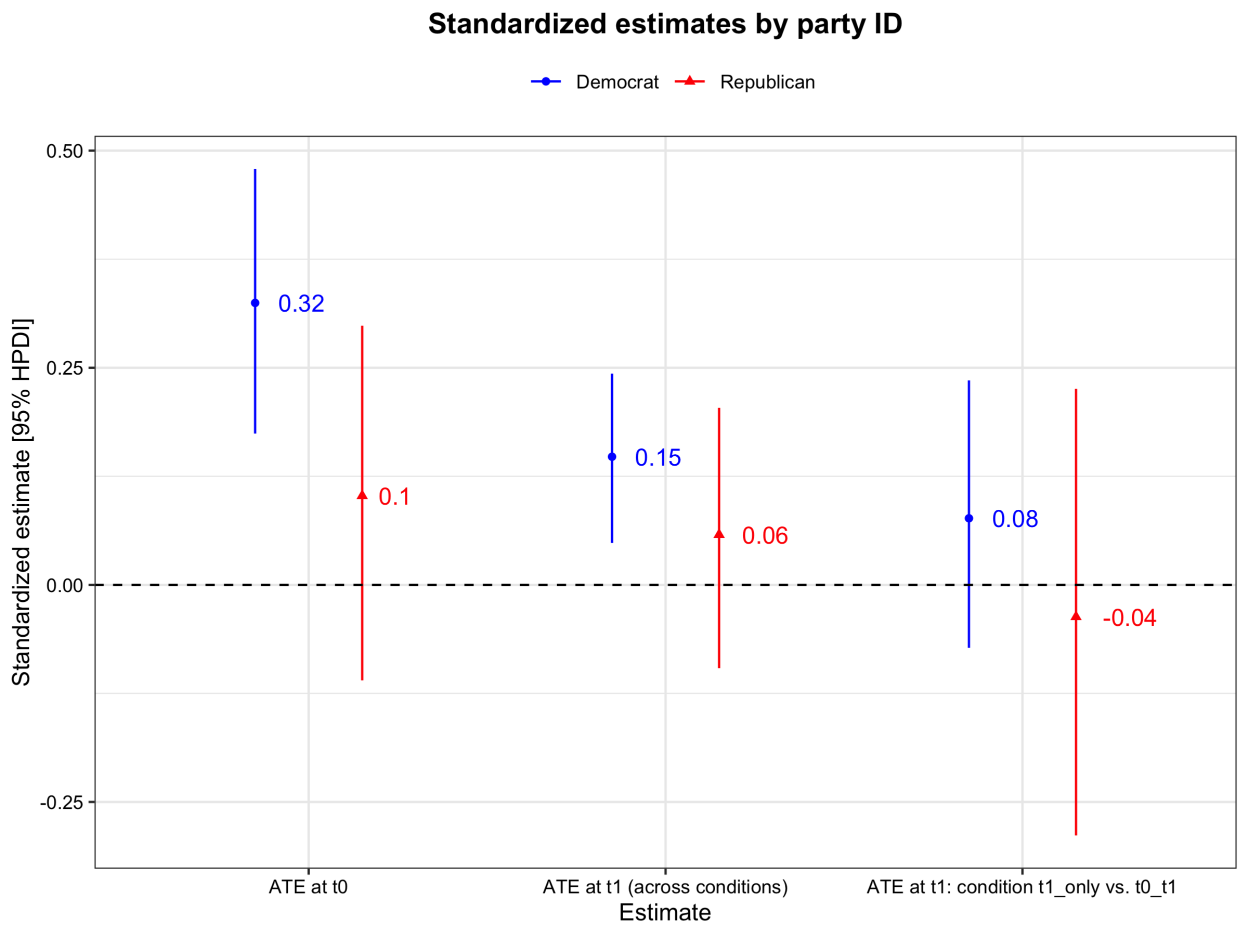

Figure A10: Standardized estimates by partisanship.

3.6 Persistence ratio

As per the main text, we also estimated a formal ratio of the party cue ATE at t1 vs. at t0. To estimate this ratio, we require a model that simultaneously estimates the party cue ATE at both t0 and t1. The posterior distribution of each of the ATE parameters can then easily be transformed into a ratio and summarized. Accordingly, we fit a new model to the full dataset (i.e., all of the data, from both t0 and t1).

The outcome variable is designated as Y_i, which corresponds to the observed policy opinion data, re-coded such that larger numbers indicate a more party-consistent response (as before). We standardize the outcome variable by the mean and SD in the control (no cue) group at wave t0. There are three dummy variables: the first, `cue`, takes a value of 1 if the ith observation was assigned to the party cue treatment at t0, and 0 otherwise. Thus, the parameter on this dummy, λ_1, tells us the party cue ATE at t0. The second dummy variable, `wave_t1`, takes a value of 0 if the ith observation is from wave t0, and a value of 1 if it is from wave t1. The third dummy variable is equal to `cue` $\times$ `wave_t1`, for the interaction between the first two dummies. The parameter on this interaction term, λ_3, therefore tells us whether the ATE is larger (positive values) or smaller (negative values) at t1 relative to t0. To estimate the persistence ratio we simply compute $\frac{\lambda_1+\lambda_3}{\lambda_1}$.

The model allows all parameters to vary across *policy issues* and *respondents*. The priors on all parameters are vague and weakly-informative, allowing the data to speak for themselves. The formal specification of the model is as follows, where i indexes observations, J indexes the vector of parameters for policy issues, K indexes the vector of parameters for respondents, and $\mathbf{R}_J$ and $\mathbf{R}_K$ are 4×4 correlation matrices:

$$
\begin{aligned}
Y_i &\sim \text{Normal}(\mu_i, \sigma) \\
\mu_i &= \phi + \lambda_1 \texttt{cue}_i + \lambda_2 \texttt{wave\_t1}_i + \lambda_3 \texttt{(cue}\times\texttt{wave\_t1)}_i \\
\phi &= \alpha + \alpha_{J[i]} + \alpha_{K[i]} \\
\lambda_1 &= \beta_1 + \beta_{1J[i]} + \beta_{1K[i]} \\
\lambda_2 &= \beta_2 + \beta_{2J[i]} + \beta_{2K[i]} \\
\lambda_3 &= \beta_3 + \beta_{3J[i]} + \beta_{3K[i]} \\
\begin{bmatrix} \alpha_J \\ \beta_{1J} \\ \beta_{2J} \\ \beta_{3J} \end{bmatrix} &\sim \text{MVNormal}\left(\begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \mathbf{S}_J \right) \\
\begin{bmatrix} \alpha_K \\ \beta_{1K} \\ \beta_{2K} \\ \beta_{3K} \end{bmatrix} &\sim \text{MVNormal}\left(\begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \mathbf{S}_K \right) \\
\mathbf{S}_J &= \begin{pmatrix} \sigma_{\alpha_J} & 0 & 0 & 0 \\ 0 & \sigma_{\beta_{1J}} & 0 & 0 \\ 0 & 0 & \sigma_{\beta_{2J}} & 0 \\ 0 & 0 & 0 & \sigma_{\beta_{3J}} \end{pmatrix} \mathbf{R}_J \begin{pmatrix} \sigma_{\alpha_J} & 0 & 0 & 0 \\ 0 & \sigma_{\beta_{1J}} & 0 & 0 \\ 0 & 0 & \sigma_{\beta_{2J}} & 0 \\ 0 & 0 & 0 & \sigma_{\beta_{3J}} \end{pmatrix} \\
\mathbf{S}_K &= \begin{pmatrix} \sigma_{\alpha_K} & 0 & 0 & 0 \\ 0 & \sigma_{\beta_{1K}} & 0 & 0 \\ 0 & 0 & \sigma_{\beta_{2K}} & 0 \\ 0 & 0 & 0 & \sigma_{\beta_{3K}} \end{pmatrix} \mathbf{R}_K \begin{pmatrix} \sigma_{\alpha_K} & 0 & 0 & 0 \\ 0 & \sigma_{\beta_{1K}} & 0 & 0 \\ 0 & 0 & \sigma_{\beta_{2K}} & 0 \\ 0 & 0 & 0 & \sigma_{\beta_{3K}} \end{pmatrix} \\
\alpha &\sim \text{Normal}(0, 0.5) \\
\beta_1, \beta_2, \beta_3 &\sim \text{Normal}(0, 0.5) \\
\sigma_{\alpha_J}, \sigma_{\beta_{1J}}, \sigma_{\beta_{2J}}, \sigma_{\beta_{3J}} &\sim \text{Exponential}(5) \\
\sigma_{\alpha_K}, \sigma_{\beta_{1K}}, \sigma_{\beta_{2K}}, \sigma_{\beta_{3K}} &\sim \text{Exponential}(5) \\
\sigma &\sim \text{Exponential}(1) \\
\mathbf{R}_J, \mathbf{R}_K &\sim \text{LKJcorr}(4).
\end{aligned}
$$

A summary of the model results is shown in Table A15.

Table A15: Results from persistence ratio model.

Group	Term	Estimate	Est.Error	l-95% CI	u-95% CI	Rhat	Bulk_ESS	Tail_ESS
fixed	Intercept	-0.03	0.16	-0.35	0.30	1.00	1679.00	2232.00
fixed	cue	0.27	0.06	0.16	0.38	1.00	3260.00	4140.00
fixed	wave_t1	-0.01	0.04	-0.09	0.07	1.00	5123.00	5390.00
fixed	cue:wave_t1	-0.14	0.06	-0.25	-0.03	1.00	4908.00	5549.00
residual	sigma	0.83	0.01	0.81	0.85	1.00	6224.00	5670.00
item	sd(Intercept)	0.36	0.12	0.20	0.66	1.00	3030.00	4322.00
item	sd(cue)	0.06	0.05	0.00	0.19	1.00	3055.00	3126.00
item	sd(wave_t1)	0.03	0.03	0.00	0.11	1.00	4627.00	3570.00
item	sd(cue:wave_t1)	0.04	0.04	0.00	0.14	1.00	4701.00	4810.00
item	cor(Intercept,cue)	-0.10	0.29	-0.63	0.48	1.00	8767.00	5633.00
item	cor(Intercept,wave_t1)	0.03	0.29	-0.54	0.57	1.00	9773.00	6073.00
item	cor(cue,wave_t1)	0.01	0.31	-0.58	0.59	1.00	8969.00	6006.00
item	cor(Intercept,cue:wave_t1)	-0.02	0.30	-0.60	0.56	1.00	9718.00	5743.00
item	cor(cue,cue:wave_t1)	-0.01	0.31	-0.59	0.57	1.00	8856.00	6250.00
item	cor(wave_t1,cue:wave_t1)	-0.02	0.31	-0.60	0.57	1.00	7763.00	6334.00
pid	sd(Intercept)	0.46	0.02	0.42	0.51	1.00	2756.00	4528.00
pid	sd(cue)	0.51	0.03	0.44	0.58	1.00	1764.00	3919.00
pid	sd(wave_t1)	0.02	0.02	0.00	0.06	1.00	2989.00	2922.00
pid	sd(cue:wave_t1)	0.03	0.03	0.00	0.10	1.00	2255.00	3353.00
pid	cor(Intercept,cue)	-0.46	0.06	-0.58	-0.34	1.00	2152.00	3779.00
pid	cor(Intercept,wave_t1)	-0.02	0.30	-0.59	0.57	1.00	8599.00	5809.00
pid	cor(cue,wave_t1)	-0.02	0.30	-0.58	0.54	1.00	8554.00	6129.00
pid	cor(Intercept,cue:wave_t1)	0.06	0.30	-0.53	0.61	1.00	9052.00	5727.00
pid	cor(cue,cue:wave_t1)	-0.08	0.30	-0.64	0.52	1.00	9018.00	6422.00
pid	cor(wave_t1,cue:wave_t1)	-0.03	0.31	-0.61	0.55	1.00	6567.00	6169.00

Note:
Group item = policy issue. Group pid = respondent. ESS = effective sample size.

4 Sensitivity power analysis

As per the main text, we conducted a post-hoc sensitivity power analysis to assess the size of party cue effects that our sample, design and analytic approach is equipped to reliably detect. For simplicity, we conduct this power analysis in a frequentist (not Bayesian) multilevel framework. To preview its results, we find that, under reasonable assumptions about the data generating parameters, our sample and design can detect a standardized party cue ATE of somewhere between 0.15 and 0.20 with >80% power. In other words, we can detect effect sizes conventionally considered small with over 80% statistical power. The code to reproduce the power analysis can be found online, and we describe our assumptions in words here.

There are nine key parameters in the simulation:

- Number of respondents
- Number of policy issues
- Population average intercept (ϕ)
- Population average treatment effect (λ)
- Standard deviation in the outcome variable (σ)
- Standard deviation in the intercept across *respondents* (σ_{α_K})
- Standard deviation in the treatment effect across *respondents* (σ_{β_K})
- Standard deviation in the intercept across *policy issues* (σ_{α_J})
- Standard deviation in the treatment effect across *policy issues* (σ_{β_J})

The number of respondents and policy issues are fixed by our sample and design to 773 (respondents) and 5 (issues), respectively. The population average intercept is equal to the mean in the control (no cue) group in our design. Thus, we set this parameter, ϕ, to zero, because we standardized the outcome variable by the control group for our primary empirical analysis. We set λ, the population average treatment effect, to 0.2, assuming an overall true effect of exposure to the party cue on opinions that is conventionally considered small; and we set the residual standard deviation in the outcome variable (σ) to 1. Because σ is approximately the standard deviation of the outcome variable, λ is interpretable as a "standardized effect size".

The remaining four parameters govern the assumed true variation in the intercept and treatment effect across respondents and policy issues, respectively. To set these parameter values, we take our cue (pun intended) from recent work that sought to estimate the variation in party cue effects across policy issues, using a design similar to ours (Tappin 2020; Tappin & McKay 2021). Though said work did not analyze a standardized outcome variable - putting their estimates of these four parameters on a different absolute scale than ours - we can nevertheless infer plausible values for the parameters by relative comparisons between them. For example, in this previous work the standard deviation in the party cue treatment effect across policy issues is such that, for the large majority of policy issues, the treatment effect is expected to be positive. This matches our intuition: in-party cues very rarely cause opinion updating in the *opposite* direction (a negative treatment effect). Thus, given our assumption of a population average treatment effect of 0.2, we set the SD in this effect across policy issues to 0.1 ($\sigma_{\beta_J} = 0.1$). This implies that party cues are expected to have an effect in the range of 0.2 $\pm$ 0.2 (i.e., the mean $\pm$ 2 SDs) for the majority of policy issues; satisfying the condition that the effect is rarely negative, but may nevertheless sometimes be minimal or even zero.

Moving onto σ_{α_J}, the standard deviation in the intercept across policy issues. In the previous work, this parameter is typically larger than σ_{β_J}. This makes sense: we expect variation across policy issues in the extent to which the average opinion is party-consistent (at baseline) to probably be greater than variation in the average *increase* in party-consistent opinions under exposure to the party cue. For example, on some policy issues voters' baseline opinions may be strongly inconsistent with their in-party leaders', while for other issues the two are closely aligned. In the previous work, σ_{α_J} is as much as three times larger than σ_{β_J}. Thus, we set σ_{α_J} to 0.3 (i.e., $\sigma_{\beta_J}\times 3$). This leaves the standard deviations in the intercept and treatment effect across *respondents*; σ_{α_K} and σ_{β_K}, respectively. We set σ_{α_K} to 0.4, since it is typically slightly larger than the equivalent parameter for policy issues, which we set above to 0.3. We set σ_{β_K} to 0.2, which assumes that the standard deviation in treatment effects across respondents is larger than that across policy issues. While this doesn't follow directly from previous work (where it is typically smaller than that across policy issues), it is implied by the primary empirical analysis of this paper. Furthermore, erring on the side of a larger SD actually makes our power estimate more conservative, because larger between-unit variation in the parameter of interest typically inflates the standard error on the population parameter of interest (Yarkoni 2020); in this case, the party cue average treatment effect, λ.

Finally, for computational simplicity, we omit the parameter that captures the correlations between intercept and treatment effect within each population of respondents and policy issues. We are comfortable doing this because said parameter is highly unlikely to substantively affect any of the simulation results. We run the simulation 2000 times, and confirm that it recovers the parameter values we set above. The key simulation result is that our statistical power on a party cue average treatment effect (λ) of 0.2 is approximately 90%.

5 References

Gomila, R., & Clark, C. S. (2020). "Missing data in experiments: Challenges and solutions." *Psychological Methods*. https://psyarxiv.com/mxenv

Tappin, Ben M. 2020. "Estimating the Between-Issue Variation in Party Elite Cue Effects." *PsyArXiv*. preprint. https://osf.io/p48zb (October 8, 2020).

Tappin, Ben M., and Ryan McKay. 2021. "Estimating the Causal Effects of Cognitive Effort and Policy Information on Party Cue Influence." https://psyarxiv.com/tdk3y/ (February 14, 2021).

Yarkoni, T. 2020. "The generalizability crisis." *Behavioral and Brain Sciences*.