

When Democracy Is Worth the Effort: Seeing the Other Side as Willing Strengthens Pro-Democratic Motivation

Olaf Borghi^{1,2*}, Manos Tsakiris^{1,2}, Mariana von Mohr^{1,2,3}

¹Department of Psychology, Royal Holloway, University of London

²Centre for the Politics of Feelings, School of Advanced Study, University of London

³School of Psychology, University of Kent

*Correspondence concerning this article should be addressed to Olaf Borghi, Department of Psychology, Royal Holloway, University of London, Egham Hill, Egham, Surrey TW20 0EX, United Kingdom.

Email: olaf.borghi@rhul.ac.uk

Author Note

Olaf Borghi  <https://orcid.org/0000-0001-9724-2839>

Manos Tsakiris  <https://orcid.org/0000-0001-7753-7576>

Mariana von Mohr  <https://orcid.org/0000-0003-0671-0735>

Data, analysis code, and materials for both studies are openly available on the Open Science Framework at <https://doi.org/10.17605/OSF.IO/F6CGS>. Study 1 (<https://doi.org/10.17605/OSF.IO/ZYEVV>) and Study 2 (<https://doi.org/10.17605/OSF.IO/KXP2R>) were preregistered before data collection. A preprint of this manuscript is available at https://doi.org/10.31234/osf.io/9jf8s_v2.

This research was conducted as part of the Interdisciplinary Perspectives on the Politics of Adolescence and Democracy (IP-PAD) Doctoral Network. IP-PAD is funded by the European Union under the Horizon Europe MSCA Doctoral Networks programme (Call HORIZON-MSCA-2021-DN-01, Grant Agreement No. 101072992), with participation from the UKRI Horizon Europe guarantee scheme (UKRI EPSRC EP/X026825/1). Manos Tsakiris is also supported by a NOMIS Foundation Grant for the Centre for the Politics of Feelings. The funders had no role in study design, data collection, analysis, interpretation, or the decision to submit for publication.

We have no known conflicts of interest to disclose.

We thank Jo Cutler for helpful comments on the computational models, and Ryan McKay, Luis Sebastian Contreras-Huerta, and Ryan Ratnam for valuable feedback on an early draft of the manuscript. Their input does not imply endorsement of the views expressed here, and any remaining errors are our own.

Abstract

Democracy takes effort, and these efforts do not always benefit one's own political side. Across two preregistered studies (total N = 916), we used computational modeling of effortful behavior to measure US partisans' motivation to strengthen democratic engagement in rural (traditionally Republican-leaning), urban (traditionally Democratic-leaning), or all US areas. Study 1 showed that participants were generally willing to make pro-democratic efforts, but less so when efforts benefited out-group areas, particularly among more affectively polarized participants. We hypothesized that this asymmetry is sustained by misperceptions about out-partisans' willingness to engage in cross-partisan pro-democratic efforts. Study 2 showed that most participants (87%) underestimated out-party effort willingness. Importantly, correcting this misperception increased motivation to support out-group areas without significantly reducing in-group motivation. These findings suggest that cross-partisan pro-democratic motivation is common but partly obscured by misperceptions about out-partisans' willingness, and that making this willingness visible might help sustain the collective effort democracy depends on.

Public Significance Statement

Democracies are most resilient when citizens support them regardless of which political side benefits. Across two preregistered studies, we found that US partisans were often willing to exert effort to strengthen democratic participation, but less so when doing so benefited areas linked to the opposing party. This reluctance was partly sustained by misperceptions: most partisans underestimated how willing their political opponents were to make pro-democratic efforts across party lines. Correcting this misperception increased motivation to support areas linked to the opposing party without weakening support for one's own side. These findings show that partisans are willing to make costly efforts for democracy even across divides, and that making this willingness visible may help sustain the collective effort democracy depends on.

Keywords: democracy, effort, computational modeling, misperceptions, partisanship

Introduction

Democracy takes effort. Engaging in civic life, going to vote, staying informed, and helping others participate in democratic life all require time and effort, both mental and physical. But democratic effort often does not selectively benefit oneself or one's political group. Mobilizing voter engagement in rural as opposed to urban areas, for example, may strengthen representation and thereby increase the legitimacy of democratic outcomes, while also disproportionately benefiting more conservative candidates. Democracies are most resilient when citizens support them even when doing so benefits the other side (Levitsky & Ziblatt, 2018). Yet when democratic support becomes conditional on who benefits, norms can erode, as each side can come to view its own withdrawal as a justified response to the other's (Braley et al., 2023). For whom, then, and under what circumstances, is democracy worth the effort?

Neuroeconomic models of effort-based decision making formalize when and why people choose to exert effort. Humans generally avoid effort, and the subjective value of an action declines as effort increases (Kool & Botvinick, 2018; Kurzban, 2016; Westbrook & Braver, 2015; Zipf, 1949). This trade-off can be computationally modeled to distinguish motivation from decision noise (Contreras-Huerta et al., 2020; Lockwood et al., 2017). Critically, who benefits from an action further shapes this trade-off. Effort discounting—whereby people view exertion of effort as a cost, reducing motivation for demanding tasks—is often steeper when actions benefit others rather than oneself (Contreras-Huerta et al., 2020; Cutler et al., 2025; Depow et al., 2022; Lockwood et al., 2017, 2022). In politics, the question of who benefits is especially consequential, as actions often carry clear implications for which groups stand to gain or lose (Nadeau et al., 2023). Citizens may thus be less motivated to support democracy when efforts benefit both sides equally or the other side more, making the conditionality of pro-democratic effort a potential vulnerability of democratic systems.

One factor that likely sustains this conditionality is misperceptions of the other side. Partisans consistently overestimate how extreme, hostile, and antidemocratic their opponents are (Braley et al., 2023; Druckman et al., 2022; Levendusky & Malhotra, 2016; Mernyk et al., 2022; Pasek et al., 2022), and those who feel most negatively toward the out-party tend to hold the strongest misperceptions and show the greatest in-group favoritism (Lees & Cikara, 2020; Moore-Berg et al., 2020). Correcting misperceptions can increase support for democratic norms and reduce tolerance for undemocratic practices (Braley et al., 2023; Mernyk et al., 2022; Voelkel et al., 2024). Yet it remains unclear whether such corrections extend to costly, effort-based forms of *actual pro-democratic behavior*, and whether they can

shift motivation to support democratic engagement when the beneficiaries are one's political opponents rather than one's own side. This question is particularly important for highly polarized societies where citizens may hold particularly strong misperceptions and may be the least willing to incur effort costs for their political out-group or for democracy at large. It is precisely in such polarized contexts that there may be more room for correction of such misperceptions and the greatest need for cross-partisan democratic support.

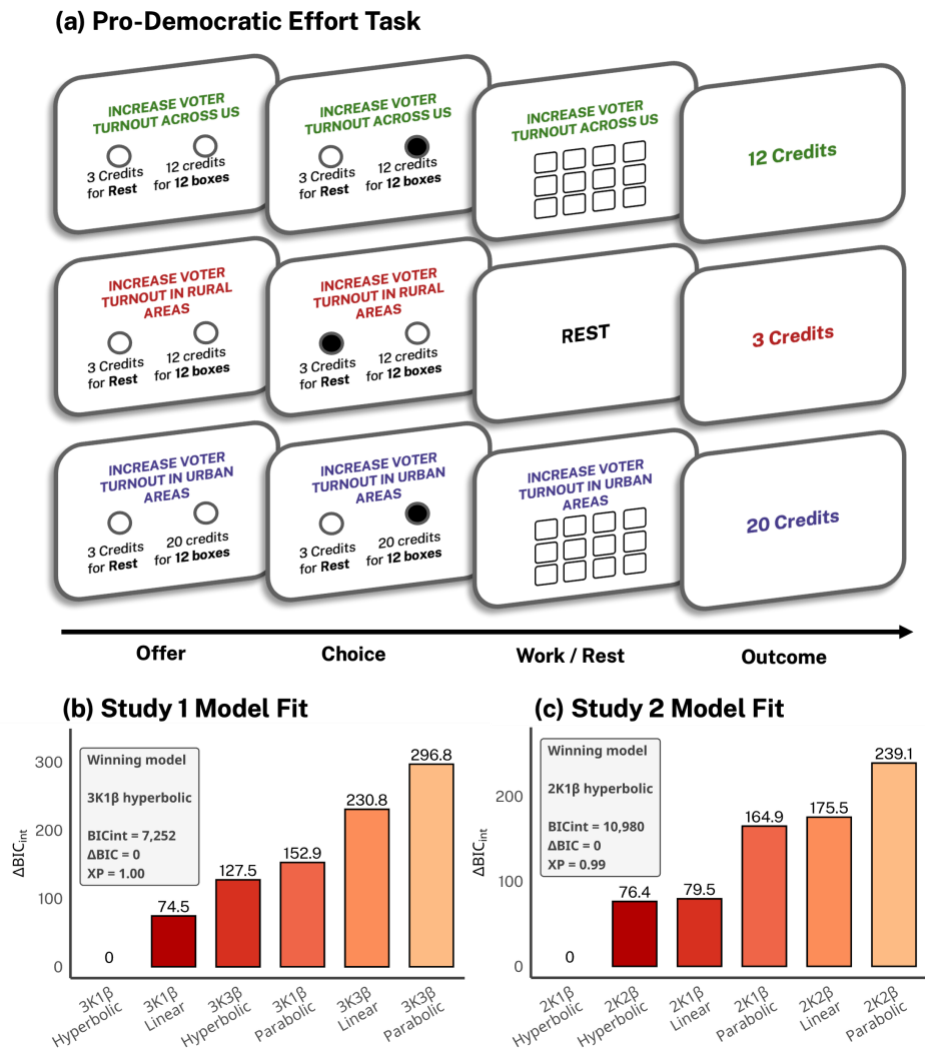
We therefore ask two key questions. First, how much and under what circumstances are partisans willing to exert effort for democracy? Specifically, we look at how much effort people are willing to exert when society as a whole benefits, compared with when their own party or their out-party stands to gain more from these efforts. We hypothesize that citizens are more willing to make pro-democratic efforts when their supported party benefits from them, and less willing when they benefit everyone equally or disproportionately benefit their out-party. We expect this in-group favoritism to be stronger among more affectively polarized citizens, who feel very positive about their own party and negative about their out-party. Second, we hypothesize that such in-group favoritism is, at least in part, sustained by underestimations of out-party supporters' willingness to make pro-democratic efforts for the other side. We therefore ask whether correcting these misperceptions can increase the motivation to exert effort in support of democracy beyond one's own side, helping to break a cycle in which each side's belief that the other will not contribute justifies its own withdrawal.

To address these questions, we focused on the highly polarized two-party context of US politics and introduced the Pro-Democratic Effort Task. This behavioral paradigm, adapted from established effort tasks (Cutler et al., 2025; Lockwood et al., 2017), measures willingness to exert physical effort to support democratic engagement under different partisan beneficiary conditions (Figure 1). On each trial, participants chose between resting for a small fixed reward or working for larger rewards at varying effort levels. Participants were informed that rewards benefited democratic engagement programs that, while non-partisan and aimed at strengthening civic participation among US citizens, targeted different areas: urban areas, described as having historically leaned more Democratic; rural areas, described as having historically leaned more Republican; or all areas across the US. Participants were thus explicitly informed of the rural–urban partisan divide, in which rural areas have become increasingly aligned with Republicans and urban areas with Democrats (Bishop & Cushing, 2008; Mettler & Brown, 2022; Nadeem, 2024), but also of the broader democratic benefits that their efforts would carry. Critically, this design allowed us to fit computational models of

effort discounting to the choice data, distinguishing variation in latent motivation from variation in decision noise (Contreras-Huerta et al., 2020; Westbrook & Braver, 2015).

In Study 1 we validated the task and examined partisan differences in pro-democratic effort. We estimated beneficiary-specific effort-discounting parameters (K) for each geographical condition, where higher values indicate steeper discounting of reward by effort and thus lower motivation. We then recoded these parameters relative to participants' partisanship into motivation for the in-group (urban for Democrats; rural for Republicans), the out-group (rural for Democrats; urban for Republicans), and everyone (all areas). We tested the preregistered hypothesis that participants would show greater motivation to exert effort for in-group than out-group areas and examined political and psychological correlates of pro-democratic effort willingness.

In Study 2 we used a between-subjects experiment to test whether correcting misperceptions about out-party effort willingness could reduce in-group favoritism in pro-democratic effort. Before completing the task, participants estimated how willing out-party supporters had been to exert effort for areas associated with their own party (e.g., Democrats estimated how often the average Republican had chosen to work for urban areas and vice versa). They were then randomly assigned to receive accurate feedback showing that the average out-party member had in fact often chosen to exert such effort (on about 60% of trials, based on our Study 1), or no feedback. We then tested participants' willingness to make efforts to engage in-group or out-group areas using the Pro-Democratic Effort Task.

Figure 1*Pro-Democratic Effort Task and Computational Modeling of Effort Discounting*

Note. The task structure is shown in (a). On each trial, participants chose between a low-reward rest option (3 credits) and higher-effort work options yielding higher rewards (4, 12, or 20 credits). The effort level varied with the number of boxes to be clicked (65%, 80%, or 95% of an individually calibrated maximum). Trials were framed by beneficiary condition (Vote All, Vote Rural, or Vote Urban); Study 2 used the same task without the vote-all condition. Model comparisons for Study 1 and Study 2 are shown in (b) and (c), respectively, as ΔBIC_{int} relative to the best-fitting model, for candidate models varying in discounting function (linear, hyperbolic, parabolic) and parameterization.

Method

Participants

We conducted two preregistered studies. Both studies used stratified quotas to approximate representativeness for US Democrats and Republicans, matching the 2022 American National Election Studies pilot study (American National Election Studies, 2022) on gender and ethnicity across age groups within each party.

Study 1

We recruited 300 US participants (150 Democrats, 150 Republicans). Following our preregistered exclusion criteria, we excluded participants who failed at least one of two embedded attention checks ($n = 10$), missed more than half of the trials in the Pro-Democratic Effort Task ($n = 3$), or identified as Independent despite pre-screening ($n = 1$). While not preregistered, we additionally had to remove two duplicate responses from the same participants, for which we retained only the first submission. This yielded a final sample of $N = 284$ (138 Democrats, 146 Republicans; $M_{\text{age}} = 52.3$ years, $SD = 15.9$, range 19–87; 60.6% women; 84.2% White). The median study duration was 22.3 min, and participants were paid £3 (approximately \$4). Full demographics are reported in Table S2.

Study 2

We recruited 667 US participants (330 Democrats, 330 Republicans, and 7 over-recruits), aiming for an analyzable sample of approximately $N = 600$ after exclusions. Following our preregistered exclusion criteria, we excluded participants who indicated they would not pay attention ($n = 3$), failed a pre-treatment attention check ($n = 2$), missed more than half of the trials in the Pro-Democratic Effort Task ($n = 15$), or identified as Independent or another party despite pre-screening ($n = 15$). This yielded a final sample of $N = 632$ (326 Democrats, 306 Republicans; $M_{\text{age}} = 51.0$ years, $SD = 16.3$, range 18–91; 62.0% women; 82.6% White). The median study duration was 15.5 min, and participants were paid £2 (approximately \$2.65). Full demographics are reported in Tables S3–S4.

Sample Size Rationale

Sample sizes for both studies were determined via preregistered power simulations using the DeclareDesign package (Blair et al., 2019), targeting at least 80% power ($\alpha = .05$, two-sided) after anticipated exclusions. For Study 1, simulations indicated that $N = 200$ would achieve 80% power to detect small-to-medium within-subject differences of interest of Cohen's $d = 0.20$ and correlations of $r = .20$ between effort parameters and survey measures. We therefore preregistered to recruit $N = 300$, yielding approximately 90% power and ensuring the analyzable sample would remain well above this threshold. For Study 2,

simulations indicated that $N = 500$ would achieve approximately 80% power to detect a between-subjects average treatment effect of Cohen's $d = 0.25$, the smallest effect of interest based on a previous study with similar design (Voelkel et al., 2024). We therefore preregistered to recruit $N = 660$, yielding well above 80% power after anticipated exclusions.

Procedure

Participants for both studies were recruited via Prolific and completed the study online in Qualtrics. In Study 1, participants completed the Pro-Democratic Effort Task (see Figure 1), followed by a battery of survey measures including political variables, mental and physical health indicators, and demographics. Participants were told that their effort in the task generated credits that would translate into real donations, but, as in similar effort paradigms (Cutler et al., 2025), they were not informed of the rate at which credits converted into money. Because no real organizations target voter engagement in only urban or only rural areas, the task contained a necessary element of deception, and we decided not to vary participants' real payment by task performance or beneficiary. All participants therefore received the same (maximum) bonus of \$0.20. This was not disclosed until after the task, when participants were informed of the bonus and given a choice to either donate it to Vote.org (a non-partisan voter registration organization supporting engagement across both rural and urban areas) or keep it as a bonus payment via Prolific, in addition to the £3 baseline compensation. These choices were honored and bonuses were either paid out or donated as each participant decided.

In Study 2, participants were randomly assigned to either a misperception correction condition or a control condition using balanced randomization within each partisan group, implemented in Qualtrics. Participants first completed pre-treatment survey measures including demographics and political variables. Before completing the Pro-Democratic Effort Task, they estimated how willing supporters of the opposing party were to exert effort for areas associated with their own party (e.g., Democrats estimated how often the average Republican chose to work for urban areas in our previous study), using a 5-point scale: *Almost never (0–19%)*, *Rarely (20–39%)*, *Sometimes (40–59%)*, *Often (60–79%)*, or *Almost always (80–100%)*. In the correction condition, participants received feedback indicating that out-party supporters had in fact chosen to exert such effort often (approximately 60% of the time). This figure was based on the model-based out-group choice rate to work of the average Democrat or Republican participant in Study 1. In the control condition, participants received no feedback. All participants then completed the Pro-Democratic Effort Task. As in Study 1, participants were not informed of the credit-to-money conversion rate, and the task involved

the same deceptive element, with no actual non-partisan organization supporting voter engagement in only rural or only urban areas. For Study 2, we thus donated the maximum per-participant task earning (\$0.20) to Vote.org, and participants were debriefed about this deceptive element and informed that their earnings had been donated in full.

Ethical approval for both studies was obtained from the Research Ethics Committee at Royal Holloway, University of London (Reference ID: 768), including the use of a minor deceptive element (the framing of region-specific voter engagement programs). All participants provided informed consent and were fully debriefed about this element at the end of each study.

Pro-Democratic Effort Task

On each trial of the Pro-Democratic Effort Task (see Figure 1), participants chose between a low-reward rest option (3 credits for resting 10 seconds) and a higher-reward work option that required a different level of effort (65%, 80%, or 95% of an individually calibrated maximum) and provided larger rewards (4, 12, or 20 credits). Effort was operationalized as clicking different numbers of boxes on the screen within 10 seconds, an established paradigm that reliably captures individual differences in effort-based decision-making (Contreras-Huerta et al., 2020; Cutler et al., 2025; Lockwood et al., 2017).

Before receiving task instructions, participants completed two calibration rounds in which they clicked as many boxes as possible within the time limit, with the higher count serving as their individual maximum effort threshold (minimum: 13 boxes). Participants then read the task instructions and completed five practice trials and two comprehension questions, which had to be answered correctly before proceeding to the main task. During the task, participants had 4 seconds to decide whether to work or rest, with no credits awarded for missed trials. If participants chose to work, they had to meet the box-clicking target within 10 seconds to receive the reward, and failure resulted in no credits.

Participants were instructed that the generated credits would benefit one of several programs supporting democratic engagement, each with a distinct regional focus. The instructions explicitly described all programs as non-partisan while noting that, owing to their regional focus, the programs might reach different communities. The instructions then stated the historical partisan lean of each area directly: the Vote Urban Program targeted urban areas, such as major cities and metropolitan regions, described as having historically shown higher support for Democratic candidates; the Vote Rural Program targeted rural areas, described as having historically shown higher support for Republican candidates; and, in Study 1 only, the

Vote All Program targeted voter turnout across all US communities, described as historically balanced in partisan support.

In Study 1, participants completed 27 trials in random order (9 per beneficiary: 3 effort levels $\times$ 3 reward levels). In Study 2, the Vote All condition was omitted and participants completed 18 trials (9 per beneficiary).

Survey Measures

Both studies included a battery of survey measures assessing partisanship, affective polarization, and additional political and demographic variables. Study 1 additionally included measures of mental and physical health. All measures are described in the Supplementary Materials and provided in their full form on the OSF project page.

Most central to the present analyses, partisanship was assessed with a single item asking participants which political position best described them (Democrat, Republican, Independent, or Other). Affective polarization was measured with four items assessing how favorably (valence) and how intensely (arousal) participants felt toward Democrats and Republicans, each rated on a 101-point scale. Following recent methodological developments (Kasper et al., 2025), we computed each participant's arousal-weighted mean valence and then the arousal-weighted spread of valence around this mean, with higher values indicating greater affective polarization.

Analytic Strategy

Computational models of effort discounting were fitted in MATLAB R2025b (The MathWorks Inc., 2025). Statistical models of choice data and model parameters were conducted in R Version 4.5.2 (R Core Team, 2025).

Computational Modeling

We fitted hierarchical computational models estimating beneficiary-specific effort-discounting parameters (K) and a decision consistency parameter (β). On each trial, the subjective value (SV) of the work option was computed according to a candidate discount function, and choice probabilities were modeled using a softmax function comparing the SV of work against the fixed rest alternative (3 credits). This modeling framework was adapted from the Pro-Environmental Effort Task and followed the same hierarchical maximum a posteriori (MAP) fitting procedure and model comparison strategy (Cutler et al., 2025; Lockwood et al., 2017). Full details of the fitting procedure are provided in the Supplementary Materials.

In Study 1, we compared six candidate models with three beneficiary-specific K parameters ($3K$), varying in discount function (linear: $SV = R - K * E$; hyperbolic: $SV = R/(1$

+ $K * E$); parabolic: $SV = R - K * E^2$) and in whether β was shared across beneficiaries (1β) or beneficiary-specific (3β). In Study 2, we fitted analogous models with two beneficiary-specific K parameters ($2K$). Model comparison relied on integrated BIC and random-effects Bayesian model selection, yielding exceedance probabilities for each candidate model; results are shown in Figure 1. Model-identifiability and parameter-recovery simulations supported the recoverability of the winning models and their parameters from the task structure (Figures S2 and S5).

In both studies, the winning model was the hyperbolic model with beneficiary-specific K parameters and a single shared β ($3K1\beta$ in Study 1; $2K1\beta$ in Study 2), indicating that effort discounting differed across beneficiaries whereas decision consistency was best captured by a single shared parameter. From the winning models, participant-level K parameters were extracted and recoded relative to partisanship: for Democrats, $K_{\text{ingroup}} = K_{\text{urban}}$ and $K_{\text{outgroup}} = K_{\text{rural}}$; for Republicans, the reverse. Higher K values indicate steeper effort discounting and lower motivation; lower K values indicate shallower discounting and greater motivation. In-group favoritism was operationalized as $K_{\text{outgroup}} - K_{\text{ingroup}}$, with higher values reflecting greater relative motivation to exert effort for in-group over out-group areas.

Statistical Models

In Study 1, trial-level choices to work or rest were modeled using a generalized linear mixed model (GLMM) with a logit link function, fitted with the afex package (Singmann et al., 2024). The model included fixed effects for beneficiary, effort, reward, and all their interactions. Following our preregistered protocol, we began with the maximal random effects structure (Barr et al., 2013) and simplified in cases of non-convergence by sequentially trying alternative optimizers, removing correlations between random effects, and removing random slopes with variance near zero. The final converging model included random intercepts and random slopes for beneficiary: $\text{decision} \sim \text{beneficiary} \times \text{effort} \times \text{reward} + (1 + \text{beneficiary} | \text{participant})$. Fixed effects were assessed using likelihood ratio tests and interpreted via model-based predictions and average comparisons computed with the marginalesffects package (Arel-Bundock et al., 2024).

All regression models predicting K parameters and survey outcomes were fitted using the estimatr package (Blair et al., 2025) with heteroscedasticity-consistent (HC2) standard errors. In addition, to test whether the misperception correction in Study 2 selectively affected out-group versus in-group effort discounting, we fitted a linear mixed model using afex, with K as the outcome in long format, and beneficiary (in-group vs. out-group), condition (correction vs. control), and their interaction as predictors. Following our

preregistered protocol, we again started with a maximal random effects structure including a by-participant random slope for beneficiary. Because each participant contributed only two K values (one per beneficiary), the by-participant random slope for beneficiary was not identifiable, and we simplified to a random intercept-only model: $K \sim \text{beneficiary} \times \text{condition} + (1 \mid \text{participant})$. Continuous focal predictors and K -based outcomes were z-standardized for regression analyses, so that reported coefficients can be interpreted as standardized effect sizes. Trial-level work decisions were analyzed as binary outcomes.

In follow-up analyses in Study 2, we additionally tested for non-linear moderation of the treatment effect by affective polarization using a generalized additive model (Wood, 2017), and quantified evidence for the null hypothesis of no moderation via Bayesian model comparison with Jeffreys–Zellner–Siow (JZS) priors (Morey & Rouder, 2024).

Statistical Inference

For all preregistered hypotheses, we used $\alpha = .05$ (two-sided) and report 95% confidence intervals alongside p values. For directional hypotheses, a significant p value was a necessary but not sufficient condition for support, as the effect also had to be in the predicted direction. All analyses not specified in the preregistration are described as exploratory.

Transparency and Openness

Both studies were preregistered on the Open Science Framework before data collection, including hypotheses, sampling and exclusion rules, and analysis plans (Study 1: <https://doi.org/10.17605/OSF.IO/ZYEVVR>; Study 2: <https://doi.org/10.17605/OSF.IO/KXP2R>). Data, analysis code, and all study materials for both studies are openly available at <https://doi.org/10.17605/OSF.IO/F6CGS>. The repository contains the raw and processed data, the MATLAB code used for computational model fitting, and the R scripts used for all statistical analyses, as well as survey and experimental materials (also see Supplementary Materials).

Generative AI

We used generative AI (Claude Opus 4.8, Anthropic, 2026; ChatGPT-5.5, OpenAI, 2026) under full human supervision for code optimization and documentation, and for proofreading and editing of the manuscript. We have carefully reviewed and edited all code and text and take full responsibility for the final manuscript and materials. No other AI-assisted technologies were used in this research.

Deviations from Preregistration

We report one major and three minor deviations from our Study 1 preregistration (see Table S1 for details). The major deviation concerns the preregistered model space. Due to an oversight, our preregistration included models with a single shared K parameter ($1K$) and partially shared parameters ($2K$) alongside fully beneficiary-specific models ($3K$), and did not include parabolic discounting functions. Because only models with separate K parameters for each beneficiary permit testing our hypotheses, we restricted the comparison to $3K$ models and added parabolic discounting functions, in line with prior modeling approaches in similar designs (Cutler et al., 2025; Lockwood et al., 2017). As minor deviations, we applied heteroscedasticity-consistent (HC2) standard errors in Study 1, although these were preregistered only for Study 2, to ensure consistent and appropriately conservative inference across both studies (Hayes & Cai, 2007). We also z-scored all continuous variables in Study 1 rather than the preregistered percentage scoring for questionnaire measures, matching Study 2 and ensuring consistent interpretation of regression coefficients across analyses. Finally, we removed two duplicate submissions from Study 1, retaining only the first submission per participant, which the preregistration had not anticipated. There were no deviations from the Study 2 preregistration.

Results

Detailed sample and variable descriptives for both studies, including pairwise correlations among all measures, are provided in Tables S2–S5 and S10–S11 and Figures S1 and S4.

Study 1

In Study 1, we validated the task as a measure of pro-democratic effort across partisan lines and investigated the correlates of pro-democratic efforts for different beneficiaries.

Behavioral Analysis

Following our preregistered protocol, we modeled trial-level choices to work or rest using a generalized linear mixed model (GLMM). All three within-subject factors significantly predicted the decision to work: beneficiary ($\chi^2(2) = 65.67, p < .001$); effort ($\chi^2(2) = 116.39, p < .001$); and reward ($\chi^2(2) = 252.64, p < .001$), alongside a significant beneficiary $\times$ reward interaction ($\chi^2(4) = 12.18, p = .016$; see Figure 2 and Tables S6–S7).

Model-based contrasts confirmed our preregistered hypotheses: participants were significantly more likely to work for in-group areas (82.5%, 95% CI [78.5, 86.5]) than for out-group areas (61.6%, 95% CI [55.1, 68.1]; $\Delta = 20.9$ percentage points, 95% CI [15.4, 26.3], $p < .001$), and more likely to work for all areas across the US (81.8%, 95% CI [77.5, 86.1]) than for out-group areas ($\Delta = 20.2$ percentage points, 95% CI [15.1, 25.2], $p < .001$). However, willingness to work for in-group areas did not significantly differ from willingness to work for all areas across the US ($\Delta = 0.7$ percentage points, 95% CI [−2.6, 4.1], $p = .678$).

These results reveal a partisan gap in decisions to make pro-democratic efforts. Participants were most willing to work for their own side or both sides equally, and least willing to work for out-group areas. At the same time, willingness to work for pro-democratic efforts was high (more than half of trials) even for out-group areas.

Computational Modeling


As participants were sensitive to both effort and reward (Figure 2), we used hierarchical computational models to estimate separate effort-discounting parameters (K) for each beneficiary, where higher K parameters index steeper discounting or lower motivation, and lower K parameters index shallower discounting or higher motivation. Among six candidate models varying in discount function and parameterization, a hyperbolic model with three area-specific K parameters and a single shared decision-consistency parameter β (3K1 β) provided the best fit (Figure 1b).

Consistent with the behavioral findings, there was steeper effort discounting for out-group than in-group areas (Figure 2) suggesting that participants were less motivated to exert

effort for the out-group. Robust linear regressions controlling for locality of residence, age, gender, education, and general effort capacity (individually calibrated maximum box-clicking rate) confirmed that Republicans discounted efforts for rural areas (i.e., their in-group area) less than Democrats ($b = -0.52$, 95% CI $[-0.77, -0.28]$, $p < .001$), whereas Democrats discounted efforts for urban areas less ($b = 0.41$, 95% CI $[0.17, 0.64]$, $p < .001$).

Figure 2

Behavioral Choices and Effort Discounting Across Conditions



Note. Choice behavior in Study 1 (top) and Study 2 (bottom) is shown as a function of (a) effort level (1 = 65%, 2 = 80%, 3 = 95%) and (b) available reward (4, 12, or 20 credits), alongside (c) effort-discounting (K) parameters by beneficiary condition. In (a) and (b), lines represent mean proportions of choices to work, and shaded bands represent 95% confidence intervals. In (c), bars represent group means, error bars indicate 95% confidence intervals, and points represent individual participants. Significance brackets reflect paired-samples t tests, with Cohen's d values from the same tests. *ns* = not significant; *** $p < .001$.

Validation of Modeling Parameters

To validate that the modeling parameters captured motivation relevant to real-world pro-democratic decisions, we examined participants' choice to donate or keep the money at

the end of the study. Because all participants received the same maximum bonus (\$0.20), and the option to donate or keep it was not disclosed until after the task (see Procedure), this choice was independent of task earnings and could not have influenced task behavior. Participants could either donate the bonus to a non-partisan organization operating across the US or keep it in addition to their baseline compensation. Of the 284 participants, 126 (44.4%) chose to donate and 158 (55.6%) chose to keep the bonus. Supporting our preregistered validation, participants with higher motivation for the all-areas condition (lower K_{all}) were significantly more likely to donate ($b = -0.42$, 95% CI $[-0.64, -0.19]$, $p < .001$), indicating that the task captured motivation that translated into a real-world pro-democratic choice.



Partisan Differences and Political and Psychological Correlates

Using the same robust linear regression specification with controls as before, we then tested whether Democrats and Republicans differed in effort discounting for the all-areas condition (K_{all}) or their in-group favoritism ($K_{outgroup} - K_{ingroup}$, where higher values indicate steeper discounting for out-group relative to in-group areas).

While Democrats and Republicans did not significantly differ in effort discounting for all areas ($b = 0.17$, 95% CI $[-0.07, 0.42]$, $p = .164$), Democrats showed significantly stronger in-group favoritism than Republicans ($b = -0.34$, 95% CI $[-0.58, -0.10]$, $p = .006$; Figure 3).

In addition, affective polarization was positively associated with greater in-group favoritism ($b = 0.25$, 95% CI $[0.13, 0.36]$, $p < .001$; Figure 3). Other preregistered correlations between effort-discounting parameters and survey measures (democratic support, voting importance, and indicators of physical and mental health) were not statistically significant or small in magnitude, with exploratory Bayesian correlations indicating moderate evidence for the null for most associations (Tables S8–S9 and Figure S3).

In summary, we validated the Pro-Democratic Effort Task as a measure of pro-democratic efforts. US partisans showed clear in-group favoritism in effort discounting, an effect that was more pronounced among more affectively polarized participants. Yet the average participant showed substantial willingness to work, even for areas associated with the political out-group, choosing to work on 61.6% of trials benefiting it. Study 2 tested whether partisans expect this willingness from their political opponents, and whether correcting misperceptions about out-party effort willingness could reduce in-group favoritism.

Figure 3*In-Group Favoritism by Partisanship and Affective Polarization***Study 1****Study 2**

Note. In-group favoritism ($K_{\text{outgroup}} - K_{\text{ingroup}}$) in Study 1 (top) and Study 2 (bottom) is shown in (a) as distributions, box plots, and individual data points for Democrat and Republican participants, and in (b) as its association with affective polarization, with the shaded band representing the 95% confidence interval around the regression line. Bracketed annotations report coefficients from robust linear regressions controlling for age, gender, education, locality, and effort capacity. Boxes represent interquartile ranges with median lines, and whiskers extend to 1.5 times the interquartile range. Asterisks denote significance level: $** p < .01$; $*** p < .001$.

Study 2

In Study 2, participants first estimated how willing out-party supporters were to exert effort for areas associated with their own party (e.g., Democrats estimated how often the average Republican had chosen to work for urban areas, and vice versa). Participants were


then randomly assigned to either a misperception correction condition ($n = 317$), in which they received feedback that out-party supporters had in fact chosen to exert effort often (approximately 60% of the time, based on the findings from Study 1), or to a control condition ($n = 315$), where they received no feedback. Participants then completed the same Pro-Democratic Effort Task as in Study 1, but with only two beneficiary conditions (in-group and out-group areas; Figure 1).

Baseline Misperceptions

Consistent with our preregistration, most participants (~87%) systematically underestimated the willingness of the political out-group to make pro-democratic efforts for areas associated with the participant's own party (Figure 4). On average, participants' estimates fell 1.57 scale points below the correct answer of "Often, 60–79%", placing the average estimate closest to "Rarely, 20–39%" ($b = -1.57$, 95% CI $[-1.64, -1.50]$, $p < .001$), with no significant difference in pre-treatment underestimation between the correction and control conditions ($b = 0.08$, 95% CI $[-0.07, 0.23]$, $p = .294$). As predicted, underestimation was also stronger among more affectively polarized participants ($b = -0.17$, 95% CI $[-0.25, -0.10]$, $p < .001$).

Figure 4

Perceptions of Out-Party Willingness and Effects of Misperception Correction



Note. The distribution of Study 2 participants' estimates of how often the average out-party supporter had chosen to exert effort for areas associated with the participant's own party is shown in (a); the dashed vertical line indicates the sample mean. Effort-discounting (K) parameters for in-group and out-group beneficiaries by condition, restricted to participants who underestimated out-party effort willingness at baseline, are shown in (b). Bars represent

condition means, error bars represent 95% confidence intervals, and points represent individual participants; bracketed annotations report between-condition differences from preregistered linear mixed models. Higher K values indicate steeper effort discounting (lower motivation to exert effort for a given beneficiary). *ns* = not significant; *** $p < .001$.

Correcting Misperceptions

As in Study 1, we then estimated effort-discounting parameters (K) for in-group and out-group beneficiaries using hierarchical computational modeling (winning model: $2K1\beta$ -hyperbolic; see Figure 1). These K parameters served as our primary outcome variables. Following our preregistration, we restricted our analyses to participants who underestimated out-party effort willingness ($N = 550$, 87% of total sample; of which $n = 273$ were in the treatment and $n = 277$ in the control condition), because the treatment was expected to be relevant primarily for participants who held the targeted misperception (see Table S12 for comparisons between participants who over-, under-, or accurately estimated). We expected that learning that out-party supporters were more willing than expected to exert effort for areas associated with participants' own party would reduce in-group favoritism by decreasing participants' tendency to discount out-group efforts more steeply than in-group efforts.

Supporting our preregistered hypotheses, participants in the correction condition showed significantly weaker in-group favoritism in effort discounting ($K_{\text{outgroup}} - K_{\text{ingroup}}$) than those in the control condition ($b = -0.21$, 95% CI $[-0.37, -0.04]$, $p = .016$; Figure 4). To determine whether this reflected changes in out-group effort discounting, in-group effort discounting, or both, we fitted a preregistered linear mixed model with K as the outcome, including beneficiary (in-group vs. out-group), condition (correction vs. control), and their interaction as predictors, with a random intercept for participant (see Method). The beneficiary $\times$ condition interaction was significant ($b = -0.20$, 95% CI $[-0.38, -0.01]$, $p = .042$), indicating that the correction selectively reduced effort discounting for out-group areas ($b = -0.27$, 95% CI $[-0.44, -0.10]$, $p = .001$), with no significant effect on in-group effort discounting ($b = -0.08$, 95% CI $[-0.24, 0.09]$, $p = .377$). Results were comparable when modeling participants' choices (see Tables S13–S14). In other words, correcting misperceptions about out-party effort willingness selectively increased participants' motivation to exert effort for areas associated with the opposing party, with no significant reduction in their motivation to support their own party's areas.

Heterogeneity in Treatment Effects

As preregistered, we tested whether the treatment effect on in-group favoritism was moderated by affective polarization, either because more polarized participants might be more resistant to updating their views of the out-party (predicting weaker effects at higher polarization) or because they hold stronger misperceptions that leave more room for correction (predicting stronger effects). We found no significant interaction between the correction and affective polarization on in-group favoritism ($b = 0.06$, 95% CI $[-0.11, 0.22]$, $p = .508$). Corroborating this null finding, an exploratory generalized additive model suggested no evidence for a non-linear interaction (Figure S6), and a Bayesian model comparison using Jeffreys–Zellner–Siow (JZS) priors (Morey & Rouder, 2024) yielded moderate-to-strong evidence for the null (no interaction) across different prior specifications (medium: $BF_{01} = 6.88$; wide: $BF_{01} = 8.31$; ultrawide: $BF_{01} = 11.64$; Table S15). Together, these analyses indicate that the correction worked similarly across levels of affective polarization.

Affective Polarization and Partisan Differences

Beyond the experimental effects, we examined whether key correlational patterns from Study 1 were replicated in Study 2. Consistent with our preregistration and replicating Study 1, higher affective polarization was associated with greater in-group favoritism across both treatment conditions ($b = 0.27$, 95% CI $[0.20, 0.35]$, $p < .001$; Figure 3). In an exploratory analysis, we also replicated the finding that Democrats showed significantly stronger in-group favoritism than Republicans ($b = -0.30$, 95% CI $[-0.45, -0.15]$, $p < .001$; Figure 3).

Finally, an exploratory analysis indicated that participants' underestimation of out-party effort willingness was associated with their own in-group favoritism. The more participants underestimated (more negative discrepancy from the correct answer), the greater was their in-group favoritism in effort discounting ($b = -0.18$, 95% CI $[-0.30, -0.06]$, $p = .003$).

To summarize, US partisans substantially underestimated how willing their political opponents were to exert pro-democratic effort for areas associated with their own party, and those with stronger underestimation showed greater in-group favoritism in their own effort discounting. Providing accurate information about out-party effort willingness significantly reduced this gap by selectively increasing motivation to make efforts for out-group areas, and the effect was consistent across levels of affective polarization.

Discussion

Who exerts effort for democracy, for whom, and under what conditions? We found that US partisans showed clear in-group favoritism, exerting less effort for democratic engagement in areas traditionally supporting the opposing party, with this gap widening among more affectively polarized individuals. However, participants were also highly willing to make efforts benefiting all areas and commonly made pro-democratic efforts even when these benefited out-group areas. Importantly, while participants systematically underestimated the willingness of out-partisans to make such efforts, the correction of such misperceptions increased motivation to engage out-group areas without significantly undermining in-group motivation. Motivation to support democratic engagement thus appeared to be present but conditional, partly restrained by misperceptions about out-partisans rather than absent altogether. But as the findings of Study 2 show, this pattern can be successfully altered in a positive way.

Our approach extended neuroeconomic models of effort-based decision-making (Contreras-Huerta et al., 2020; Lockwood et al., 2017; Westbrook & Braver, 2015; Zipf, 1949) to the political domain, modeling the trade-off between effort and reward in pro-democratic action. This allowed us to isolate motivation to support democratic causes and compare it across beneficiaries within the same individual. Across both studies, motivation was lower for areas leaning toward the opposing party than for areas leaning toward one's own. Importantly, this effect was present even as we controlled for participants' own locality of residence, indicating that it reflected partisan rather than geographic divisions, although the two are increasingly intertwined (Mettler & Brown, 2022). Complementing prior work on prosocial apathy, whereby people are less willing to exert effort for others than for themselves (Lockwood et al., 2017), we thus documented a form of democratic apathy directed at efforts that strengthen democracy but politically benefit the out-group more than one's own side—an asymmetry that scaled with participants' level of affective polarization.

Notably, in-group favoritism was also stronger among Democrats than Republicans in both studies. We did not predict this partisan difference, and it should be interpreted cautiously. One speculative explanation is that both studies were conducted at a time when Republicans held the US presidency and majorities in both chambers of Congress, so that Democrats may have perceived pro-democratic effort as more urgently needed for their own side. If so, this finding might weaken or reverse when Democrats hold power. Future work could test this directly by administering the task on a more localized level in states governed by different parties, or at the national level should the balance of power shift.

In addition, even with in-group favoritism clearly present, pro-democratic motivation was not confined to partisan interests. In Study 1, willingness to exert effort for all-areas initiatives did not significantly differ from willingness to support in-group areas, suggesting that pro-democratic motivation can extend beyond narrow partisan interests when benefits are framed as collective rather than partisan. Yet, motivation for out-group areas, while lower, was not absent, suggesting that participants recognized the shared democratic benefit it could carry.

Why, then, did participants make less effort for areas associated with the out-party? One possibility is that they underestimate how willing out-partisans are to make such efforts, which in turn might reduce their own willingness, resulting in a self-reinforcing antidemocratic spiral (Braley et al., 2023). Previous work has shown that partisans overestimate out-partisans' hostility and support for antidemocratic practices, and that correcting such misperceptions can improve democratic attitudes (Graham & Svolik, 2020; Voelkel et al., 2024). We found that indeed most participants underestimated how willing out-party supporters were to make pro-democratic efforts for areas associated with the participant's own party. Correcting this misperception selectively raised participants' own motivation for out-group areas, without significantly undermining motivation for in-group areas, and this effect was evident across levels of affective polarization. Extending prior work, these results highlight that misperceptions are not only attitudinally consequential but also shape the motivation to incur effort costs for democratic goods, at both a behavioral and a computational level. Democratic motivation therefore tracks not only partisan commitments but also meta-perceptions about how others might act.

At the same time, misperception corrections do not always produce robust effects in more contested information environments (Druckman, 2023). Our findings do not imply that simple informational interventions will uniformly improve democratic functioning, but rather that corrections can be effective when they target concrete, behaviorally relevant expectations. In addition, while the Pro-Democratic Effort Task offers a controlled and computationally tractable measure of motivation, it remains a proxy for real-world democratic behavior. However, providing convergent evidence that the task captured consequential motivation, participants who exerted more effort for all-areas were more likely to donate their task earnings bonus to a non-partisan voter-engagement organization rather than keep them for themselves, with this option revealed only after the task had been completed (Study 1). Nonetheless, future work should test whether these patterns extend to naturally occurring civic engagement, such as volunteering or cross-partisan collaboration.

Constraints on Generality

Our inferences are intended to generalize to US adults who identify as Democrats or Republicans, with Independents and supporters of other parties excluded by design. Samples in both studies were recruited through Prolific and stratified to approximate benchmarks from an ANES probability sample on gender, ethnicity, and age within each party. While Prolific provides comparatively high data quality among opt-in online panels (Stagnaro et al., 2026), it is not a probability sample of US partisans, and absolute levels of effort and misperception should not be read as population estimates. We further operationalized effort as physically clicking boxes within a time limit, and partisan benefit through the political rural–urban divide. While we expect that core findings hold for other forms of effort (e.g., cognitive) and other ways to manipulate which party efforts benefit, we did not test this directly, and future work is also required to examine whether findings hold in multiparty systems, for other group divides, or for real-world effortful civic engagement such as volunteering. Finally, Study 1 was conducted in October and November 2025 and Study 2 in March 2026, when Republicans held the presidency and both chambers of Congress. Some findings may depend on this political context and should be tested under a different balance of power.

Conclusion

Together, our findings point to a coordination problem in democratic engagement. Many forms of pro-democratic effort generate collective benefits but require individual costs, and willingness to act depends not only on one's own preferences but also on expectations about others' contributions (Fehr & Gächter, 2000). When citizens believe out-partisans are unwilling to support shared democratic goods, they may withhold effort even when they would otherwise act. This constraint is real but not fixed: willingness to act across party lines was greater than participants assumed and correcting that misperception reduced in-group favoritism without significantly undermining support for one's own side. On this account, the erosion of cross-partisan cooperation may be driven partly by inaccurate beliefs about others' willingness to contribute. Making this willingness visible may be one pathway to strengthening the collective effort democracy depends on.

References

- American National Election Studies. (2022). *ANES 2022 pilot study* (Version December 14, 2022) [Data set]. <https://electionstudies.org/data-center/2022-pilot-study>
- Anthropic. (2026). *Claude* (Opus 4.8 version) [Large language model]. <https://claude.ai>
- Arel-Bundock, V., Greifer, N., & Heiss, A. (2024). How to interpret statistical models using marginaeffects for R and Python. *Journal of Statistical Software*, 111(9), 1–32. <https://doi.org/10.18637/jss.v111.i09>
- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68(3), 255–278. <https://doi.org/10.1016/j.jml.2012.11.001>
- Bishop, B., & Cushing, R. G. (2008). *The big sort: Why the clustering of like-minded America is tearing us apart*. Houghton Mifflin Harcourt.
- Blair, G., Cooper, J., Coppock, A., & Humphreys, M. (2019). Declaring and diagnosing research designs. *American Political Science Review*, 113(3), 838–859. <https://doi.org/10.1017/S0003055419000194>
- Blair, G., Cooper, J., Coppock, A., Humphreys, M., & Sonnet, L. (2025). *estimatr: Fast estimators for design-based inference* [Manual]. <https://doi.org/10.32614/CRAN.package.estimatr>
- Braley, A., Lenz, G. S., Adjodah, D., Rahnama, H., & Pentland, A. (2023). Why voters who value democracy participate in democratic backsliding. *Nature Human Behaviour*, 7(8), 1282–1293. <https://doi.org/10.1038/s41562-023-01594-w>
- Contreras-Huerta, L. S., Pisauro, M. A., & Apps, M. A. J. (2020). Effort shapes social cognition and behaviour: A neuro-cognitive framework. *Neuroscience & Biobehavioral Reviews*, 118, 426–439. <https://doi.org/10.1016/j.neubiorev.2020.08.003>
- Cutler, J., Contreras-Huerta, L. S., Todorova, B., Nitschke, J., Michalaki, K., Koppel, L., Gkinopoulos, T., Vogel, T. A., Lamm, C., Västfjäll, D., Tsakiris, M., Apps, M. A. J., & Lockwood, P. L. (2025). Psychological interventions that decrease psychological distance or challenge system justification increase motivation to exert effort to mitigate climate change. *Communications Psychology*, 3(1), 148. <https://doi.org/10.1038/s44271-025-00332-4>
- Depow, G. J., Lin, H., & Inzlicht, M. (2022). Cognitive effort for self, strangers, and charities. *Scientific Reports*, 12(1), 15009. <https://doi.org/10.1038/s41598-022-19163-y>

- Druckman, J. N. (2023). Correcting misperceptions of the other political party does not robustly reduce support for undemocratic practices or partisan violence. *Proceedings of the National Academy of Sciences*, 120(37), e2308938120.
<https://doi.org/10.1073/pnas.2308938120>
- Druckman, J. N., Klar, S., Krupnikov, Y., Levendusky, M., & Ryan, J. B. (2022). (Mis)estimating affective polarization. *The Journal of Politics*, 84(2), 1106–1117.
<https://doi.org/10.1086/715603>
- Fehr, E., & Gächter, S. (2000). Cooperation and punishment in public goods experiments. *American Economic Review*, 90(4), 980–994. <https://doi.org/10.1257/aer.90.4.980>
- Graham, M. H., & Svolik, M. W. (2020). Democracy in America? Partisanship, polarization, and the robustness of support for democracy in the United States. *American Political Science Review*, 114(2), 392–409. <https://doi.org/10.1017/S0003055420000052>
- Hayes, A. F., & Cai, L. (2007). Using heteroskedasticity-consistent standard error estimators in OLS regression: An introduction and software implementation. *Behavior Research Methods*, 39(4), 709–722. <https://doi.org/10.3758/BF03192961>
- Kasper, J., Bakker, B. N., Lelkes, Y., & Schumacher, G. (2025). *Beyond valence: Arousal as a core dimension of affective polarization*. PsyArXiv.
https://doi.org/10.31234/osf.io/5dfbj_v1
- Kool, W., & Botvinick, M. (2018). Mental labour. *Nature Human Behaviour*, 2(12), 899–908. <https://doi.org/10.1038/s41562-018-0401-9>
- Kurzban, R. (2016). The sense of effort. *Current Opinion in Psychology*, 7, 67–70.
<https://doi.org/10.1016/j.copsyc.2015.08.003>
- Lees, J., & Cikara, M. (2020). Inaccurate group meta-perceptions drive negative out-group attributions in competitive contexts. *Nature Human Behaviour*, 4(3), 279–286.
<https://doi.org/10.1038/s41562-019-0766-4>
- Levendusky, M. S., & Malhotra, N. (2016). (Mis)perceptions of partisan polarization in the American public. *Public Opinion Quarterly*, 80(S1), 378–391.
<https://doi.org/10.1093/poq/nfv045>
- Levitsky, S., & Ziblatt, D. (2018). *How democracies die*. Crown.
- Lockwood, P. L., Hamonet, M., Zhang, S. H., Ratnavel, A., Salmony, F. U., Husain, M., & Apps, M. A. J. (2017). Prosocial apathy for helping others when effort is required. *Nature Human Behaviour*, 1(7), 0131. <https://doi.org/10.1038/s41562-017-0131>
- Lockwood, P. L., Wittmann, M. K., Nili, H., Matsumoto-Ryan, M., Abdurahman, A., Cutler, J., Husain, M., & Apps, M. A. J. (2022). Distinct neural representations for prosocial

- and self-benefiting effort. *Current Biology*, 32(19), 4172–4185.e7.
<https://doi.org/10.1016/j.cub.2022.08.010>
- Mernyk, J. S., Pink, S. L., Druckman, J. N., & Willer, R. (2022). Correcting inaccurate metaperceptions reduces Americans' support for partisan violence. *Proceedings of the National Academy of Sciences*, 119(16), e2116851119.
<https://doi.org/10.1073/pnas.2116851119>
- Mettler, S., & Brown, T. (2022). The growing rural-urban political divide and democratic vulnerability. *The Annals of the American Academy of Political and Social Science*, 699(1), 130–142. <https://doi.org/10.1177/00027162211070061>
- Moore-Berg, S. L., Ankori-Karlinsky, L.-O., Hameiri, B., & Bruneau, E. (2020). Exaggerated meta-perceptions predict intergroup hostility between American political partisans. *Proceedings of the National Academy of Sciences*, 117(26), 14864–14872.
<https://doi.org/10.1073/pnas.2001263117>
- Morey, R. D., & Rouder, J. N. (2024). *BayesFactor: Computation of Bayes factors for common designs* [Manual]. <https://doi.org/10.32614/CRAN.package.BayesFactor>
- Nadeau, R., Daoust, J.-F., & Dassonneville, R. (2023). Winning, losing, and the quality of democracy. *Political Studies*, 71(2), 483–500.
<https://doi.org/10.1177/00323217211026189>
- Nadeem, R. (2024, April 9). Party affiliation of US voters in urban, rural and suburban communities. *Pew Research Center*.
<https://www.pewresearch.org/politics/2024/04/09/partisanship-in-rural-suburban-and-urban-communities/>
- OpenAI. (2026). *ChatGPT* (5.5 version) [Large language model]. <https://chatgpt.com>
- Pasek, M. H., Ankori-Karlinsky, L.-O., Levy-Vene, A., & Moore-Berg, S. L. (2022). Misperceptions about out-partisans' democratic values may erode democracy. *Scientific Reports*, 12(1), 16284. <https://doi.org/10.1038/s41598-022-19616-4>
- R Core Team. (2025). *R: A language and environment for statistical computing* (Version 4.5.2) [Computer software]. R Foundation for Statistical Computing. <https://www.R-project.org/>
- Singmann, H., Bolker, B., Westfall, J., Aust, F., & Ben-Shachar, M. S. (2024). *afex: Analysis of factorial experiments* [Manual]. <https://CRAN.R-project.org/package=afex>
- Stagnaro, M. N., Druckman, J. N., Berinsky, A. J., Arechar, A. A., Willer, R., & Rand, D. G. (2026). Representativeness and response validity across nine opt-in online samples.

Nature Human Behaviour. Advance online publication.

<https://doi.org/10.1038/s41562-026-02438-z>

The MathWorks Inc. (2025). *MATLAB version: 25.2.0 (R2025b)* [Computer software].

<https://www.mathworks.com>

Voelkel, J. G., Stagnaro, M. N., Chu, J. Y., Pink, S. L., Mernyk, J. S., Redekopp, C., Ghezze, I., Cashman, M., Adjodah, D., Allen, L. G., Allis, L. V., Baleria, G., Ballantyne, N., Van Bavel, J. J., Blunden, H., Braley, A., Bryan, C. J., Celniker, J. B., Cikara, M., ... Willer, R. (2024). Megastudy testing 25 treatments to reduce antidemocratic attitudes and partisan animosity. *Science*, 386(6719), eadh4764.

<https://doi.org/10.1126/science.adh4764>

Westbrook, A., & Braver, T. S. (2015). Cognitive effort: A neuroeconomic approach.

Cognitive, Affective, & Behavioral Neuroscience, 15(2), 395–415.

<https://doi.org/10.3758/s13415-015-0334-y>

Wood, S. N. (2017). *Generalized additive models: An introduction with R*. Chapman and Hall/CRC.

Zipf, G. K. (1949). *Human behavior and the principle of least effort*. Addison-Wesley Press.