

What Do We Learn From a Decade of Candidate Choice Conjoint Experiments?*

Alexander Coppock, Matthew Blyth, Marco M. Aviña, and Semra Sevi[†]

October 19, 2026

Abstract

Candidate choice conjoint experiments are now among the most widely deployed designs in political science, but their lessons are scattered across hundreds of separate investigations. We present a meta-reanalysis of 243 such studies (reported in 144 papers) conducted in 52 countries among 434,079 respondents who chose between hypothetical candidates in 2,176,525 pairs. Political attributes like policy positions and partisanship exert substantially larger effects than demographic attributes such as race, gender, age, and religion, a result that informs enduring theoretical debates over

*We thank Danielle Jing, Natália Pimenta, and Sonia Meroian for detailed and helpful feedback, and participants at the 2025 APSA Annual Meeting and the 2026 NYU CESS Experimental Political Science Conference for their comments.

[†]Alexander Coppock is Associate Professor of Political Science at Northwestern University. Matthew Blyth is a Tobin Center pre-doctoral fellow at Yale University. Marco M. Aviña is a doctoral candidate in Government and Social Policy at Harvard University. Semra Sevi is Assistant Professor of Political Science at the University of Toronto. Corresponding author: alex.coppock@northwestern.edu. As a companion to this paper, an interactive web application that allows users to explore the generalizability of the many possible comparisons available in the corpus of candidate choice experiments is available at <https://anonymouspoliticalscientist.shinyapps.io/afcp-explorer/>. Only after an initial round of data collection and full conference draft did we integrate AI tools into our workflow, the main bulk of which was to clean, reproduce, reanalyze, and cross-check the results of hundreds of conjoint experiments. Appendix R describes our uses of AI and the checks we applied to its output. We take full responsibility for every estimate, citation, and claim in the paper.

substantive and descriptive representation. The effects of most attributes shrink when policy information is displayed alongside them. In US studies, respondents generally favor co-partisans over demographically similar candidates. When party and policy conflict, however, policy wins: partisans typically favor the out-partisan who shares their policy views over the co-partisan who does not.

Keywords: candidate choice conjoint experiments; meta-reanalysis; descriptive representation; substantive representation

How voters choose among candidates for office in free and fair elections is a central question in political science, as it underpins the main democratic mechanism by which citizens are theorized to hold representatives accountable and translate policy preferences into government action (Dahl, 2008; Przeworski, Stokes, and Manin, 1999). Suppose voters did not care about a candidate’s policy positions, but instead only their race, their religion, their gender, their social class, their political party: any resulting absence of substantive representation (representatives advancing their constituents’ policy interests) would compromise the electoral connection between voters and their representatives (Pitkin, 1967). Suppose, by contrast, voters did not care about a candidate’s background but instead *only* their policy positions: any resulting absence of descriptive representation (representatives sharing salient social characteristics with their constituents), especially for historically marginalized groups, would erode democratic legitimacy, political inclusion, and trust (Pitkin, 1967; Mansbridge, 1999; Phillips, 1995). Actual voter behavior is likely to fall between these poles, and in subtle ways, since voters can use social group membership to make informed guesses about policy positions and since policy stances themselves can signal “authentic” social group membership.

Empirically assessing how voters in fact choose among candidates in real elections is inherently difficult because candidate attributes are correlated due to causes both observed and unobserved. Differences in vote shares between men and women candidates, for example, reflect not only the effects of gender, but also the effects of other characteristics that co-vary with it, including experience, political affiliations, policy positions, and age. Leading theories of candidate emergence suggest that individuals self-select into politics in structured ways (Black, 1972; Rohde, 1979; Jacobson and Kernell, 1983; Fowler and McClure, 1989; Maestas et al., 2006; Lawless and Fox, 2010). An observational comparison of the men and women who end up standing for election therefore cannot separate the effects of gender from the effects of selection; selection confounds observational comparisons of the many other attributes voters

might weigh in just the same way.

In response to this threat of omitted variable bias, political scientists have turned to candidate choice conjoint experiments (hereafter “conjoint”), introduced to political science by Hainmueller, Hopkins, and Yamamoto (2014). In the standard design, respondents are presented with a series of pairs of hypothetical candidate profiles that vary along a small number of randomized attributes and are asked to choose between them. These experiments aim to simulate the experience of choosing among real candidates at the ballot box; how much conjoint can tell us about real elections depends on how well they simulate them.

Since their introduction to political science, conjoint have been deployed at a tremendous scale. The rapid expansion of this design has generated a large and diverse body of studies, spanning a wide range of countries, samples, and candidate profiles. Even without explicit coordination, researchers have fielded structurally similar designs again and again. That repetition provides a rare opportunity to test whether findings generalize across contexts, and it is what the present study takes up.

Here we synthesize the experimental record using the tools of “meta-reanalysis” (Galos and Coppock, 2023). In a meta-reanalysis, original study datasets are cleaned and standardized, then reanalyzed using common estimation procedures before the resulting estimates are combined in a meta-analytic framework. To date, meta-reanalysis has been applied to experiments on gender-based hiring discrimination (Galos and Coppock, 2023), preferences over prospective immigrants (Aviña et al., 2026), efforts to reduce expressive responding (Graham, 2026), randomized religious messages and endorsements (Sarkar and Coppock, 2026), and party cues (Coppock and Galos, 2025). Whereas most areas of political science rarely apply the same research design over and over, candidate choice emphatically does. And while this repetition is rightly or wrongly deplored by some who might sigh “not another conjoint,” the repeated application of the design yields a special view into the generalizability of conjoint experiments across settings.

Earlier work has already pooled conjoint results, one attribute at a time. Focusing on candidate gender, Schwarz and Coppock (2022) estimate an average effect of being a woman (relative to a man) candidate of 1.8 percentage points on vote choice. Eshima and Smith (2022) find that candidates in the oldest age category receive approximately 7 percentage points fewer votes than those in the second-youngest category. With respect to race, van Oosten, Mügge, and van der Pas (2024) show that average effects are close to zero overall, but that shared racial identity among minority respondents produces a substantial premium of 7.9 percentage points. Relatedly, Mikkelsen (2025) documents a temporal trend in which white Democrats’ support for Black candidates has increased by roughly half a percentage point per year since 1988. Finally, Simon and Turnbull-Dugarte (2025) find that holding a university degree increases candidates’ vote share by 5 percentage points. We recover each of these core findings (we note correspondences explicitly in the results section) while substantially expanding the underlying corpus of experiments analyzed in each case.

Beyond confirming and extending our existing knowledge of candidate attributes in isolation, the present study also places candidate attributes in direct comparison with one another. We collect together the effects of demographic attributes (age, gender, race, religion), political attributes (party, policy), and experience attributes (education, prior political office, occupation), using common measures and estimation strategies. With these estimates in hand, we can describe the relative effects of each attribute, and further, how voters navigate tradeoffs across them.

We highlight three main lessons from this broad synthesis of candidate choice conjoint experiments. First, political variables affect vote choice more than demographic variables on average, while experience variables usually fall in between. Second, attribute effects are systematically smaller in studies that supply policy information than in those that do not, a pattern that holds for demographic, experience, and political attributes alike. This cross-study contrast is descriptive, as any comparison across studies in a meta-regression must

be, but the corpus also contains a few studies that randomize whether policy information is shown, and these yield qualitatively similar estimates on a causal footing. Third, conjoint experiments are especially amenable to the study of how voters navigate explicit tradeoffs. Exploiting matchups in which one candidate matches the respondent on demographics but not politics and the other on politics but not demographics, we find that voters tend to prioritize political alignment over demographic similarity. The sharpest instance pits party labels against policy positions, with policy congruence inferred from respondent partisanship. When these cues conflict, Democrats choose the co-partisan only 41.2% of the time, and Republicans 48.4%. Both estimates therefore favor candidates expressing the respondent’s party-typical policy positions over those sharing their party label, although only the Democratic result is clearly distinguishable from indifference.

The special virtue of reanalyzing such a large corpus of randomized experiments is that we can directly assess the generalizability of each of these three lessons. By examining the distribution of effects across contexts, we find that in all cases we reject the null hypothesis of exact homogeneity across settings: quite evidently, the precise numerical estimates vary from study to study. That said, we can characterize how well the qualitative conclusions hold across contexts by estimating the sign and magnitude generalizability of effects. For example, the co-partisan advantage is particularly robust: 99 percent of study-level estimates can be affirmed to have the same direction. In the case of candidate gender, however, only 34 percent of study-level estimates can be affirmed to favor women, yet 90 percent can be affirmed to lie within 10 percentage points of the pooled average, so studies agree that the gender effect is small even when they cannot agree on its sign.

Theories of candidate choice

In one view of electoral selection, voters choose the candidate they expect to best advance their policy preferences, whether those positions are stated outright or inferred from party, background, and social identity. In another, voters derive utility from supporting candidates who share their social identities, independent of policy considerations. In this section, we briefly elaborate these two views and connect them to the empirical patterns conjoint experimentation might reveal.

Classical theories of candidate choice (e.g., Downs, 1957) conceptualize elections as a utility maximization problem in which voters select the candidate whose policy positions are closest to their own. Under this framework, voters who prefer liberal policies tend to choose liberal candidates over conservative candidates, and voters who prefer conservative policies tend to do the reverse. Although this choice can become complex when voters evaluate candidates across multiple policy dimensions, must infer candidates' positions from incomplete information, or themselves hold weak or internally inconsistent policy preferences (Davis, Hinich, and Ordeshook, 1970; Calvert, 1985; Enelow and Hinich, 1981; Shepsle, 1972; Alvarez, 1997), the core implication remains the same: electoral choice is fundamentally about selecting among competing bundles of policy commitments.

Within this classical view, non-policy candidate attributes, such as race, gender, age, religion, occupation, education, or prior experience, can influence voter support, but only indirectly. These attributes matter insofar as voters use them as heuristics or shortcuts to infer the policies a candidate is likely to support or implement if elected (Cutler, 2002; Popkin, 1991). Put differently, demographic characteristics and other non-policy cues affect vote choice not because voters intrinsically value those traits, but because they provide probabilistic cues about candidates' political orientations and capacities in environments where policy information is limited or uncertain.

A substantial body of empirical work supports this inference-based account. For example, studies show that voters often infer that Black candidates are more liberal than white candidates on average, particularly in the absence of clear partisan or ideological information (Wu and Huber, 2025; Karpowitz et al., 2021; Sigelman et al., 1995; Nelson, Sanbonmatsu, and McClerking, 2007; Karl and Ryan, 2016). These inferred differences shrink or disappear once clear ideological cues, such as party labels, are provided. Similarly, a large literature shows that voters tend to perceive women candidates as more liberal than men candidates, both when party labels are present and when they are absent (Abramson and Gillespie, 2024; Koch, 2000, 2002; Huddy and Terkildsen, 1993; Sanbonmatsu, 2002; McDermott, 1997). On this account, one mechanism through which race and gender shape vote choice is the policy expectations voters infer from them.

This inference-based mechanism also provides a bridge between descriptive and substantive representation. Under this classical view, descriptive representation can facilitate substantive representation precisely because voters expect candidates who share their social background to better understand, prioritize, and advocate for their group’s policy concerns. Thus, even when voters are ultimately motivated by policy outcomes, they may prefer descriptively similar candidates as a means of achieving substantive representation, especially under informational constraints. A related tradition in the study of Black political behavior emphasizes collective interests and social norms: Dawson (1995) argues that the Black utility heuristic, often referred to as “linked fate” (Gay, Hochschild, and White, 2016), leads Black Americans to evaluate political choices in terms of their perceived consequences for the group as a whole.

An alternative view, which we treat as equally plausible, holds that voters value descriptive representation for its own sake, independent of any policy inferences. Often rooted in social identity theory (Tajfel and Turner, 2004), this perspective emphasizes that individuals derive psychological benefits from supporting in-group members and may prefer candidates

who share salient aspects of their identity, such as race, gender, age, religion, or class background, because doing so affirms belonging, shared perspective, or the symbolic standing of the group (Wallace, 2014; Lublin and Wright, 2024; Manzano and Sanchez, 2010; Schildkraut, 2013; Stokes-Brown and Dolan, 2010; Dassonneville and McAllister, 2018). Under this view, demographic-identity attributes should exert stable and independent effects on vote choice, even when policy information is available, including in settings where voters face competing cues.

This identity-based logic may also extend to partisanship. Voters may prefer in-party candidates over out-party candidates not because of policy agreement, but due to a generalized preference for in-group members over out-group members (Iyengar and Westwood, 2015). In this sense, partisanship can function not only as a signal of policy alignment, but also as a social identity (Green, Palmquist, and Schickler, 2004).

Observational Equivalence and Mechanisms

Adjudicating between these accounts is complicated by an observational equivalence problem. For example, in a meta-reanalysis of candidate choice conjoint experiments focused on white Democratic voters, Mikkelsen (2025) documents a steady increase in support for Black candidates. This pattern could reflect a genuine rise in taste-based preferences for descriptive representation (a preference for Black candidates as such), but it could also reflect changing policy preferences or shifting inferences about candidates' policy positions. Nonexperimental evidence admits similar ambiguity: changes in evaluations of women and racial-minority representatives may reflect intrinsic preferences, or changing beliefs about competence, policy positions, or group advocacy (Weissman, 2025). We therefore remain agnostic about the underlying mechanisms, both theoretically and empirically. Our study can tell us how large each effect is and how much it varies across contexts but it cannot tell us whether a preference is held for its own sake or because of the policies voters expect.

This tension mirrors debates in the study of discrimination in other domains. Scholars distinguish between “statistical” discrimination, where decision-makers rely on group stereotypes to infer (correctly or not) unobserved characteristics, and “taste-based” discrimination, where decision-makers hold intrinsic preferences for or against particular groups (Phelps, 1972; Guryan and Charles, 2013). Statistical discrimination parallels the inference-based account of voter behavior, while taste-based discrimination parallels the identity-affinity account. Despite early optimism that factorial experimental designs could distinguish between these mechanisms, subsequent work shows that such distinctions remain empirically elusive (Heckman, 1998; Neumark, 2012). Conjoint experiments therefore reveal how voters respond to attributes, but not the motivations underlying those responses.

The Role of Information and Cross-Pressures

If demographic and partisan cues function as heuristics for unobserved policy positions, their effects should systematically attenuate when policy information is made explicit (Cutler, 2002; Popkin, 1991; Thomander, 2024; Malesky and Schuler, 2020). This expectation is consistent with evidence that voters rely more heavily on stereotype-based and other non-policy cues in lower-information electoral settings, where partisan cues, prior elected experience, or clear policy signals are less available (Thomander, 2024; Malesky and Schuler, 2020). Alternatively, if voters value descriptive similarity intrinsically, the effects of demographic attributes should persist even in the presence of clear policy information.

A related question is how voters resolve tradeoffs when they are “cross-pressured” by competing identity and political cues (Huddy, Mason, and Aarøe, 2015; Roccas and Brewer, 2002; Klar, 2013). A classical account predicts that voters will typically favor the in-partisan candidate, since party signals the policy agenda a candidate would pursue (Rahn, 1993; Lupia and McCubbins, 1998; Popkin, 1991; Downs, 1957). An identity-affinity account suggests that shared demographic ties may sometimes pull voters away from their partisan group.

However, partisanship can function both as a policy signal and as a social identity, and these roles need not align. When a candidate’s party label and stated positions conflict, the choice reveals how respondents weigh those two cues; it does not separately identify the role of partisan identity or respondents’ own policy agreement. Likewise, a tradeoff between demographic similarity and co-partisanship identifies which cue prevails, not the motivations behind that preference.

Design considerations for candidate choice conjoint experiments

The basics of candidate choice conjoint experimentation are well covered elsewhere, but we recapitulate the main points here. Survey respondents are recruited either through random sampling from a panel (rare) or through opt-in websites (common). After subjects’ pre-treatment covariates are measured, they are introduced to the conjoint choice environment. In the standard design (the only variant we consider in this paper), subjects are presented with two candidates per task and are asked to make a binary choice between them on the basis of a small set of attributes, each of which has a randomized set of levels. Subjects typically complete multiple tasks. Figure 1 shows the distribution of the number of tasks and attributes in our corpus. The median number of tasks is 5 and the median number of attributes is 8. Five studies randomize whether an attribute is presented or how many attributes appear, so they have no fixed attribute count and are omitted from the right panel of Figure 1; all 243 appear in the left panel.

Typically, the data produced by conjoint experiments are reshaped so that the candidate profile is the unit of analysis. Researchers then estimate an ordinary least squares regression of the binary outcome on candidate features with standard errors clustered at the respondent level. The coefficients returned by this regression are unbiased estimates of the Average Marginal




Figure 1: Number of tasks and number of attributes

Note: Distributions cover conjoints from all countries, not US studies only. Left panel: the number of binary choice tasks presented to each respondent, for all 243 conjoint experiments in the corpus; the terminal bin collects the 15 studies that present more than ten. Right panel: the number of randomized candidate attributes per profile, for the 238 studies with a fixed attribute count; the five studies that randomize whether an attribute is presented, or how many appear, are omitted. Study-by-study task counts are in Appendix Table C.2.

Component Effect (AMCE), the standard estimand in conjoint experiments (Hainmueller, Hopkins, and Yamamoto, 2014). The AMCE for a pair of attribute levels is the average change in the probability that a profile is chosen when one level is swapped for the other, averaging over the randomization distribution of every other attribute in the design. An AMCE of 0.03 for candidate gender, for example, says that making a candidate a woman rather than a man raises the chance that profile is picked by three percentage points, averaged across all the other candidates a respondent might have seen. Whether the AMCE itself is conceptually the “right” estimand has been the subject of serious methodological debate. Abramson, Koçak, and Magazinnik (2022) point out that a positive AMCE does not indicate a strict preference for a candidate feature, since it can be influenced by minority preferences held more strongly than those of the majority. Bansak et al. (2023) respond that this property is a feature, not a bug, since it correctly accounts for the preference intensities that presumably influence real-world vote choice.

In a separate line of critique, Abramson et al. (2026) describe how the AMCE of level 1 versus level 2 of an attribute can depend on the comparison of levels 1 and 2 to the other levels of an attribute. These authors therefore recommend an estimand they call the “Average Feature Choice Probability” (AFCP), which is defined with respect to tasks that include only profiles that differ on the relevant attribute levels and excludes all tasks in which one or both of the candidates has an attribute level not included in the comparison. We adopt the AFCP as our estimand. As shown in Abramson et al. (2026), for binary attributes (comparisons of men candidates to women candidates, for example), the AFCP is equivalent to the AMCE up to rescaling; the distinction matters more for multi-level attributes (race, age categories) where the AMCE would additionally incorporate indirect comparisons.

The rescaling depends on how a design pairs the two profiles. When each profile’s levels are drawn independently of the other’s, $AMCE = AFCP - 0.50$. Some studies that otherwise follow the standard design instead force the two profiles to differ on an attribute: in every task

one candidate is a Democrat and the other a Republican, and which profile carries which party is randomized. The literature has no settled name for this arrangement. Authors describe it in words (“One candidate was always a Democrat and the other a Republican,” Leeper and Robison, 2020), and the restrictions discussed by Hainmueller, Hopkins, and Yamamoto (2014) rule out implausible combinations within a profile rather than governing how profiles are paired. We call it *forced opposition*. Twenty-one of our 243 studies, from 16 papers, force opposition on at least one of the binary attributes we analyze, most often party (ten studies) and race (six). Under forced opposition every task is an A-versus-B task, so a win for the A profile is a loss for the B profile, and the AMCE equals $2 \times (\text{AFCP} - 0.50)$, twice its value under independent draws. The AFCP is unaffected, since it is defined on A-versus-B tasks whatever the design does with the rest. When pooling across designs that pair profiles differently, this invariance is a further point in the AFCP’s favor.

Random assignment balances the other attributes across the A and B profiles in expectation, not in any particular sample. In the 2,248 woman-versus-man tasks of Atkeson and Hamel (2020), for example, 21.4% of the women but 19.5% of the men were school guidance counselors, and 19.1% of the women but 20.7% of the men were business owners. The raw AFCP absorbs such imbalances along with the effect of gender; adjusting for them moves that study’s estimate from 0.566 to 0.557. Appendix P re-estimates every study-level AFCP with a regression adjustment for the study’s other randomized attributes. The adjustment is a mild correction: across 1,815 estimates the median change is 0.18 percentage points, and 93% of estimates move by less than half of their standard error. We report the raw estimates throughout.

Regardless of whether we target AMCEs or AFCPs, a key design question for conjoint experiments is the amount of information about the candidates to provide to subjects. Bansak et al. (2021a) dub this dilemma the “satisficing-masking tradeoff.” Too much information and respondents could become overloaded, taking cognitive shortcuts to answer the ques-

tions without putting much thought into them (satisficing). Not enough information and respondents could be thinking about unmentioned, relevant characteristics related to the attributes that they do see (masking). Seen this way, the Clayton et al. (2026) measurement error critique concerns satisficing. By their estimates, at least some subjects answer conjoint tasks by clicking haphazardly, typically resulting in choice probabilities closer to 50%. Abramson and Gillespie’s (2024) separate critique of conjoint experiments is about masking: their concern is that we do not learn the real effect of gender, since respondents use gender to form beliefs about unobserved attributes such as ideology or party.

Our somewhat contrary view is that neither satisficing nor masking generates bias *per se*. Instead we view each point on the tradeoff as corresponding to a different set of estimands. Under their assumptions (random assignment, noninterference, excludability), conjoint experiments generate unbiased estimates of their estimands; this property was proven for the AMCE by Hainmueller, Hopkins, and Yamamoto (2014), and it holds for other related quantities of interest, including the AFCP. If subjects are satisficing because they are overloaded with information and all choices between candidates are coin flips or deterministic button-mashing, then we correctly learn how subjects behave in an overloaded information environment. If excluded attributes are “masked” by included attributes, then we correctly learn about the total effect of the included attributes, inclusive of any inferences drawn about excluded attributes.

Some elections for local office feature nonpartisan candidates with virtually no name recognition among voters, who make their choice in the voting booth with no more information than the name of the candidate from which some demographic information can be vaguely inferred. Conjoint with few attributes might target this estimand well. Other campaigns for office successfully communicate partisanship, some policy positions, and some demographic information; these elections might be well simulated with conjoint experiments with a moderate number of attributes that cover the main sorts of information communi-

cated by the campaign to a distracted public. At the other extreme, high-salience elections such as recent US presidential contests expose voters to a much richer mix of policy and non-policy information. Conjoint with more attributes may better approximate these settings, although they also increase the risk of respondent overload.

Meta-reanalysis design

Our study is a “meta-reanalysis” (Galos and Coppock, 2023), which means we first gathered the original datasets for all included studies, cleaned and standardized them, and then calculated win rates (AFCPs) for the relevant matchups in each study.

Our goal was to assemble the candidate choice conjoint experiments we could identify and reanalyze from available replication data. Eligible studies used a paired, binary forced-choice design in which at least two candidate attributes were randomized and respondents selected between candidates for electoral office. We refer to this as the “standard design.” Studies that force opposition on an attribute meet this definition, since which profile carries which level is still randomized. We do not consider here variations on the standard design, such as designs that hold one or more of one profile’s attributes constant, vignette-based candidate choice studies, or single (or triple or quadruple) profile designs. We searched for studies using tools like Google Scholar to look through the literature in political science. We also searched online data repositories, including the Time Sharing Experiments in the Social Sciences (TESS) database, OSF, and the Harvard Dataverse. Additionally, we located some studies by following citations from other works and corresponding with the original authors. Replication data that were not available online were requested from the corresponding author. We followed the PRISMA 2020 reporting guidelines in documenting this search, screening, and inclusion process (Page et al., 2021); the corresponding diagram is presented in Appendix Figure B.3. Throughout, a study is a single conjoint design fielded on a single sample. A

paper contributes more than one study when it fields separate samples, or when it fields the same design in several countries, in which case each country enters as its own study. A paper that randomizes some feature of the design within one sample still contributes a single study. In total, we collected data from 243 different studies drawn from 144 papers and conducted in 52 countries, representing over 2,176,000 unique choice tasks and over 434,000 unique respondents. These studies include online quota panels, probability and convenience samples, student samples, and samples of political elites. The full table listing the studies, where they were conducted, the sample types, the number of respondents, tasks, and profiles is too long for the main text, so we present it in the appendix (Appendix Table C.2).

The experiments we collected considered many different attributes; we necessarily focus here on those attributes with the largest “estimate density,” by which we mean those attributes with the greatest overlap across studies. We generated two sets of attributes that we could compare, one for the set of all studies, regardless of the country of origin, and one for the set of studies conducted in the US only. As shown in Figure 2, we compare eight attributes falling into three broad families: demographic identity (gender, age, race, and religion), experience (education, prior political office, and occupation), and politics (party). In the US analyses we add policy positions to the political family. The families summarize the overall pattern, and every attribute is also reported on its own, so readers who would sort an attribute differently can read its estimates directly. Three assignments are arguable. Occupation and education are included under experience because, unlike ascribed traits such as race or gender, they reflect a life course of choices, though both also mark social identity (which is why we use education as a matching trait in the analyses of shared characteristics). Prior political office sits with experience, although a record in office also conveys policy information. Because of the extraordinary diversity of study countries, we study only average effects for the international comparisons, and do not subdivide those samples by individual difference variables. The only exception to this rule is the “in-party” versus “out-party”

comparison, which relies on a measure of subjects’ own partisan affiliation.

To make AFCPs comparable across studies, we coarsen study-specific attribute levels into common categories before estimating pairwise task-level win rates. Age is standardized to young, middle, and old. Religion is standardized to Christian, Jewish, Muslim, and none; because only seven studies randomize a Jewish candidate, the pooled comparisons are among Christian, Muslim, and none. Outside the United States, race is standardized only as majority versus minority because country-specific racial and ethnic categories do not map cleanly across settings; in the US analyses, we retain the more specific white, Black, Hispanic, and Asian categories. Education is standardized to Bachelor’s degree or more versus no Bachelor’s degree; political experience to any prior political experience versus none. Occupational backgrounds are standardized into working class, professional (chiefly lawyers, doctors, and engineers), business/managerial, educator, farmer, and military, with residual “other” categories omitted from the main comparisons. In US studies that randomize policy positions, we collapse each profile’s policy content into a single overall stance: liberal, conservative, or mixed. The full crosswalk from study-specific levels to these standardized categories is provided in the replication materials.

Among the US studies only, we can disaggregate candidate race and party somewhat and we can add a ninth attribute, policy liberalism. Because of the relative homogeneity across the US studies, we can subdivide these samples by characteristics of individuals. Here we split US samples by respondent partisanship, considering Republicans and Democrats (including leaners where specified) separately for all comparisons.

All our estimates are computed at the task level, not the profile level. We ask how frequently respondents choose a candidate of type A when given a choice between a candidate of type A and a candidate of type B. Type A and type B can vary on one dimension (Black candidates versus white candidates) or on multiple dimensions (Black Republican candidates versus white Democratic candidates). We cluster our standard errors by respondent. For




Figure 2: Attribute levels by study (international comparisons)

Note: Each row is one study and each column is one of eight candidate attributes; 238 studies randomize at least one of the eight. A dot indicates the study randomized that attribute level; dot size is proportional to the number of tasks in which the first profile carried it. Studies are divided into US and non-US panels. The studies shown here are listed in Appendix Table C.2.




Figure 3: Attribute levels by study (US comparisons)

Note: Each row is one US study and each column is one of nine candidate attributes. A dot indicates the study randomized that attribute level; dot size is proportional to the number of tasks in which the first profile carried it. The studies shown here are listed in Appendix Table C.2.

US studies, we assign respondent weights by iterative raking to 2022 American Community Survey targets for gender, age, race, and education whenever the replication data record the necessary demographics, which they do for 61 of the 86 US studies; the remaining 25 are analyzed unweighted. Six of those 25 do supply survey weights of their own, and we set them aside deliberately so that every weighted US estimate rests on the same benchmark. Outside the United States we retain whatever weights a study supplies, as 68 of the 157 non-US studies do, and analyze the rest unweighted. We do not rake non-US samples to national benchmarks because comparable targets are not consistently available across the corpus’s many national contexts.

We apply a random effects model to calculate the meta-analytic estimates. The model presumes that the true distribution of effects is characterized by a normal distribution with mean μ and standard deviation τ . We report estimates of both μ and τ . We also report the p -value of a Q-test against the null hypothesis of residual effect homogeneity across sites. We estimate sign generalizability (Egami and Hartman, 2023; Slough and Tyson, 2024), with a procedure adapted to answer what fraction of the estimates fall to the same side of 50% as the estimate of μ . We estimate magnitude generalizability as the fraction of estimates that can be affirmed to fall within 10 percentage points of the average effect estimate.

A natural concern about any meta-reanalysis is whether our reanalyzed estimates faithfully recover the effects reported in the original papers. We validate our approach by comparing our study-level AFCPs against a set of published estimates gathered from the original papers. To make the comparison, we translate published values to the AFCP scale. For a binary attribute with levels A and B, a published AMCE equals $m \times (\text{AFCP} - 0.5)$, where $m = 1 + \Pr(\text{opponent is B} \mid \text{A}) - \Pr(\text{opponent is B} \mid \text{B})$ records how the design pairs the two levels: $m = 1$ under independent draws, $m = 2$ under forced opposition, and other values for designs that make A-versus-B tasks more or less common than independent draws would. We compute m from each study’s task data and set $\text{AFCP} = 0.5 + \text{AMCE}/m$; a marginal

mean (MM, Leeper, Hobolt, and Tilley, 2020) translates as $\text{AFCP} = 0.5 + 2(\text{MM} - 0.5)/m$. A published share of choices in A-versus-B tasks is already an AFCP and needs no translation. Estimates reported in designs that use conditional or interactive specifications (e.g., the party effect conditional on the number of policy agreements shared with the candidate) are flagged as not directly comparable and excluded from the validation. We also record how each article reports the contrast, whether as a printed number or only in a coefficient plot. Across the 341 comparisons an article states numerically, our re-estimates recover the published value closely: the median absolute discrepancy is 0.5 percentage points, and 97% of comparisons fall within 5 points. For the 76 comparisons an article displays only in a figure and from which a single value can be read, agreement is looser but still close (median 0.6 points, 99% within 5 points). Appendix Figure E.4 presents both comparisons, which we take to be strong evidence that we have reanalyzed the corpus correctly. When we investigate the largest of the discrepancies, they turn out to depend on (in our view all defensible) weighting, subsetting, and construct-definition decisions made differently by original analysts and by us: the single largest, 17 points, arises because the original authors count a candidate as a co-partisan at the level of the electoral coalition while we match on the party label itself.

Beyond validation, this exercise shows how rarely the comparisons we report appear, in the form we report them, in the papers that produced them. Of the 1,705 study-comparison pairs included in our publication-reporting audit, only 341 (20.0%) are stated numerically in the original article. A further 209 appear only inside a coefficient plot, so even by that looser standard just 550 (32.3%) have a published analogue and 1,155 (67.7%) do not. Reading each of those 1,155 against their source articles shows that very little of the shortfall is a number the original authors withheld. For only 212 of them did the original produce nothing at all: 34 because no article exists, and 178 because the attribute is in the design but appears in no results display. The other 943 were analyzed and published, but at a coarser or finer

level of aggregation, for a different subsample, or as a different quantity, most often because a category we pool is one the article reports level by level. The value of working from the microdata lies less in surfacing hidden numbers than in harmonizing the level of aggregation: a literature that reported its results at many different resolutions can be brought onto a single common scale, which is what makes the estimates in this paper comparable across studies at all.

Finally, a concern that attends any attempt to learn from the published record but does so especially keenly for meta-analyses is the publication filter. We address publication bias in two ways. First, we do not condition our search on publication status; that is, we include experiments from working papers so long as they meet our design criteria. Second, because we recompute each matchup from the microdata regardless of whether the original paper reported it, an estimate does not enter the corpus by virtue of being statistically significant. Even so, we can assess the extent of publication bias with a three-parameter selection model (Vevea and Hedges, 1995), which provides a test of selection and a selection-adjusted estimate of μ (Appendix K). Under that model, only the women-versus-men comparison shows statistically significant evidence of selection, and the estimated selection favors *non*-significant results. The selection-adjusted AFCPs strongly track the unadjusted estimates closely in all cases. In the following section, we present estimates that have not been adjusted for selection.

Results

Demographic, Experience, and Political Attributes

We begin by evaluating the relative importance of demographic, experience, and political attributes in shaping candidate choice. Figure 4 summarizes the distribution of study-level estimates across all candidate conjoint experiments in our sample. Each row corresponds

to a matchup between two attribute levels, and the horizontal axis displays the vote share for the “first” candidate. The point estimate and confidence interval in each row refer to $\hat{\mu}$, the average of the distribution from which study-level estimates are drawn, while the fitted normal distribution reflects between-study variation.

Across these matchups, political attributes generally have larger effects on vote choice than demographic attributes, with experience attributes often falling between the two. The in-party advantage is the largest and most consistent contrast.

Within the demographic family, effects are generally modest and often close to indifference. Respondents prefer younger candidates to older candidates, but preferences by gender are comparatively small and race-based matchups are often near coin flips, consistent with prior work (Schwarz and Coppock, 2022; van Oosten, Mügge, and van der Pas, 2024). Appendix Figure O.20 shows that the pooled international gender and majority-minority race averages summarize real country-to-country heterogeneity, especially for gender, even though most country-level estimates remain substantively close to 50%. Fewer experiments investigate the effects of candidates’ religious affiliation, so we are less certain of these estimates. That said, Christian versus none is centered almost exactly on 50%, whereas Christian versus Muslim and Muslim versus none imply larger average differences on the available evidence.

Overall, the experience attributes show somewhat stronger effects: candidates with a college degree perform better on average, while candidates without prior political experience face a substantial penalty. Occupational background contrasts are modest: business/managerial, educator, farmer, military, and working-class candidates all cluster relatively close to professional candidates.

Political attributes exert much larger effects. Respondents choose in-party candidates over otherwise similar out-party candidates 72.5% of the time. In the main, respondents weigh political attributes far more heavily than demographic ones.

Table 1 summarizes these comparisons in terms of the estimated mean (μ) and stan-




Figure 4: Meta-analytic estimates of win rates by matchup

Note: Each row shows the distribution of study-level AFCP estimates for one matchup type across all countries. The shaded curve at left is the fitted normal distribution from the random effects model, labeled with its mean $\hat{\mu}$ and between-study SD $\hat{\tau}$ on the AFCP percentage scale; the point and interval show $\hat{\mu}$ with its 95% confidence interval. The dashed vertical line marks 50%, the point of indifference between the two candidates. Gold: matchups from studies that include explicit policy information. Blue: matchups from studies without. k at the right of each row is the number of contributing study-level estimates, which exceeds the number of studies wherever a study randomizes whether policy information is shown and so contributes an estimate in each condition. Complete numerical results for every point plotted here are in Appendix Table Q.20.

dard deviation (τ) of the estimated normal distribution. The table also reports the p -value from the Q-test against the null of homogeneity across sites; that null is rejected for every matchup. The “sign generalizability fraction” column reports the fraction of estimates that can be affirmed to lie on the same side of 50% as the estimated μ . Among the matchups estimated from many studies, sign generalizability is strongest for in-party versus out-party comparisons at 99%, indicating near-universal agreement in direction across studies; the only fractions above it are the 1.00 recorded for Christian versus Muslim and Republican versus Independent, resting on nine and eight studies respectively. Sign generalizability is weakest for the five pairwise race matchups, which sit between 2% and 10% because study-level estimates cluster close to indifference; the pooled majority-versus-minority comparison, at 33%, is the exception. For women versus men specifically, the sign generalizability fraction is only 34%. The distinction here is between agreement and affirmation. Fully 86% of the study-level women-versus-men estimates lie above 50%, but most are individually too imprecise to rule out the other side, so only about a third can be affirmed by the sequential procedure. A low sign generalizability fraction can therefore reflect imprecision rather than disagreement in estimated direction. The final column reports a complementary *magnitude* generalizability fraction: the share of study-level estimates affirmed to lie within ± 10 percentage points of $\hat{\mu}$ by the same Egami–Hartman sequential procedure (Egami and Hartman, 2023; Slough and Tyson, 2024). High magnitude generalizability means that studies agree not merely on the sign of an effect but on its rough size; Appendix Table D.3 shows that these fractions move in the expected direction as the tolerance band is tightened to ± 5 or widened to ± 15 points.

Figure 5 focuses on the US studies and splits estimates by respondent partisanship. Across most demographic and experience attributes, responses are strikingly similar across partisan groups: Republicans and Democrats evaluate candidates of different ages, education levels, and experience levels in nearly identical ways, and four of the five occupational

Table 1: Sign and magnitude generalizability by matchup (all studies)

Attribute	k	μ (SE)	τ (SE)	Q_p	Sign gen.	Mag. gen.
Demographics						
Women vs Men	218	0.53 (0.00)	0.04 (0.01)	<0.001	0.34	0.90
Young vs Old	114	0.58 (0.00)	0.05 (0.02)	<0.001	0.61	0.53
Young vs Middle	133	0.51 (0.00)	0.03 (0.01)	<0.001	0.11	0.80
Middle vs Old	95	0.58 (0.01)	0.05 (0.02)	<0.001	0.63	0.58
Majority vs Minority	106	0.53 (0.01)	0.06 (0.02)	<0.001	0.33	0.59
White vs Black	61	0.48 (0.01)	0.03 (0.02)	<0.001	0.10	0.64
White vs Hispanic	49	0.50 (0.00)	0.02 (0.01)	<0.001	0.02	0.73
White vs Asian	38	0.50 (0.01)	0.04 (0.02)	<0.001	0.05	0.61
Asian vs Hispanic	35	0.52 (0.01)	0.05 (0.03)	<0.001	0.09	0.43
Black vs Hispanic	46	0.52 (0.01)	0.03 (0.02)	<0.001	0.04	0.61
Christian vs None	22	0.49 (0.01)	0.03 (0.02)	<0.001	0.09	0.68
Christian vs Muslim	9	0.61 (0.01)	0.03 (0.03)	<0.001	1.00	1.00
Muslim vs None	9	0.37 (0.01)	0.03 (0.02)	<0.001	0.89	0.56
Experience						
College Degree vs None	55	0.57 (0.01)	0.06 (0.03)	<0.001	0.67	0.75
No vs Any Political Experience	120	0.43 (0.01)	0.06 (0.02)	<0.001	0.62	0.57
Business/Managerial vs Professional	110	0.48 (0.00)	0.03 (0.01)	<0.001	0.12	0.53
Educator vs Professional	75	0.50 (0.00)	0.03 (0.02)	<0.001	0.03	0.40
Farmer vs Professional	35	0.51 (0.01)	0.02 (0.02)	<0.001	0.03	0.37
Military vs Professional	14	0.49 (0.01)	0.02 (0.02)	0.019	0.07	0.64
Working Class vs Professional	35	0.49 (0.01)	0.05 (0.02)	<0.001	0.14	0.60
Politics						
Conservative vs Liberal Issue Positions	95	0.43 (0.01)	0.09 (0.04)	<0.001	0.55	0.37
Conservative vs Mixed Issue Positions	76	0.44 (0.01)	0.07 (0.03)	<0.001	0.49	0.53
Liberal vs Mixed Issue Positions	76	0.53 (0.01)	0.08 (0.03)	<0.001	0.37	0.53
In-party vs Out-party	79	0.73 (0.01)	0.08 (0.03)	<0.001	0.99	0.46
Republican vs Democrat	49	0.47 (0.01)	0.04 (0.02)	<0.001	0.45	0.73
Republican vs Independent	8	0.38 (0.02)	0.04 (0.03)	<0.001	1.00	0.75
Democrat vs Independent	8	0.48 (0.01)	0.03 (0.03)	0.002	0.12	0.88

Estimates cover conjoints from all countries, not US studies only. Each row summarizes one matchup across all studies in the corpus. μ is the meta-analytic mean AFCP and τ is the between-study standard deviation (REML). Q_p is the p -value from Cochran's Q test of residual homogeneity. Sign gen. is the fraction of study-level estimates affirmed to lie on the same side of 50% as $\hat{\mu}$ by the Egami–Hartman sequential test. Mag. gen. is the fraction affirmed to lie within ± 10 percentage points of $\hat{\mu}$ by the same procedure. Standard errors in parentheses.

contrasts likewise. Differences do emerge for gender and race, and for business/managerial candidates, but all are modest in magnitude. Democrats are somewhat more likely to support women candidates (54.2%) than Republicans (49.1%), and they are less likely to favor white candidates over Black candidates (46.6% versus 52.2%).

Importantly, the pooled US race averages obscure meaningful heterogeneity by respondent race. Appendix Figure O.27 shows that in white-versus-Black matchups, white respon-




Figure 5: Meta-analytic estimates of win rates by matchup (US studies, by respondent partisanship)

Note: Same design as Figure 4 but restricted to US studies, split by respondent partisanship (Democrats and Republicans). US estimates are weighted to 2022 American Community Survey benchmarks on gender, age, race, and education where the requisite variables are available; otherwise, they are unweighted. Complete numerical results for every point plotted here are in Appendix Table Q.21.

dents slightly favor the white candidate on average (50.8%), whereas non-white respondents slightly favor the Black candidate (46.4% choosing the white candidate). A separate shared-identification analysis, closest to the approach in van Oosten, Mügge, and van der Pas (2024), shows that across studies with compatible respondent and candidate race or ethnicity measures, the pooled same-race premium is small overall (3.0 percentage points), close to zero among white respondents (0.9 points), but much larger among Black respondents (9.3 points) and moderate among Hispanic respondents (5.4 points; Appendix Table O.19 and Figure O.28). Appendix Figure O.30 reproduces the trend reported by Mikkelsen (2025): among white Democrats, support for the Black candidate has increased by about 0.8 percentage points per year ($p = 0.004$). The slope among white Republicans is similar (0.7 points per year, $p = 0.06$), and the two party slopes are not distinguishable from each other.

Political attributes produce stark divergence. Republicans and Democrats overwhelmingly favor candidates aligned with their own party and policy positions: Republicans choose the Republican candidate over the Democrat 72.4% of the time, while Democrats choose that same Republican candidate only 27.2% of the time. Similarly, Republicans prefer candidates with conservative positions over liberal ones 65.6% of the time, compared with 28.7% among Democrats.

In our view, the main takeaway from Figure 5 is that demographic and experience attributes have modest effects that do not differ greatly by respondent partisanship; whatever differences there may be pale in comparison to the differences by candidate partisan affiliation. Study-by-study forest plots for every matchup, split by country grouping and, for US studies, by respondent partisanship, are collected in Appendix O (Figures O.31–O.39).

Shared Identification

In this section, we consider how strongly voters favor candidates who share their own characteristics. Figure 6 reports, for every characteristic on which we can match candidates

to respondents, the rate at which voters choose the candidate who shares their trait over one who does not. Sharing a candidate’s party is by far the strongest source of affinity: co-partisans are chosen roughly 70% of the time. Sharing age, race, or gender moves choices only a few points off 50%, and shared education is indistinguishable from a coin flip. We further find that the co-partisan premium is larger when no policy information is available, consistent with party serving as a summary cue for policy.

The small average same-race premium masks substantial heterogeneity across respondents. Appendix Figure O.29 disaggregates the same-race preference by respondent race and shows that it is sizable for non-white respondents: among Black respondents the co-racial candidate is chosen about 59% of the time, approaching the roughly 70% rate at which voters choose co-partisans, while among white respondents the same-race premium is essentially zero.

Information as a Moderator of Candidate Evaluation

We next examine how these patterns depend on the informational environment in which candidate attributes are presented. We count a study as showing policy information when at least one randomized attribute states where a candidate stands on a policy issue. Attributes that give only an ideological label, or that name the issue a candidate would prioritize without saying what position they take on it, do not count. Where a study states positions on several issues, we collapse them into a single liberal, conservative, or mixed classification. Four studies randomize whether policy information appears rather than fixing it by design. For these we estimate separately within each condition, so that the tasks without policy information contribute to the no-policy estimates and the tasks with it contribute to the policy estimates. No estimate averages over the two conditions. Figure 7 provides a compact visualization of these effects by plotting Democratic estimates (horizontal axis) against Republican estimates (vertical axis), allowing us to distinguish attributes that generate shared versus




Figure 6: Shared identification across candidate characteristics

Note: Each row reports the meta-analytic AFCP of choosing the candidate who matches the respondent on the named characteristic over the candidate who does not, estimated study by study and pooled with a random effects model. Estimates cover conjoints from all countries, not US studies only. Slabs are the predictive distributions $N(\text{mean}, \text{between-study SD})$; the point and interval give the pooled mean and its 95% confidence interval; small points are study-level estimates. Black: all studies. Gold: studies that show candidate policy positions. Blue: studies that do not. The dashed line marks 0.5. US estimates use 2022 ACS weights where available and are otherwise unweighted. Complete numerical results for the pooled estimates plotted here are in Appendix Table Q.22.

divergent responses across partisans.

Two patterns emerge. First, most demographic and experience attributes tend to produce similar evaluations across partisans, whereas political attributes generate sharp partisan divergence. Estimates for age, education, experience, and the working-class versus professional comparison lie primarily in the top-right and bottom-left quadrants, indicating that Democrats and Republicans respond in the same direction. Gender and race cluster near the 50–50 point where the two reference lines cross, reflecting their relatively small effects on vote choice for both groups. Estimates for partisanship, by contrast, lie in the top-left quadrant, indicating that Democrats and Republicans evaluate the same candidates in opposite ways.

Second, estimates from studies that include policy information (gold) are often closer to the 50–50 point than those from studies without (blue), particularly for party and several demographic comparisons, though there are exceptions. See Appendix Table F.8 for formal hypothesis tests of the difference in AFCPs by the presence of policy information in the conjoint table. We emphasize that these contrasts represent descriptive differences: in this analysis we simply describe how conjoints with and without policy information (which may differ on other dimensions as well) generate different estimates.

Fortunately, the corpus also contains a few studies that randomize policy display within an otherwise fixed design: Bansak et al. (2018) (studies 1 and 2), Funck and McCabe (2022), and Peterson (2017) (study 1). We meta-analyze these designs to isolate the causal effect of policy information on other attribute estimates. Appendix Table F.9 shows that policy information substantially attenuates the in-party advantage, as well as several age, education, and experience contrasts. The effects are not uniform across attributes, but the randomized evidence broadly reinforces the descriptive pattern.

One might worry that the attenuation in the descriptive analysis of party cues under policy information is mechanical: the presence of an additional attribute of any kind would pull effects toward 50% (Abramson, Koçak, and Magazinnik, 2022). Appendix Table G.12

shows that dilution of this kind is not a general feature of these experiments: pooling across matchups, AFCPs are essentially unrelated to the number of randomized attributes (a slope of -0.001 per attribute, $p = 0.61$). Attenuation instead concentrates in the matchup with the strongest cue: the in-party advantage shrinks by 1.1 percentage points per additional attribute, while the penalty for political inexperience moves by 0.2 points, which we cannot distinguish from zero. Appendix Table G.13 separates the two explanations directly. Conditional on the number of attributes, displaying policy positions reduces the in-party AFCP by 9.9 percentage points, most of its unconditional effect of 10.4 points, while each additional attribute of any kind contributes a further 0.6 points, which we cannot distinguish from zero. Political experience shows the same pattern: policy display attenuates the penalty by 6.1 points, and the 0.1-point contribution of each additional attribute is indistinguishable from zero. Across studies, then, the smaller in-party AFCP under policy information is associated with policy content specifically rather than with longer profiles in general.

Cross-pressured Respondents

We now turn to the most direct test of how voters resolve competing considerations. Rather than varying one attribute at a time, the analyses in this section consider tasks in which respondents choose between candidates who match on different dimensions. Figure 8 illustrates this design by comparing candidates who match respondents on party but not gender to those who match on gender but not party.

Across the gender tradeoff panels, voters usually prioritize political alignment. For example, Democratic men choose women Democrats over men Republicans 71% of the time, with Republican men and Republican women both at 71%, and Democratic women at 68%. The race panels point in the same direction, though with a thinner evidence base. When demographic similarity and partisan alignment come into conflict in these tradeoff analyses, respondents generally favor the latter. As in previous analyses, these effects depend on




Figure 7: Matchup estimates among Democratic and Republican respondents, depending on the presence of policy information

Note: Each point is one study-level ACP estimate. Estimates are restricted to US studies, the only ones in which respondents are classified as Democrats and Republicans. The Democrat estimate (x-axis) and Republican estimate (y-axis) come from the same study and matchup type. Gold: studies that include explicit policy information; blue: studies without. Faint points with whiskers are study-level estimates and their 95% confidence intervals; the solid, larger point in each color is the pooled random effects mean for that group. The annotation gives that pooled point's coordinates (D on the horizontal axis, R on the vertical) and k , the number of contributing studies, so panels resting on few studies are identifiable as such. Dashed lines at 0.5 mark the no-preference baseline. Points in the upper-left or lower-right quadrants indicate opposite signs for Democratic and Republican respondents. Complete numerical results for the pooled estimates plotted here are in Appendix Table Q.23.

the informational environment. Estimates are closer to indifference when policy information is included. This pattern is consistent with an informational explanation, although the comparison does not identify why these cues become less influential.

Figure 9 extends this analysis across 12 sets of tradeoffs. The key comparisons are in the fourth column of panels, which reflect the direct choices between demographic similarity and political alignment. Across most specifications, respondents prefer candidates who match their political side over those who match them demographically. The pattern is evident for most gender and race subgroups, although the Black Republican estimates in the direct party-versus-race tradeoff fall below 50% in both policy-information conditions and are too imprecise to distinguish from indifference. The corresponding forest plots for the remaining panels are presented in the appendix.




Figure 8: Meta-analytic estimates of party versus gender tradeoffs

Note: Each panel is one respondent subgroup defined by gender and partisanship. Estimates are restricted to US studies. Candidate A matches the respondent on party but not gender; Candidate B matches on gender but not party. The estimand is the AFCP for Candidate A. Small points are study-level estimates; large points are random effects meta-analytic means with 95% confidence intervals. Values beside the pooled estimates are the AFCP with its standard error in parentheses. Study-level estimates are plotted without labels; their numerical values appear in the estimate-level supplement. Gold: studies with policy information; blue: without. Complete numerical results for the pooled estimates plotted here are in Appendix Table Q.24.




Figure 9: Meta-analytic estimates of candidate choice tradeoffs

Note: Each panel is one combination of demographic attribute (gender or race, the two panel rows) and comparison type (the six panel columns); within every panel the rows of points are respondent subgroups defined by partisanship and by the demographic attribute. Estimates are restricted to US studies. The six column facets compare candidates who match or mismatch the respondent on political alignment and on the demographic attribute; the fourth column, where one candidate matches only on politics and the other only on the demographic attribute, is the core tradeoff and its cells are shaded light gray to mark it. Political alignment is measured using party match. Small points are study-level AFCP estimates; large points are random effects meta-analytic means with 95% confidence intervals. Gold: studies with policy information; blue: without. Complete numerical results for every point plotted here are in Appendix Table Q.25.

We conclude with estimates of how respondents choose among candidates when party labels and policy positions point in opposing directions. Because respondents' own issue positions are not measured in most studies, we proxy policy congruence using their partisanship. We first collapse each candidate's randomized policy content into a single overall stance (liberal, conservative, or mixed) using the coarsening described above; candidates with mixed or offsetting positions are set aside for this analysis. We then classify liberal positions as congruent for Democratic respondents and conservative positions as congruent for Republican respondents. The key matchup pits a candidate who shares the respondent's party but expresses positions conventionally associated with the other party against an out-partisan whose positions align with the respondent's party. Figure 10 summarizes these tradeoffs.

The figure shows the common-sense pattern that respondents generally favor co-partisans when policy direction is held constant and candidates expressing their party's typical policy positions when party labels are held constant. In the fourth facet (where party and policy push in opposite directions), we see that Democratic respondents choose the Democrat with conservative positions over the Republican with liberal positions 41.2% of the time (95% CI: 35.1–47.3%). Republican respondents choose the Republican with liberal positions over the Democrat with conservative positions 48.4% of the time (95% CI: 44.0–52.8%). The Democratic result favors the out-partisan with party-typical policy positions; the Republican estimate does not differ clearly from an even split.

These results reveal an important limit to the influence of party labels. Although co-partisanship strongly shapes candidate choice, respondents lean toward party-typical policy positions when the two cues conflict. The Democratic estimate is clearly distinguishable from indifference; the Republican estimate points in the same direction but is less conclusive. Splitting each party by respondents' self-reported ideology, which 19 of the studies behind the fourth facet measure, shows that the pooled estimates combine two very different

groups: liberal Democrats choose the co-partisan with conservative positions only 38.2% of the time (95% CI: 31.7–44.7%) and conservative Republicans choose the co-partisan with liberal positions 47.7% of the time (95% CI: 39.9–55.4%), whereas conservative Democrats choose the co-partisan 63.6% of the time (95% CI: 58.9–68.3%) and liberal Republicans 61.4% of the time (95% CI: 52.2–70.6%) (Appendix Figure L.6). In other words, the cross-pressured partisans also side with the candidate whose positions match their own ideology, so, measured against respondents’ ideology rather than their party’s typical positions, the pull of policy over party is stronger than Figure 10 suggests. Nine of these studies also asked respondents their own positions on the candidates’ issues. Defining policy congruence by those positions, Democrats choose the co-partisan who disagrees with them 41.7% of the time (95% CI: 38.2–45.2%) and Republicans 43.0% of the time (95% CI: 40.0–46.0%), so the Republican estimate, too, now favors the out-partisan who agrees (Appendix Figure M.7).

Additional Analyses

In the appendix, we present additional analyses that probe the scope and robustness of our findings. One natural question is whether attribute effects vary systematically with the type or level of office being simulated. The dominant result, reported in Appendix Tables F.10 and F.11, is null. Across nearly all attributes and both moderators, effects are statistically indistinguishable from their baseline values. The within-study evidence on office type is especially thin, however, so these appendix results should be read as suggestive rather than definitive.

Election level tells much the same story. The two age matchups are the only demographic rows whose differences reach conventional significance, and both do so in regional elections: the young-over-old preference falls by 7.0 points, which leaves it at almost exactly indifference, and the young-over-middle preference falls by 4.6 points (and by 1.1 points in local elections). Only three other differences in the table are distinguishable from zero, for




Figure 10: Meta-analytic estimates of party-policy tradeoffs

Note: Each panel defines one comparison type between candidates who match or mismatch the respondent on party and on policy position; within each panel the rows are respondent partisanship (Democrat and Republican). Estimates are restricted to US studies, the only ones in which respondents are classified this way. The lower-left panel (the fourth comparison), where one candidate matches only on party and the other only on policy, is the key party-policy tradeoff and is shaded light gray to mark it. Small points below each row label are individual study-level AFCP estimates; large points are random effects meta-analytic means with 95% confidence intervals. The dashed vertical line marks 0.5 (indifference between candidates A and B). Complete numerical results for every point plotted here are in Appendix Table Q.26.

Democrat versus Independent, conservative versus liberal issue positions, and educator versus professional, all in local elections. Partisan preferences do not vary reliably with election level. The in-party win rate is larger by 2.3 percentage points in local elections and 9.7 points in regional elections, but neither difference is significant at the 5% level; the direction matches the argument that lower-salience contests, which offer less candidate information, place more weight on partisan cues, although these data cannot confirm it. The Republican win rate against Democrats, 0.473 nationally and so already below even odds, is 3.9 points lower in local elections, a difference that is likewise not significant.

We next examine the Clayton et al. (2026) satisficing critique directly on the subset of studies for which a correction is feasible from the existing data. Their model assumes that observed responses are a mix of real choices and coin flips, so measurement error due to satisficing will pull estimates closer to 50% and a correction will reinflate them away from 50%. Using repeated identical conjoint tasks within respondent, we estimate study-specific swapping-error parameters and recompute corrected AFCPs for directly feasible subsets of studies. Appendix Table I.15 and Figure I.5 show that the correction moves estimates modestly away from 50%, but does not overturn the basic ranking in which political attributes exert larger effects than demographic-identity attributes on the feasible subset. The largest visible shift is for in-party versus out-party matchups, which become somewhat stronger after correction. Because most studies do not contain recoverable repeated identical tasks, we treat this as a subset robustness check rather than a full-corpus correction.

We then apply an intransitivity audit in the style of Abramson et al. (2026). The concern is that aggregate pairwise win rates can imply a cyclic majority relation even when individuals have transitive preferences. Take three levels of one attribute, say young, middle-aged, and old candidates. If a majority chooses young over middle-aged and another majority chooses middle-aged over old, the aggregate ranking is cyclic if a majority nevertheless chooses old over young. We call three such levels a triangle, and a triangle is testable when the corpus

supplies all three of its pairwise AFCPs. A triangle whose win rates instead run in a loop, each level beating the one before it and losing to the one after, is a cycle.

Appendix Table J.16 shows that the pooled meta-analytic rankings are transitive in nine of the ten testable triangles; the exception (low, middle, and upper income) cycles only among pooled AFCPs whose confidence intervals all include 50%. At the study level, 76 of 431 complete triangles show a cycle in the point estimates. Only two survive the requirement that all three pairwise comparisons be individually distinguishable from 50% (Appendix Table J.17 lists them), so the rest are loops among win rates too close to indifference to order at all. Most observed aggregate cycles involve imprecisely estimated comparisons near 50%, providing little evidence of stable cyclic majority relations; this audit does not test whether individuals have coherent preferences.

Finally, we audit abstention or no-choice behavior in the small subset of studies where the raw archives preserve it. Appendix Table H.14 shows that recoverable no-choice rates are often nontrivial when the option is available, but these outcomes are too sparsely and inconsistently archived to support a pooled abstention meta-analysis for the full corpus.

Discussion

So what do we learn from a decade of candidate choice conjoint experiments? Across countries, offices, samples, and time periods, three patterns stand out. First, political attributes exert substantially larger effects on candidate choice than demographic attributes, while experience attributes usually fall in between; in the international comparison and among US Democrats the party matchups are the largest single-attribute effects we estimate, and among US Republicans only the Christian-versus-Muslim matchup, estimated from three studies, is larger. Second, attribute effects are smaller in studies that make policy information explicit. The comparison is descriptive, drawn across studies, but the few experiments

that randomize policy display within a fixed design confirm that at least part of the attenuation is caused by the policy information itself. Third, respondents generally favor co-partisanship over demographic similarity; when party labels conflict with issue positions conventionally associated with respondents' parties, Democrats favor the latter, while the Republican result is inconclusive when congruence is inferred from party and favors policy when it is measured by respondents' own positions. These patterns differ in the breadth of evidence that supports them. The direction of the co-partisan advantage is highly consistent across studies, though its magnitude varies considerably. Party is also the only political attribute randomized widely enough outside the United States to enter the international comparison, whereas the policy-position contrasts come from US studies. The tradeoffs and the respondent-level heterogeneity behind the third pattern, by contrast, rest on US studies with measured respondent partisanship and race.

In short, our findings more closely resemble a classical account of candidate choice (Downs, 1957), in which political considerations take precedence over demographic affinity. Political attributes produce the largest average effects, voters generally favor co-partisans over demographically similar out-partisans, and several nonpolitical cues weaken when explicit policy information is provided. These patterns are consistent with voters using demographic characteristics to infer candidates' likely policy positions (Cutler, 2002; Popkin, 1991; Lupia and McCubbins, 1998; Kirkland and Coppock, 2018; Arnesen, Duell, and Johannesson, 2019; Abramson and Gillespie, 2024). They do not uniquely identify this mechanism, since partisanship can also function as a social identity and demographic affinity can have intrinsic value (Pitkin, 1967; Mansbridge, 1999).

An alternative perspective, rooted in social identity theory (Tajfel and Turner, 2004) and related work (Wallace, 2014; Lublin and Wright, 2024; Manzano and Sanchez, 2010; Schildkraut, 2013; Stokes-Brown and Dolan, 2010; Dassonneville and McAllister, 2018; White and Laird, 2020), emphasizes that demographic attributes may matter not only because they con-

vey information, but because voters may derive intrinsic, non-policy utility from supporting in-group candidates. On this account, shared race, gender, religion, class background, or other identities can generate affective attachments, solidarity norms, and symbolic benefits that shape vote choice even when policy considerations are held constant. Our gender- and race-based tradeoff analyses speak directly to this claim: when voters face explicitly cross-pressured choices, they generally resolve the tradeoff in favor of political congruence. This pattern does not rule out identity-based motivations, but it does suggest meaningful limits on their influence in these conjoint settings.

At the same time, we reiterate that interpreting these findings is frustrated by an observational equivalence problem. For example, increased support for candidates from historically underrepresented groups among some voters (Mikkelsen, 2025; Weissman, 2025) could arise from multiple underlying processes, including changes in policy inferences, shifts in partisan coalitions, evolving social norms, or intrinsic preferences for descriptive representation. As in debates over statistical versus taste-based discrimination (Phelps, 1972; Guryan and Charles, 2013; Heckman, 1998), distinct mechanisms can produce similar observable behavior. Our evidence therefore speaks clearly to the structure of voter choices under controlled variation, but not to the underlying motivations that generate those choices. The distinction matters for comparison with Treger et al. (2026), who measure respondents' own policy preferences in a Canadian conjoint. Our main US tradeoff infers policy congruence from partisanship, so the two analyses address related but distinct comparisons; the own-issue analysis in Appendix M, which measures congruence directly in the eight studies that asked, is the closer analogue.

Our approach also illustrates the value of meta-reanalysis for assessing generalizability in political behavior research. By studying the same phenomena across many contexts, we can estimate not only average effects but also the extent of variation, and in some cases (as in our estimates of how policy information attenuates effects) we can even explain the

correlates of that variation. Working from original study datasets rather than published estimates is crucial: 1,155 of the 1,705 study-comparison pairs in our publication audit (67.7%) have no directly comparable published analogue, usually because the original papers report results at different levels of aggregation, for different subsamples, or as different quantities. Working from original data also enabled us to estimate AFCPs (rather than rely on the more commonly reported AMCEs) and to conduct the many robustness and alternative analyses that simply would not be possible without the microdata.

Beyond these specific findings, the standardized corpus is a resource that exceeds this paper. We have prioritized the lessons we take to be most valuable and most generalizable: how voters weigh broad classes of candidate attributes, how those weights depend on the information environment, and how cross-pressured voters navigate tradeoffs. We recognize, though, that theoretical importance is often in the eye of the theorist, and that the many other questions this corpus could address are important too, regardless of where they fall in our own (or others’) “importance hierarchy.” To that end, we have developed tools that will enable other scholars to use and build on our work. We provide the cleaning, reanalysis, and meta-analysis code and the resulting standardized data, subject to restrictions on the redistribution of original study datasets; we package the order-invariant matchup-extraction and estimation machinery as a standalone R package that applies to any standard conjoint design; and we provide an interactive explorer that recomputes any requested win rate on demand from the standardized data (Appendix A describes the last two). A useful first step for the next question in candidate choice may not be a new experiment at all: a scholar with a hunch can ask what the assembled evidence already implies, and where precisely our current knowledge is thin, before fielding anything new.

References

- Abbott, Jared, and Fred DeVeaux. 2024. “Do Working-Class Candidates Activate Class-Based Voting?” *Electoral Studies* 92: 102862.
- Abramson, Scott F, Korhan Koçak, and Asya Magazinnik. 2022. “What do we learn about voter preferences from conjoint experiments?” *American Journal of Political Science* 66 (4): 1008–1020.
- Abramson, Scott F, Korhan Kocak, Asya Magazinnik, and Anton Strezhnev. 2026. “Aggregation, Interpretation, and Estimation of Preferences in Conjoint Experiments.” *SocArXiv* preprint, version 3, available at https://doi.org/10.31235/osf.io/xjre9_v3, last accessed August 16, 2026.
- Abramson, Scott F, and Morgan Gillespie. 2024. “The Effect of Gender on Perceptions of Politicians’ Ideology: Implications for Candidate-Choice Experiments.” Unpublished manuscript.
- Agerberg, Mattias. 2020. “The Lesser Evil? Corruption Voting and the Importance of Clean Alternatives.” *Comparative Political Studies* 53 (2): 253–287.
- Agerberg, Mattias. 2024. “Understanding Preferences for Descriptive Political Representation among Citizens with Immigrant Background.” *Electoral Studies* 90: 102802.
- Albert, Zachary, and Mia Costa. 2025. “Winning At All Costs? How Negative Partisanship Affects Voter Decision-Making.” *Political Behavior* 47 (3): 963–989.
- Alvarez, R Michael. 1997. *Information and elections*. University of Michigan Press.
- Amsalem, Eran, and Alon Zoizner. 2024. “The Causal Effect of Candidate Extremity on Citizens’ Preferences: Evidence from Conjoint Experiments.” *Public Opinion Quarterly* 88 (3): 859–885.
- Arceneaux, Kevin, and Rory Truex. 2023. “Donald Trump and the Lie.” *Perspectives on Politics* 21 (3): 863–879.
- Argote, Pablo, and Giancarlo Visconti. 2023. “Anti-Elite Attitudes and Support for Independent Candidates.” *PLOS ONE* 18 (10): e0292098.
- Arnesen, Sveinung, Dominik Duell, and Mikael Poul Johannesson. 2019. “Do citizens make inferences from political candidate characteristics when aiming for substantive representation?” *Electoral Studies* 57: 46–60.
- Atkeson, Lonna Rae, and Brian T. Hamel. 2020. “Fit for the Job: Candidate Qualifications and Vote Choice in Low Information Elections.” *Political Behavior* 42 (1): 59–82.

- Aviña, Marco M, Taeku Lee, Mashail Malik, Reed Rasband, Marcel F. Roman, and Priyanka Sethy. 2026. "Which Immigrants do Citizens Prefer? A Meta-Reanalysis of 100 Conjoint Experiments." *Science Advances* 12 (22): eadz2271.
- Bakker, Bert N., Gijs Schumacher, and Matthijs Rooduijn. 2021. "The Populist Appeal: Personality and Antiestablishment Communication." *The Journal of Politics* 83 (2): 589–601.
- Bansak, Kirk, Jens Hainmueller, Daniel J. Hopkins, and Teppei Yamamoto. 2018. "The Number of Choice Tasks and Survey Satisficing in Conjoint Experiments." *Political Analysis* 26 (1): 112–119.
- Bansak, Kirk, Jens Hainmueller, Daniel J Hopkins, and Teppei Yamamoto. 2021a. "Beyond the breaking point? Survey satisficing in conjoint experiments." *Political Science Research and Methods* 9 (1): 53–71.
- Bansak, Kirk, Jens Hainmueller, Daniel J. Hopkins, and Teppei Yamamoto. 2021b. "Conjoint Survey Experiments." In *Advances in Experimental Political Science*, ed. James N. Druckman, and Donald P. Green. Cambridge University Press pp. 19–41.
- Bansak, Kirk, Jens Hainmueller, Daniel J Hopkins, and Teppei Yamamoto. 2023. "Using conjoint experiments to analyze election outcomes: The essential role of the average marginal component effect." *Political Analysis* 31 (4): 500–518.
- Bansak, Kirk, and Libby Jenke. 2025. "Odd profiles in conjoint experimental designs: Effects on survey-taking attention and behavior." *Political Analysis* 33 (4): 315–338.
- Barnett, Carolyn, Alexandra Blackman, and Marwa Shalaby. 2026. "Misperceptions of Male Bias Against Female Candidates: An Obstacle to Women's Political Representation?" *Political Behavior* . Online first, published 15 October 2026.
- Berz, Jan, and Michael Jankowski. 2022. "Local Preferences in Candidate Selection: Evidence from a Conjoint Experiment among Party Leaders in Germany." *Party Politics* 28 (6): 1136–1149.
- Bhatia, Jasmine, and Steve L. Monroe. 2024. "Insecurity and Support for Female Leadership in Conflict States: Evidence from Afghanistan." *British Journal of Political Science* 54 (3): 771–792.
- Bischof, Daniel, and Roman Senninger. 2026. "Can Simple Language Affect Voters' Political Knowledge and Their Beliefs About Politicians?" *The Journal of Politics* 88 (4): 1419–1436.
- Black, Gordon S. 1972. "A theory of political ambition: Career choices and the role of structural incentives." *American political science review* 66 (1): 144–159.

- Blackman, Alexandra Domike, and Marlette Jackson. 2021. "Gender Stereotypes, Political Leadership, and Voting Behavior in Tunisia." *Political Behavior* 43 (3): 1037–1066.
- Breitenstein, Sofia. 2019. "Choosing the crook: A conjoint experiment on voting for corrupt politicians." *Research & Politics* 6 (1): 2053168019832230.
- Calvert, Randall L. 1985. "Robustness of the multidimensional voting model: Candidate motivations, uncertainty, and convergence." *American Journal of Political Science* 29 (1): 69–95.
- Cammett, Melani, Christiana Parreira, Dominika Kruszewska-Eduardo, and Sami Atallah. 2022. "Commitment to the "National" in Post-Conflict Countries: Public and Private Security Provision in Lebanon." *Journal of Conflict Resolution* 66 (7–8): 1235–1262.
- Campbell, Rosie, Philip Cowley, Nick Vivyan, and Markus Wagner. 2019a. "Legislator Dissent as a Valence Signal." *British Journal of Political Science* 49 (1): 105–128.
- Campbell, Rosie, Philip Cowley, Nick Vivyan, and Markus Wagner. 2019b. "Why Friends and Neighbors? Explaining the Electoral Appeal of Local Roots." *The Journal of Politics* 81 (3): 937–951.
- Carey, John, Katherine Clayton, Gretchen Helmke, Brendan Nyhan, Mitchell Sanders, and Susan Stokes. 2022. "Who will defend democracy? Evaluating tradeoffs in candidate support among partisan donors and voters." *Journal of Elections, Public Opinion and Parties* 32 (1): 230–245.
- Carlson, Elizabeth. 2015. "Ethnic Voting and Accountability in Africa: A Choice Experiment in Uganda." *World Politics* 67 (2): 353–385.
- Carnes, Nicholas, and Noam Lupu. 2016. "Do Voters Dislike Working-Class Candidates? Voter Biases and the Descriptive Underrepresentation of the Working Class." *American Political Science Review* 110 (4): 832–844.
- Carter, Christopher L. 2021. "The Representational Effects of Communal Property: Evidence from Peru's Indigenous Groups." *Comparative Political Studies* 54 (12): 2191–2225.
- Castanho Silva, Bruno, Fabian Guy Neuner, and Christopher Wratil. 2023. "Populism and Candidate Support in the US: The Effects of "Thin" and "Host" Ideology." *Journal of Experimental Political Science* 10 (3): 438–447.
- Cho, John, Mia Costa, and Yusaku Horiuchi. 2026. "Descriptive or Partisan Representation? Examining Trade-Offs for Asian Americans." *British Journal of Political Science* 56.
- Christensen, Henrik Serup, Theodora Järvi, Mikko Mattila, and Åsa von Schoultz. 2021. "How Voters Choose One Out of Many: A Conjoint Analysis of the Effects of Endorsements on Candidate Choice." *Political Research Exchange* 3 (1): 1892456.

- Claassen, Christopher. 2026. "Voting for Democracy? Testing Political-Cultural Constraints on Democratic Backsliding." *Public Opinion Quarterly*. Advance article, published online 11 September 2026.
- Clayton, Katherine, Yusaku Horiuchi, Aaron R. Kaufman, Gary King, and Mayya Komisarchik. 2026. "Correcting measurement error bias in conjoint survey experiments." *American Journal of Political Science* p. ajps.70076. Published online July 23, 2026.
- Coffé, Hilde, Sparsha Saha, and Ana Catalano Weeks. 2025. "Quotas or parity? How the framing of positive action measures impacts public support and backlash." *Journal of European Public Policy* 32 (10): 2409–2438.
- Coppock, Alexander, and Diana Roxana Galos. 2025. "Toward the generalizability of the party cue effect: a meta-reanalysis of hundreds of experimental contrasts." Conditionally accepted, *American Journal of Political Science*.
- Costa, Mia. 2021. "Ideology, Not Affect: What Americans Want from Political Representation." *American Journal of Political Science* 65 (2): 342–358.
- Cutler, Fred. 2002. "The simplest shortcut of all: Sociodemographic characteristics and electoral choice." *Journal of Politics* 64 (2): 466–490.
- Dahl, Malte, and Jacob Nystrup. 2021. "Confident and cautious candidates: Explaining under-representation of women in danish municipal politics." *European Journal of Political Research* 60 (1): 199–224.
- Dahl, Robert A. 2008. *Polyarchy: Participation and opposition*. Yale university press.
- Dai, Yaoyao, and Alexander Kustov. 2024. "The (In)effectiveness of Populist Rhetoric: A Conjoint Experiment of Campaign Messaging." *Political Science Research and Methods* 12 (4): 849–856.
- Dassonneville, Ruth, and Ian McAllister. 2018. "Gender, political knowledge, and descriptive representation: The impact of long-term socialization." *American Journal of Political Science* 62 (2): 249–265.
- Davis, Otto A, Melvin J Hinich, and Peter C Ordeshook. 1970. "An expository development of a mathematical model of the electoral process." *American political science review* 64 (2): 426–448.
- Dawson, Michael C. 1995. *Behind the mule: Race and class in African-American politics*. Princeton University Press.
- Doherty, David, Conor M. Dowling, and Michael G. Miller. 2019. "Do Local Party Chairs Think Women and Minority Candidates Can Win? Evidence from a Conjoint Experiment." *The Journal of Politics* 81 (4): 1282–1297.

- Downs, Anthony. 1957. "An economic theory of political action in a democracy." *Journal of political economy* 65 (2): 135–150.
- Driscoll, Amanda, and Michael J. Nelson. 2023. "The Costs of Court Curbing: Evidence from the United States." *The Journal of Politics* 85 (2): 609–624.
- Egami, Naoki, and Erin Hartman. 2023. "Elements of external validity: Framework, design, and analysis." *American Political Science Review* 117 (3): 1070–1088.
- Enelow, James, and Melvin J Hinich. 1981. "A new approach to voter uncertainty in the Downsian spatial model." *American journal of political science* 25 (3): 483–493.
- Eshima, Shusei, and Daniel M. Smith. 2022. "Just a Number? Voter Evaluations of Age in Candidate-Choice Experiments." *The Journal of Politics* 84 (3): 1856–1861.
- Ferrari, Diogo, and Brianna A. Smith. 2025. "Status Threat, Partisanship, and Voters' Conservative Shift Toward Right-Wing Candidates." *Journal of Experimental Political Science* 12 (3): 324–338.
- Finseraas, Henning, Oliver Heath, Peter Egge Langsæther, and Kaat Smets. 2026. "How Group Appeals Shape Group–Party Linkages in Two Political Systems." *The Journal of Politics* 88 (3): 1264–1269.
- Finseraas, Henning, and Rino Heim. 2026. "What Drives Voter Bias Against Muslim Politicians? An Experimental Examination." Working paper, Norwegian University of Science and Technology. Draft dated August 8, 2025. Replication data: Harvard Dataverse, doi:10.7910/DVN/A1OKBZ.
- Fowler, Linda L, and Robert D McClure. 1989. *Political ambition: Who decides to run for Congress*. Yale University Press.
- Franchino, Fabio, and Francesco Zucchini. 2015. "Voting in a Multi-Dimensional Space: A Conjoint Analysis Employing Valence and Ideology Attributes of Candidates." *Political Science Research and Methods* 3 (2): 221–241.
- Funck, Amy S, and Katherine T McCabe. 2022. "Partisanship, information, and the conditional effects of scandal on voting decisions." *Political Behavior* 44 (3): 1389–1409.
- Gagnon, Etienne, Fumi Ikeda, and Kenneth Mori McElwain. 2026. "Preferring National Elites or Local Candidates: A Conjoint Analysis of Voter Heuristics." *The Journal of Politics* .
- Galos, Diana Roxana, and Alexander Coppock. 2023. "Gender composition predicts gender bias: A meta-reanalysis of hiring discrimination audit experiments." *Science Advances* 9 (18): eade7979.

- Gatto, Malu A. C., and Marco Radojevic. 2025. "Choosing Women: Party Elites' Preferences in the Candidate Selection Process." *British Journal of Political Science* 55: e36.
- Gay, Claudine, Jennifer Hochschild, and Ariel White. 2016. "Americans' Belief in Linked Fate: Does the Measure Capture the Concept?" *Journal of Race, Ethnicity, and Politics* 1 (1): 117–144.
- Gherghina, Sergiu, and Inga A.-L. Saikkonen. 2024. "Partisanship and Tolerance for Clientelism: Evidence from a Conjoint Experiment in Romania." *Political Studies* 72 (3): 1250–1266.
- Gidengil, Elisabeth, Dietlind Stolle, and Olivier Bergeron-Boutin. 2022a. "COVID-19 and Support for Executive Aggrandizement." *Canadian Journal of Political Science* 55 (2): 342–372.
- Gidengil, Elisabeth, Dietlind Stolle, and Olivier Bergeron-Boutin. 2022b. "The Partisan Nature of Support for Democratic Backsliding: A Comparative Perspective." *European Journal of Political Research* 61 (4): 901–929.
- Gonthier, Frédéric, and Tristan Guerra. 2022. "From the People, Like the People, or For the People? Candidate Appraisal Among the French Yellow Vests." *Political Psychology* 43 (5): 969–989.
- Gothreau, Claire, and Lasse Laustsen. 2026. "Women and Left-Wing Citizens Prefer Women Candidates: Testing Consistency and Psychological Processes Across Twenty Diverse Countries." *Public Opinion Quarterly* 90: 989–1016.
- Goyal, Tanushree. 2020. "Local Female Representation as a Pathway to Power: A Natural Experiment in India." Working paper, May 1, 2020. SSRN 3590118.
- Graham, Matthew H. 2026. "Partisan Expressive Responding: Lessons from Two Decades of Research." *American Political Science Review*.
- Graham, Matthew H, and Milan W Svobik. 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.
- Grahn, Michal, and Theodor Thisell. 2025. "Are Minorities in Politics Held to a Higher Standard? Experimental and Observational Evidence from Candidate Selection." *Journal of European Public Policy* 32 (9): 2282–2306.
- Green, Donald P., Bradley Palmquist, and Eric Schickler. 2004. *Partisan Hearts and Minds: Political Parties and the Social Identities of Voters*. Yale University Press.
- Green, Jon, Brian Schaffner, and Sam Luks. 2023. "Strategic discrimination in the 2020 democratic primary." *Public Opinion Quarterly* 86 (4): 886–898.

- Greene, Kenneth F. 2024. "Money Can't Buy You Love: Partisan Responses to Vote-Buying Offers." *American Journal of Political Science* 68 (2): 644–660.
- Gross, Martin, Christina-Marie Juen, and Christian Stecker. 2026. "The Cordon Dilemma: How Voters Evaluate Cooperation With the Far-Right in Parliament." *Comparative Political Studies*. Online first, published 16 May 2026.
- Guryan, Jonathan, and Kerwin Kofi Charles. 2013. "Taste-based or statistical discrimination: The economics of discrimination returns to its roots." *The Economic Journal* 123 (572): F417–F432.
- Haines, Pavielle, and Seth Masket. 2024. "'You Had Better Mention All of Them:' Race and Gender Effects in Election Loss Narratives." *Political Research Quarterly* 77 (1): 417–431.
- Hainmueller, Jens, Daniel J. Hopkins, and Teppei Yamamoto. 2014. "Causal Inference in Conjoint Analysis: Understanding Multidimensional Choices via Stated Preference Experiments." *Political Analysis* 22 (1): 1–30.
- Hassell, Hans J. G., and Neil Visalvanich. 2024. "Perceptions of Electability: Candidate (and Voter) Ideology, Race, and Gender." *Political Behavior* 46 (4): 2075–2098.
- Heckman, James J. 1998. "Detecting Discrimination." *Journal of Economic Perspectives* 12 (2): 101–116.
- Hjortskov, Morten, and Simon Calmar Andersen. 2024. "Lexicographic Preferences in Candidate Choice: How Party Affiliation Dominates Gender and Race." *British Journal of Political Science* 54 (4): 1404–1423.
- Hobolt, Sara B., and Toni Rodon. 2020. "Cross-Cutting Issues and Electoral Choice: EU Issue Voting in the Aftermath of the Brexit Referendum." *Journal of European Public Policy* 27 (2): 227–245.
- Holman, Mirya R, Jennifer L Merolla, and Elizabeth J Zechmeister. 2016. "Terrorist threat, male stereotypes, and candidate evaluations." *Political Research Quarterly* 69 (1): 134–147.
- Hopkins, Daniel J. 2014. "Do Voting Criteria Differ across the Levels of the U.S. Federal System?" TESS proposal. Unpublished candidate conjoint study.
- Horiuchi, Yusaku, Daniel M. Smith, and Teppei Yamamoto. 2020. "Identifying Voter Preferences for Politicians' Personal Attributes: A Conjoint Experiment in Japan." *Political Science Research and Methods* 8 (1): 75–91.
- Huddy, Leonie, Lilliana Mason, and Lene Aarøe. 2015. "Expressive partisanship: Campaign involvement, political emotion, and partisan identity." *American Political Science Review* 109 (1): 1–17.

- Huddy, Leonie, and Nayda Terkildsen. 1993. "Gender stereotypes and the perception of male and female candidates." *American journal of political science* 37 (1): 119–147.
- Hunter, Andrew. 2024. "The Populist Impulse: Cognitive Reflection, Populist Attitudes and Candidate Preferences." *Electoral Studies* 92: 102868.
- Islam, Md Mujahedul, and Peter John Loewen. 2026. "Ethical Distance and Electoral Accountability: Voters Punish Politicians Whose Ethical Views Diverge from Societal Consensus." *Journal of Elections, Public Opinion and Parties* pp. 1–30.
- Iyengar, Shanto, and Sean J Westwood. 2015. "Fear and loathing across party lines: New evidence on group polarization." *American journal of political science* 59 (3): 690–707.
- Jacobson, Gary C, and Samuel Kernell. 1983. *Strategy and choice in congressional elections*. New Haven, CT: Yale University Press.
- Joo, Minhye, Danielle Casarez Lemi, Jennifer L. Merolla, and Abbylin H. Sellers. 2026. "Only for Some Women: Experimental Evidence on Descriptive Representation and Political Engagement." *Political Behavior* . Online first, 21 April 2026.
- Justesen, Mogens K., Sigrid Koob, and Sina Smid. 2025. "Clientelism and programmatic redistribution: Evidence from a conjoint survey experiment in Brazil." *World Development* 195: 107124.
- Kage, Rieko, Frances McCall Rosenbluth, and Seiki Tanaka. 2019. "What Explains Low Female Political Representation? Evidence from Survey Experiments in Japan." *Politics & Gender* 15 (2): 285–309.
- Kang, Woo Chang, Jill Sheppard, Feodor Snagovsky, and Nicholas Biddle. 2021. "Candidate Sex, Partisanship and Electoral Context in Australia." *Electoral Studies* 70: 102273.
- Karl, Kristyn L, and Timothy J Ryan. 2016. "When are stereotypes about black candidates applied? An experimental test." *Journal of Race, Ethnicity, and Politics* 1 (2): 253–279.
- Karpowitz, Christopher F, Tyson King-Meadows, J Quin Monson, and Jeremy C Pope. 2021. "What leads racially resentful voters to choose black candidates?" *The Journal of Politics* 83 (1): 103–121.
- Kato, Gento, Fan Lu, and Masahisa Endo. 2025. "The Preference-Expectation Gap in Support for Female Candidates: Evidence from Japan." *Public Opinion Quarterly* 89 (1): 217–228.
- Kennedy, Megan, Michael J. Nelson, and Erin Heidt-Forsythe. 2025. "Gender Stereotypes and Candidate Qualifications in Judicial Elections." *Political Research Quarterly* 78 (1): 182–197.

- Kim, Jeong Hyun, and Yesola Kweon. 2024. "Double penalty? How candidate class and gender influence voter evaluations." *Research & Politics* 11 (1): 20531680241226511.
- Kirkland, Patricia A., and Alexander Coppock. 2018. "Candidate Choice Without Party Labels: New Insights from Conjoint Survey Experiments." *Political Behavior* 40 (3): 571–591.
- Klar, Samara. 2013. "The influence of competing identity primes on political preferences." *The Journal of Politics* 75 (4): 1108–1124.
- Koch, Jeffrey W. 2000. "Do citizens apply gender stereotypes to infer candidates' ideological orientations?" *The Journal of Politics* 62 (2): 414–429.
- Koch, Jeffrey W. 2002. "Gender stereotypes and citizens' impressions of House candidates' ideological orientations." *American Journal of Political Science* 46 (2): 453–462.
- Kustov, Alexander, and Yaoyao Dai. 2026. "What is populism good for? An experimental test of mobilization effects." *Research & Politics* 13 (1): 20531680261434109.
- Laterzo, Isabel G. 2024. "Progressive Ideology and Support for Punitive Crime Policy: Evidence from Argentina and Brazil." *Comparative Political Studies* 57 (6): 999–1034.
- Laterzo-Tingley, Isabel, and Leah Christiani. 2024. "The "Tough-on-Crime" Left: Race, Gender, and Elections of Law-and-Order Democrats." *Research & Politics* 11 (3): 20531680241261765.
- Lavezzolo, Sebastián, Luis Ramiro, and Pablo Fernández-Vázquez. 2021. "The Will for Reason: Voter Demand for Experts in Office." *West European Politics* 44 (7): 1506–1531.
- Lawless, Jennifer L, and Richard L Fox. 2010. *It still takes a candidate: Why women don't run for office*. Cambridge University Press.
- Le Foulon, Carmen, and Catherine Reyes-Housholder. 2021. "Candidate Sex, Corruption and Vote Choice." *Electoral Studies* 69: 102270.
- Leeper, Thomas J, and Joshua Robison. 2020. "More important, but for what exactly? The insignificant role of subjective issue importance in vote decisions." *Political Behavior* 42 (1): 239–259.
- Leeper, Thomas J., Sara B. Hobolt, and James Tilley. 2020. "Measuring Subgroup Preferences in Conjoint Experiments." *Political Analysis* 28 (2): 207–221.
- Leininger, Arndt, and Tilko Swalve. 2026. "What's in a name? Name recognition, place-based identities, and candidate success." *The Journal of Politics* .
- Lemi, Danielle Casarez. 2021. "Do voters prefer just any descriptive representative? The case of multiracial candidates." *Perspectives on Politics* 19 (4): 1061–1081.

- Levy, Gabriella. 2022. "Evaluations of violence at the polls: Civilian victimization and support for perpetrators after war." *The Journal of Politics* 84 (2): 783–797.
- Littmann, Brian, Nadia Natasha James Nielsen, and Lasse Laustsen. 2026. "Terrorist threat decreases citizens' relative preferences for women candidates: Insights from a Danish conjoint experiment." *European Journal of Political Research* pp. 1–14.
- Loewen, Peter John, and Ludovic Rheault. 2021. "Voters Punish Politicians with Depression." *British Journal of Political Science* 51 (1): 427–436.
- Lublin, David, and Matthew Wright. 2024. "Diversity Matters: The Election of Asian Americans to U.S. State and Federal Legislatures." *American Political Science Review* 118 (1): 380–400.
- Lupia, Arthur, and Mathew D. McCubbins. 1998. *The Democratic Dilemma: Can Citizens Learn What They Need to Know?* Cambridge: Cambridge University Press.
- Madrid, Raul, Jr., Jennifer L. Merolla, Aldo Yanez Ruiz, and Jean Reith Schroedel. 2022. "The Relevance of Religion for Political Office: Voter Bias Toward Candidates from Different Religious Backgrounds." *Political Behavior* 44 (2): 981–1001.
- Maestas, Cherie D., Sarah Fulton, L. Sandy Maisel, and Walter J. Stone. 2006. "When to Risk It? Institutions, Ambitions, and the Decision to Run for the U.S. House." *American Political Science Review* 100 (2): 195–208.
- Magni, Gabriele, and Andrew Reynolds. 2021. "Voter preferences and the political underrepresentation of minority groups: Lesbian, gay, and transgender candidates in advanced democracies." *The Journal of Politics* 83 (4): 1199–1215.
- Malesky, Edmund, and Paul Schuler. 2020. "Single-Party Incumbency Advantage in Vietnam: A Conjoint Survey Analysis of Public Electoral Support." *Journal of East Asian Studies* 20 (1): 25–52.
- Manento, Cory, and Paul F Testa. 2022. "In party we trust? Voter support for party-backed candidates in primary elections." *Political Behavior* 44 (4): 1633–1656.
- Mansbridge, Jane. 1999. "Should Blacks Represent Blacks and Women Represent Women? A Contingent "Yes". " *The Journal of politics* 61 (3): 628–657.
- Manzano, Sylvia, and Gabriel R Sanchez. 2010. "Take one for the team? Limits of shared ethnicity and candidate preferences." *Political Research Quarterly* 63 (3): 568–580.
- Mares, Isabela, and Giancarlo Visconti. 2020. "Voting for the Lesser Evil: Evidence from a Conjoint Experiment in Romania." *Political Science Research and Methods* 8 (2): 315–328.
- Mares, Isabela, Kenneth Scheve, and Christina Toenshoff. 2026. "Compensation, Beliefs in State Intervention, and Support for the Energy Transition." *Comparative Political Studies* 59 (2): 272–307.

- Martini, Sergio, and Francesco Olmastroni. 2021. "From the lab to the poll: The use of survey experiments in political research." *Italian Political Science Review/Rivista Italiana di Scienza Politica* 51 (2): 231–249.
- McDermott, Monika L. 1997. "Voting cues in low-information elections: Candidate gender as a social information variable in contemporary United States elections." *American journal of political science* 41 (1): 270–283.
- Mercea, Dan, and Felipe G. Santos. 2025. "Policy over Protest: Experimental Evidence on the Drivers of Support for Movement Parties." *Perspectives on Politics* 23 (3): 901–923.
- Mikkelborg, Anna Caroline. 2025. "White Democrats' Growing Support for Black Politicians in the Era of the "Great Awakening"." *American Political Science Review* 119 (4): 1902–1920.
- Miller, David R., and Jeffrey Ziegler. 2024. "Preferential abstention in conjoint experiments." *Research & Politics* 11 (4): 20531680241299329.
- Mo, Cecilia Hyunjung. 2015. "The Consequences of Explicit and Implicit Gender Attitudes and Candidate Quality in the Calculations of Voters." *Political Behavior* 37 (2): 357–395.
- Mummolo, Jonathan, Erik Peterson, and Sean Westwood. 2021. "The limits of partisan loyalty." *Political Behavior* 43 (3): 949–972.
- Nelson, Thomas E, Kira Sanbonmatsu, and Harwood K McClerking. 2007. "Playing a different race card: Examining the limits of elite influence on perceptions of racism." *The Journal of Politics* 69 (2): 416–429.
- Neumark, David. 2012. "Detecting Discrimination in Audit and Correspondence Studies." *Journal of Human Resources* 47 (4): 1128–1157.
- Neuner, Fabian G., and Christopher Wratil. 2022. "The Populist Marketplace: Unpacking the Role of "Thin" and "Thick" Ideology." *Political Behavior* 44 (2): 551–574.
- Noble, Benjamin S., and Taylor N. Carlson. 2025. "CueAnon: What QAnon Signals About Congressional Candidates and What it Costs Them." *Political Behavior* 47 (1): 435–456.
- Nyholt, Niels. 2024. "Why Do Voters Prefer Local Candidates? Evidence from a Danish Conjoint Survey Experiment." *Political Behavior* 46 (4): 2313–2332.
- Ofosu, George Kwaku, Merete Bech Seeberg, and Michael Wahman. 2026. "Messages Matter: How Voter Education Campaigns Affect Citizens' Willingness to Vote for Women." *The Journal of Politics* 88 (1): 145–161.
- Ono, Yoshikuni, and Barry C Burden. 2019. "The contingent effects of candidate sex on voter choice." *Political Behavior* 41 (3): 583–607.

- Ono, Yoshikuni, and Masahiro Yamada. 2020. "Do Voters Prefer Gender Stereotypic Candidates? Evidence from a Conjoint Survey Experiment in Japan." *Political Science Research and Methods* 8 (3): 477–492.
- Ornstein, Joseph T., Amanda J. Heideman, Bryant J. Moy, and Kaylyn Jackson Schiff. 2025. "Hometown Advantage: Voter Preferences for Community Embeddedness in Local Contests." *Journal of Experimental Political Science* 12 (3): 267–281.
- Page, Matthew J., Joanne E. McKenzie, Patrick M. Bossuyt, Isabelle Boutron, Tammy C. Hoffmann, Cynthia D. Mulrow, Larissa Shamseer, Jennifer M. Tetzlaff, Elie A. Akl, Sue E. Brennan et al. 2021. "The PRISMA 2020 statement: an updated guideline for reporting systematic reviews." *Systematic Reviews* 10 (1): 89.
- Panel, Sophie, Emilien Paulis, Jean-Benoit Pilet, Sebastien Rojon, and Davide Vittori. 2024. "The Lure of Technocrats: A Conjoint Experiment on Preferences for Technocratic Ministers in Six European Countries." *Political Behavior* 46 (3): 1961–1984.
- Pedersen, Helene Helboe. 2026. "Legislator or Representative? Politicians' Tasks According to Voters." *British Journal of Political Science* 56: e4.
- Peterson, Erik. 2017. "The role of the information environment in partisan voting." *The Journal of Politics* 79 (4): 1191–1204.
- Peterson, Erik, and Gabor Simonovits. 2018. "The Electoral Consequences of Issue Frames." *The Journal of Politics* 80 (4): 1283–1296.
- Phelps, Edmund S. 1972. "The Statistical Theory of Racism and Sexism." *American Economic Review* 62 (4): 659–661.
- Phillips, Anne. 1995. *The politics of presence*. Clarendon Press.
- Pitkin, Hanna Fenichel. 1967. *The Concept of Representation*. 1 ed. University of California Press.
- Popkin, Samuel L. 1991. *The reasoning voter: Communication and persuasion in presidential campaigns*. University of Chicago Press.
- Portmann, Lea. 2022. "What Makes a Successful Candidate? Political Experience and Low-Information Cues in Elections." *The Journal of Politics* 84 (4): 2049–2063.
- Przeworski, Adam, Susan C Stokes, and Bernard Manin. 1999. *Democracy, accountability, and representation*. Vol. 2 Cambridge University Press.
- Rahn, Wendy M. 1993. "The role of partisan stereotypes in information processing about political candidates." *American Journal of Political Science* 37 (2): 472–496.
- Rains, Emily, and Erik Wibbels. 2023. "Informal Work, Risk, and Clientelism: Evidence from 223 Slums across India." *British Journal of Political Science* 53 (1): 1–24.

- Reher, Stefanie. 2025. "Voting for Disabled Candidates." *The Journal of Politics* 87 (2): 790–794.
- Rehmert, Jochen. 2022. "Party Elites' Preferences in Candidates: Evidence from a Conjoint Experiment." *Political Behavior* 44 (3): 1149–1173.
- Robinson, Thomas S. 2023. "When Do Voters Respond to Campaign Finance Disclosure? Evidence from Multiple Election Types." *Political Behavior* 45 (4): 1309–1332.
- Roccas, Sonia, and Marilynn B Brewer. 2002. "Social identity complexity." *Personality and social psychology review* 6 (2): 88–106.
- Rohde, David W. 1979. "Risk-bearing and progressive ambition: The case of members of the United States House of Representatives." *American Journal of Political Science* 23 (1): 1–26.
- Saha, Sparsha, and Ana Catalano Weeks. 2022. "Ambitious women: Gender and voter perceptions of candidate ambition." *Political Behavior* 44 (2): 779–805.
- Saikkonen, Inga A.-L., and Henrik Serup Christensen. 2023. "Guardians of Democracy or Passive Bystanders? A Conjoint Experiment on Elite Transgressions of Democratic Norms." *Political Research Quarterly* 76 (1): 127–142.
- Sanbonmatsu, Kira. 2002. "Gender stereotypes and vote choice." *American Journal of Political Science* 46 (1): 20–34.
- Sances, Michael W. 2018. "Ideology and vote choice in US mayoral elections: Evidence from Facebook surveys." *Political Behavior* 40 (3): 737–762.
- Sarkar, Radha, and Alexander Coppock. 2026. "The effects of religious messages and endorsements on political attitudes: A meta-reanalysis." *American Political Science Review* pp. 1–22.
- Şaşmaz, Aytuğ. 2026. "Organizational Cohesion and Unequal Political Selection: Evidence from Tunisia's Secular–Islamist Competition." *Perspectives on Politics* pp. 1–23. Online first, published 29 April 2026.
- Schildkraut, Deborah J. 2013. "Which birds of a feather flock together? Assessing attitudes about descriptive representation among Latinos and Asian Americans." *American Politics Research* 41 (4): 699–729.
- Schroeder, Emma, and Alejandra Campos. 2026. "The Influence of Candidate Traits on Vice Presidential Nominee Preference." *Politics, Groups, and Identities* 14 (2): 333–356.
- Schuler, Paul. 2024. "Gender and Clientelism: Do Expectations of Patronage Penalize Women Candidates in Legislative Elections?" *Comparative Political Studies* 57 (14): 2348–2375.

- Schwarz, Susanne, and Alexander Coppock. 2022. “What Have We Learned About Gender From Candidate Choice Experiments? A Meta-analysis of 67 Factorial Survey Experiments.” *Journal of Politics* 84 (2): 655–668.
- Segovia, Carolina. 2021. “Decidiendo por quién votar. Evidencia experimental del efecto de las emociones en el voto.” *Colombia Internacional* (107): 3–28.
- Senninger, Roman, and Daniel Bischof. 2023. “Do voters want domestic politicians to scrutinize the European Union?” *Political Science Research and Methods* 11 (2): 410–418.
- Sevi, Semra. 2026. “Who Makes a Good Candidate? Evidence from Voters and Politicians.” Working paper, University of Toronto. Prepared for the 2026 EGEN Conference.
- Sevi, Semra, Charles Crabtree, and André Blais. 2024. “Do Voters Punish Women Politicians More?” *American Politics Research* 52 (6): 701–708.
- Sevi, Semra, Riya Cherian, Jérémie Drouin, Aidee Nunez-Colin, Bradley Wood-MacLean, and Rebecca Zhang. 2025. “Immigration, Abortion, and the Limits of Policy Voting: Evidence from a Candidate Choice Conjoint.” Working paper.
- Shepsle, Kenneth A. 1972. “The strategy of ambiguity: Uncertainty and electoral competition.” *American Political science review* 66 (2): 555–568.
- Short, Nicholas, Sophie Hill, and Jacob Brown. 2026. “What Is Ideological Capture and How Do We Measure It? Using Antitrust Reform to Understand Expert–Public Cleavages.” *Perspectives on Politics* 24 (2): 680–700.
- Sigelman, Carol K, Lee Sigelman, Barbara J Walkosz, and Michael Nitz. 1995. “Black candidates, white voters: Understanding racial bias in political perceptions.” *American Journal of Political Science* 39 (1): 243–265.
- Simon, Elizabeth, and Stuart J Turnbull-Dugarte. 2025. “Education as identity? A meta-analysis of education-based in-group preferences in candidate choice experiments.” *The Journal of Politics* 87 (2): 807–811.
- Singh, Shane P. 2022. “Does Compulsory Voting Affect How Voters Choose? A Test Using a Combined Conjoint and Regression Discontinuity Analysis.” *Comparative Political Studies* 55 (12): 2119–2143.
- Slough, Tara, and Scott A Tyson. 2024. *External Validity and Evidence Accumulation*. Cambridge University Press.
- Smith, Amy Erica, and Taylor C. Boas. 2024. “Religion, Sexuality Politics, and the Transformation of Latin American Electorates.” *British Journal of Political Science* 54 (3): 816–835.

- Stanley, Ben. 2026. "Democratic Hypocrisy in Practice: A Panel Study of Revealed Preferences for Liberal Democracy in Poland." *East European Politics* 42 (3): 451–470.
- Stokes-Brown, Atiya Kai, and Kathleen Dolan. 2010. "Race, gender, and symbolic representation: African American female candidates as mobilizing agents." *Journal of Elections, Public Opinion and Parties* 20 (4): 473–494.
- Tajfel, Henri, and John C Turner. 2004. "The social identity theory of intergroup behavior." In *Political psychology*. Psychology Press pp. 276–293.
- Teele, Dawn Langan, Joshua Kalla, and Frances Rosenbluth. 2018. "The Ties That Double Bind: Social Roles and Women's Underrepresentation in Politics." *American Political Science Review* 112 (3): 525–541.
- Thomander, Sierra Davis. 2024. "Who Wins Where? Gender Stereotypes of Qualifications and Offices in Entry-Level Elections." Working paper.
URL: <https://sierradavisthomander.com/research>
- Thomas, Anjali, Sayan Banerjee, Charles Hankla, and Arindam Banerjee. 2024. "Ethnonationalist Gender Norms: How Parties Shape Voter Attitudes toward Female Candidates in India." *American Journal of Political Science* 68 (3): 1057–1074.
- Tirado Castro, Alejandro, and Susan Banducci. 2026. "Not just women for women: How gendered affinities affect candidate support." *Research & Politics* 13 (2): 20531680261439753.
- Treger, Claretta, Thomas Galipeau, Thomas Bergeron, Sarah Lachance, Natasha Goel, Md Mujahedul Islam, Blake Lee-Whiting, Beatrice Magistro, and Peter John Loewen. 2026. "Party or Policy? The Role of Policy Partisanship in Voter Decision-Making." *Political Behavior* .
- van Oosten, Sanne, Liza Mügge, and Daphne van der Pas. 2024. "Race/Ethnicity in Candidate Experiments: a Meta-Analysis and the Case for Shared Identification." *Acta Politica* 59 (1): 19–41.
- Vecchiato, Alessandro, and Kevin Munger. 2025. "Introducing the Visual Conjoint, with an Application to Candidate Evaluation on Social Media." *Journal of Experimental Political Science* 12 (1): 57–71.
- Ventura, Tiago, Sandra Ley, and Francisco Cantú. 2024. "Voting for Law and Order: Evidence from a Survey Experiment in Mexico." *Comparative Political Studies* 57 (4): 551–583.
- Vera, Sofia. 2025. "Cleaning up Politics: Anti-Corruption Appeals in Electoral Campaigns." *British Journal of Political Science* 55: e108.
- Vevea, Jack L., and Larry V. Hedges. 1995. "A General Linear Model for Estimating Effect Size in the Presence of Publication Bias." *Psychometrika* 60 (3): 419–435.

- Visconti, Giancarlo. 2022. "After the Flood: Disasters, Ideological Voting and Electoral Choices in Chile." *Political Behavior* 44 (4): 1985–2004.
- Vittori, Davide, Sébastien Rojon, and Jean-Benoît Pilet. 2026. "Citizens' Views of Decision-Makers' Roles: A Conjoint Experiment in Fifteen Countries." *World Politics* 78 (1): 132–167.
- Wakefield, Derek. 2025. "It's the Economy: The Effect of Economic Policy Appeals on Latino Independents." *Political Behavior* 47 (4): 1571–1594.
- Walder, Maxime, Stefanie Bailer, and Nathalie Giger. 2025. "Does It Matter What You Do (Or Only Who You Are)? On the Effects of Parliamentarians' Behavior on Vote Choice." *Legislative Studies Quarterly* 50 (4): e70044.
- Wallace, Sophia J. 2014. "Examining Latino support for descriptive representation: The role of identity and discrimination." *Social Science Quarterly* 95 (2): 311–327.
- Ward, Peter, and Steven Denney. 2026. "Partisan voters in party systems with ephemeral parties: Evidence from South Korea." *Party Politics* 32 (4): 611–625.
- Weaver, Julie Anne. 2021. "Electoral Dis-Connection: The Limits of Reelection in Contexts of Weak Accountability." *The Journal of Politics* 83 (4): 1462–1477.
- Weissman, Anna. 2025. "Descriptive Representation in an Era of Polarization." *The Journal of Politics* 87 (3): 872–888.
- White, Ismail K, and Chryl N Laird. 2020. "Steadfast democrats: How social forces shape Black political behavior." In *Steadfast Democrats*. Princeton University Press.
- Wood, Abby K. 2023. "Voters Use Campaign Finance Transparency and Compliance Information." *Political Behavior* 45 (4): 1553–1579.
- Wu, Jennifer D., and Gregory A. Huber. 2025. "How and When Candidate Race Affects Inferences About Ideology and Group Favoritism." *Political Science Research and Methods* 13 (3): 527–544.
- Wunsch, Natasha, Marc S. Jacob, and Laurenz Derksen. 2025. "The Demand Side of Democratic Backsliding: How Divergent Understandings of Democracy Shape Political Choice." *British Journal of Political Science* 55: e39.
- Wüest, Reto, and Jonas Pontusson. 2022. "Voter Preferences as a Source of Descriptive (Mis)Representation by Social Class." *European Journal of Political Research* 61 (2): 398–419.