Partisanship in the Marriage Market

Jaemin Woo[*] Camilla Adams[†]

September 2026

Latest version

Abstract

We study how partisanship shapes who partners with whom and who remains single. In the United States, the share of adults aged 25–45 who are married or cohabiting has fallen, and the couples that do form are strongly sorted by politics. We combine a discrete choice experiment with an equilibrium matching model to separate whom people prefer, who is available, and whom they meet. The experiment measures partner preferences. When comparing otherwise identical profiles, respondents choose a same-party partner an estimated 81 percent of the time. By contrast, they choose a partner with the same college-degree status 54 percent of the time on average. We then use these preferences and regional population data to fit the matching model to observed couple patterns and singlehood rates. The estimated model explains party sorting primarily through partisan aversion. Education sorting is explained primarily by meeting opportunities segmented by education, since the estimated preferences over a partner's education cannot account for the observed sorting on their own. In a counterfactual, returning partisan aversion to its mid-1990s level raises the married-or-cohabiting share of adults aged 25–45 by 1.51 percentage points, equal to 29.8 percent of the measured 5.08-point decline from 1995 to 2025.

[*]Brown University, Department of Economics. Email: jaemin_woo@brown.edu
[†]Brown University, Department of Economics. Email: camilla_adams@brown.edu

1 Introduction

In the United States, the share of adults aged 25–45 who are married or cohabiting has fallen, and the couples that do form are strongly sorted by politics. The married-or-cohabiting share fell from 66.75 percent in 1995 to 61.67 percent in 2025 (Flood et al., 2023). Nearly four in ten adults in this age range now live without a spouse or cohabiting partner (Ruggles et al., 2024). About three-quarters of married or cohabiting Americans share their partner's political affiliation, compared with 42 percent under random pairing.[1] As dislike and distrust of the opposing party have grown (Iyengar et al., 2019; Finkel et al., 2020), how has partisanship shaped who partners with whom and who remains single?

Observed sorting alone cannot tell us whether partisanship changes only who partners with whom or also how many partnerships form. People who prefer a co-partisan may find one, leaving the total number of partnerships unchanged. Differences in who is available or whom people meet can produce similar sorting: Democratic and Republican women and men are present in unequal numbers, and potential partners meet in socially segmented settings. Observed couples therefore confound partner preferences, population composition, and meeting opportunities (Currarini et al., 2009; Goldman et al., 2026). Separating these forces lets us predict how both sorting and partnership formation would change if partisan aversion and party composition returned to their earlier levels.

We separate the roles of preferences, population composition, and meeting opportunities by combining a discrete choice experiment with an equilibrium matching model. The experiment measures partner preferences independently of observed couples. Population data measure the numbers of potential partners by region, sex, party, and education. Observed couple patterns and singlehood rates then allow us to estimate the model, including how preferences from the experiment translate into match surplus and how meeting opportunities are segmented. The estimated model explains party sorting primarily through partisan aversion and education sorting primarily through differences in meeting opportunities. In the model, partisan aversion also reduces partnership formation, while historical changes in party composition have little aggregate effect.

The experiment reveals a much stronger preference for matching on party than on education. Between otherwise identical hypothetical partners, the predicted probability of choosing a same-party partner over a cross-party partner is 81 percent. The corresponding probability for matching on college-degree status is 54 percent. Party alignment matters about as much as agreement over having children. These estimates come from 423 Democrats and Republicans aged 25–45, each of whom completes 12 choices between hypothetical long-term partners. Party varies independently of displayed education, income, attractiveness, religiosity, and family plans. Respondents are instructed to treat unlisted characteristics as equal.

The preference for a same-party partner persists when the cross-party partner earns substantially

[1] Authors' calculation using 935 respondent–partner pairs from the 2024 American National Election Studies (ANES) and the main respondent's pre-election weight. Affiliations use each partner's own report. Leaners are classified with their preferred party and non-leaning Independents separately. The random-pairing benchmark uses these observed party shares.

more and when stated issue positions are held fixed. Even when an otherwise comparable cross-party profile earns twice as much, the same-party profile has a predicted choice share of 70 percent. In a separate module, we vary party labels independently of four issue positions. We use respondents' beliefs about each party's positions to combine the label and issue penalties. The residual label accounts for 27 percent of this combined penalty among Democratic respondents and 67 percent among Republican respondents. It captures what the label conveys beyond the displayed issues, which may include other beliefs as well as dislike of the opposing party. Respondents with stronger measured party preferences also tend to have more politically similar actual partners.

The matching model combines preferences on both sides of the market with partner availability and meeting opportunities to determine who partners with whom and who remains single. We extend the transferable-utility framework of Choo and Siow (2006) and Galichon and Salanié (2022) to allow two meeting pools within each of four census regions. Each potential match combines the woman's and man's experimental valuations of party and education. A common estimated factor converts those valuations into match surplus while preserving their relative magnitudes. We estimate this factor and the separation between meeting pools from how often women and men of each party and education group partner with one another. Observed singlehood rates determine each group's baseline value of partnership. Adults can also choose an opposite-sex partnership with someone outside the modeled party or age groups.[2]

The estimated meeting pools are divided almost entirely by education, even when pool assignment can also vary by party. Across the four regions, the college-centered pool is 79 percent college graduates and 21 percent nongraduates. The other pool is almost entirely nongraduates. Within each education group, pool assignment is nearly identical for Democrats and Republicans. Strong party preferences account for party sorting, while the estimated meeting structure explains the education sorting that weak education preferences alone cannot explain.

Returning partisan aversion and party composition to their mid-1990s benchmarks increases cross-party matching, with almost all of the predicted partnership response coming from lower aversion. The historical benchmark pools ANES surveys from 1992–2000, centered on 1996. To construct the aversion counterfactual, we reduce each type's cross-party penalty by 40 percent, in proportion to the reduction needed to return the ANES in-party minus out-party feeling-thermometer gap to its earlier level. We also apply historical changes in party composition within sex-by-education groups, holding meeting opportunities fixed.

The aversion counterfactual raises the married-or-cohabiting share of adults aged 25–45 by 1.51 percentage points. That gain equals 29.8 percent of the measured 5.08-point decline from 1995 to 2025. This comparison holds today's party composition, meeting opportunities, and other determinants fixed; it does not reconstruct the full historical change.

Equalizing the numbers of college-educated women and men lowers singlehood in a separate model counterfactual. The exercise shifts the same proportion of non-college Democratic and Re-

[2]The model covers Democrats and Republicans aged 25–45. "Marriage market" refers broadly to marriage and cohabitation. Existing same-sex partnerships and singles seeking only same-sex partners remain fixed in the counterfactuals. Outside-partnership payoffs are held fixed while the number choosing those partnerships can respond.

publican men into college-educated types within each region until college-educated women and men are equal in number. Party affiliations and regional party totals remain unchanged, as do the total population and each type's estimated preferences and meeting opportunities. This education imbalance therefore matters for aggregate partnership formation, even though the historical change in party composition has little effect.

External evidence supports the findings of strong party preferences and education-based meeting segmentation. The ranking of partner attributes in the experiment resembles that in the American Perspectives Survey, where family plans, religion, and politics rank above education (Cox et al., 2020). The model also closely matches statewide party sorting in held-out North Carolina voter records. In our analysis of the American Community Survey, education sorting is stronger in states and commuting zones where adults with and without a bachelor's degree are more residentially separated, even after accounting for local education composition. Workplaces are also substantially more segregated by education than by party (Dillon et al., 2025; Frake et al., 2026).

We contribute to the economics of assortative mating by tracing party and education sorting to different sources. The implications of assortative mating for inequality and mobility depend in part on why similar people partner with one another (Becker, 1973; Mare, 1991; Schwartz and Mare, 2005; Fernández et al., 2005; Greenwood et al., 2014; Eika et al., 2019; Chiappori et al., 2025; Goldman et al., 2026). Strong preferences and segmented meeting opportunities can generate similar sorting but different responses to changes in the marriage market (Currarini et al., 2009). We combine the independent measurement of preferences used in experimental and online-dating studies (Fisman et al., 2006, 2008; Hitsch et al., 2010; Belot and Francesconi, 2013; Low, 2024) with an equilibrium analysis of preferences and meeting opportunities (Ciscato, 2025; Goldman et al., 2026). The contrast between party and education helps distinguish these forces: our model must explain strong sorting on both dimensions despite much weaker experimental preferences for educational similarity.

The analysis extends research on political homophily to the partnerships forgone because of partisan aversion. Existing work documents political similarity between spouses and examines whether it arises through partner choice or convergence within relationships (Alford et al., 2011; Iyengar et al., 2018; Hersh and Ghitza, 2018; Hudde and Grunow, 2025). Online-dating studies and experiments show that a potential partner's stated party affiliation affects romantic evaluation and partner choice (Huber and Malhotra, 2017; Nicholson et al., 2016; Easton and Holbein, 2021; Sleiman et al., 2025; Taflinger and Hudde, 2026). Such preferences can redirect partnerships toward co-partisans without reducing the number of partnerships. Our equilibrium analysis quantifies both the shift toward same-party partnerships and the effect on singlehood.

Experimentally measured preferences allow us to separate partisan aversion from other determinants of match surplus in an equilibrium matching model. Standard transferable-utility models recover joint surplus from observed pairings (Choo and Siow, 2006; Galichon and Salanié, 2022; Chiappori et al., 2017), combining the contributions of observed and unobserved partner characteristics under assumptions on unobserved heterogeneity (Gualdani and Sinha, 2023). Hitsch et al. (2010) bring preferences estimated from online dating behavior into a matching model, while Andrew and Adams (2025) combine hypothetical choice data with a structural model of family decisions. With partner preferences measured experimentally and partner availability observed in population data, our model uses couple patterns to estimate meeting segmentation and quantify how each force affects partnership formation.

The analysis proceeds from the matching framework and experimental evidence to estimation and counterfactual outcomes. Section 2 develops the matching framework. Sections 3 and 4 present the experiment and preference estimates. Section 5 maps those estimates into match surplus. Section 6 describes the external data, and Section 7 presents estimation and validation. Section 8 reports the counterfactual analysis. Section 9 concludes.

2 Matching Framework

We extend the transferable-utility matching framework of Choo and Siow (2006) and Galichon and Salanié (2022) in two ways. First, we use a discrete choice experiment to identify preferences over a potential partner's party and education, components that observed couple counts and singlehood rates do not separately identify. Second, we divide each regional market into two education-centered meeting pools, allowing observed sorting to reflect both preferences and differences in whom people meet.

In our model, women and men are grouped into types defined by party and education, and each person has preferences over partner types. A match generates a joint surplus, which transferable utility allows the partners to divide between them. In equilibrium, transfers adjust until each person prefers their current outcome to the alternatives and no woman–man pair could both gain by leaving their current outcomes and forming a partnership together. The equilibrium determines who partners with whom and who remains single.[3]

A standard application recovers the joint surplus of each type pairing from observed couple counts and singlehood rates. Joint surplus is the value a match creates for both partners together, relative to both remaining single. When a pairing creates high surplus, more couples of that pairing form and fewer people of the two types stay single. Comparing how often a pairing occurs with how often its types stay single therefore reveals the surplus it creates. However, this comparison identifies the total surplus for each pairing, not the separate contributions of party and education. The standard application also treats the market as one common pool, in which every woman can in principle partner with every man.

We identify preferences over a potential partner's party and education in a discrete choice experiment. In this experiment, respondents choose between hypothetical partner profiles that vary in party, education, income, attractiveness, religiosity, and family plans. The estimates isolate preferences for party and education while holding the other displayed traits fixed. Section 3 describes the

[3]This is a static model that represents a cross-section of partnership outcomes rather than relationship formation and dissolution over time. It models different-sex partnerships and holds same-sex matching fixed.

experiment, and Section 5 maps the preference estimates into the party and education components of joint surplus.

We model opportunities to meet, which we call contact, through two meeting pools within each market. A meeting pool is a set of people who can potentially partner with one another. A market in our application is a census region, the level at which we observe how many people of each type there are and how many couples they form. Schools, workplaces, and neighborhoods are often segregated along socioeconomic lines, including education and class, so people may disproportionately meet partners who resemble them (Kalmijn, 1998; Goldman et al., 2026). A model with one common meeting pool would attribute the resulting sorting to preferences. Our meeting pools allow the model to attribute part of the resulting sorting to contact.

We study Democratic and Republican adults aged 25–45 and model woman–man partnerships. A *modeled partnership* pairs a Democratic or Republican woman in this age range with a Democratic or Republican man in the same age range. An *outside partnership* is with an Independent or someone outside the age range. A person may remain single, form a modeled partnership, or choose an outside partnership.

Same-sex matching is held fixed in the counterfactuals. Singles open to both women and men participate in the modeled woman–man market, where the counterfactuals reallocate people among modeled partnerships, outside partnerships, and singlehood. We report outcomes for all Democratic and Republican adults aged 25–45.

Separating preferences from contact requires information beyond observed matches. A type pairing may be common because the two types prefer one another or because they are more likely to meet. The experiment measures preferences independently of observed matches. Combining those estimates with observed couple counts and singlehood rates allows the calibration to distinguish preference-driven sorting from contact-driven sorting. Section 7 presents the calibration.

2.1 Markets, Types, and Meeting Pools

We treat each of the four census regions, indexed by r, as a separate matching market containing a continuum of women and men aged 25–45.[4] Within each gender, a type is defined by party and education:

$$i,\ j \in \{\mathrm{D}, \mathrm{R}\} \times \{\mathrm{BA}, \mathrm{NoBA}\}.$$

Here, D denotes Democratic affiliation and R Republican affiliation. BA denotes a bachelor's degree or more, and NoBA denotes less than a bachelor's degree. We use i for women's types and j for men's types. For any type t, let $P(t) \in \{\mathrm{D}, \mathrm{R}\}$ and $E(t) \in \{\mathrm{BA}, \mathrm{NoBA}\}$ denote its party and education components.

For sex $g \in \{f, m\}$, type t, and region r, let N^r_{gt} denote the total mass of Democratic and Republican adults aged 25–45. This mass comprises singles who seek only same-sex partners ($L^{\mathrm{SS},r}_{gt}$), singles open to an opposite-sex partner (L^r_{gt}), adults in same-sex partnerships ($C^{\mathrm{SS},r}_{gt}$), adults in

[4] The four regions are the Northeast, Midwest, South, and West. We treat region as fixed within the model and do not model cross-region migration.

modeled partnerships (C^r_{gt}), and adults in outside partnerships (O^r_{gt}). Accordingly,

$$N^r_{gt} = L^{\mathrm{SS},r}_{gt} + L^r_{gt} + C^{\mathrm{SS},r}_{gt} + C^r_{gt} + O^r_{gt}. \tag{1}$$

The model determines partnership choices for the following mass:

$$N^{\mathrm{eq},r}_{gt} = L^r_{gt} + C^r_{gt} + O^r_{gt}. \tag{2}$$

The matching equilibrium reallocates these adults among modeled partnerships, outside partnerships, and singlehood; $L^{\mathrm{SS},r}_{gt}$ and $C^{\mathrm{SS},r}_{gt}$ remain fixed. Each modeled woman–man couple contributes one woman to the female partnered total and one man to the male partnered total. Let $n^r_i = N^{\mathrm{eq},r}_{fi}$ and $m^r_j = N^{\mathrm{eq},r}_{mj}$ denote the masses of type-i women and type-j men included in the equilibrium in region r.

The data sources have distinct roles. The experiment measures preferences for party and education. The other profile attributes enter the choice model as controls but do not define types in the matching model. The CES provides regional counts of single and partnered people by type. The ACS provides partners' age and sex. HCMST records partners' party and education and provides the couple distributions used to estimate sorting. ANES observes how people identify when their partner reports them as undecided, Independent, or other. We use that comparison to estimate how many HCMST partners in this group are Democrats, Independents, or Republicans.

We represent contact with two unobserved meeting pools within each regional market. Each adult included in the equilibrium is assigned to one pool and can form a modeled partnership only with someone assigned to the same pool. Pool assignment therefore determines which Democratic and Republican partner types are available to each person. A given type pairing creates the same joint surplus in either pool, so the pools only affect partner availability. We solve the matching equilibrium separately within each pool and add the pool-specific outcomes to obtain regional counts of modeled couples, outside partnerships, and single people.

Baseline pool assignment is probabilistic and depends only on education. Section 7.3 discusses this restriction and how we assess it. Conditional on education, assignment is independent of party and individual partner preferences. One pool is centered on BA types and the other on NoBA types. Let $\pi \in [1/2, 1]$ denote the share of BA types assigned to the BA-centered pool. To keep the assignment rule symmetric, we assign the same share of NoBA types to the NoBA-centered pool. The remaining share $1-\pi$ of each education group enters the other pool. The same π applies across regions and sexes. At $\pi = 1/2$, every type is divided equally between the pools, so both pools have the composition of the full regional market and the model reduces to the standard one-pool case. At $\pi = 1$, BA and NoBA types occupy separate pools.

2.2 Equilibrium Within a Meeting Pool

Fix region r and meeting pool $h \in \{1, 2\}$. Each person selects among the three outcomes defined above: singlehood, a modeled partnership within the same pool, or an outside partnership.

People of the same observed type may value the same partner type differently. For example, two college-educated Democratic men may derive different idiosyncratic utility from a college-educated Democratic woman. We represent this within-type heterogeneity with idiosyncratic tastes. Let $\varepsilon_{j\ell}$ denote woman ℓ's taste for a type-j partner, and let $\eta_{i\ell'}$ denote man ℓ''s taste for a type-i partner. We impose separability: each person has a distinct taste for each partner type, rather than for each potential partner, and the two partners' tastes enter match surplus additively (Galichon and Salanié, 2022).

A type-i woman ℓ receives $\mathcal{U}_{j\ell}^{rh}$ from a type-j partner, while a type-j man ℓ' receives $\mathcal{V}_{i\ell'}^{rh}$ from a type-i partner:

$$\mathcal{U}_{j\ell}^{rh} = U_{ij}^{rh} + \varepsilon_{j\ell}, \tag{3}$$

$$\mathcal{V}_{i\ell'}^{rh} = V_{ij}^{rh} + \eta_{i\ell'}. \tag{4}$$

The systematic payoff U_{ij}^{rh} is common to type-i women evaluating type-j men, while V_{ij}^{rh} is common to type-j men evaluating type-i women. Both are determined in equilibrium. Each woman also draws $\varepsilon_{0\ell}$ for singlehood and $\varepsilon_{o\ell}$ for an outside partnership; each man draws the corresponding $\eta_{0\ell'}$ and $\eta_{o\ell'}$, where o indexes the outside alternative. All taste draws are independent and identically distributed Type-I extreme value with unit scale:[5]

$$\varepsilon_{j\ell},\, \varepsilon_{0\ell},\, \varepsilon_{o\ell},\, \eta_{i\ell'},\, \eta_{0\ell'},\, \eta_{o\ell'} \overset{\text{iid}}{\sim} \text{Gumbel}(0,1).$$

An equilibrium is stable if each person selects the highest-payoff available alternative and no woman–man pair could both gain by leaving their current outcomes and forming a modeled partnership with each other. Under transferable utility, a transfer changes only how the joint surplus is divided. We absorb the transfer into U_{ij}^{rh} and V_{ij}^{rh}. We set $U_{i0}^{rh} = 0$ and $V_{0j}^{rh} = 0$, so a single woman's payoff is $\varepsilon_{0\ell}$ and a single man's payoff is $\eta_{0\ell'}$. Let ζ_{gt}^{r} denote the systematic payoff from an outside partnership for sex g and type t in region r. This payoff is common across meeting pools. An outside partnership pays $\zeta_{fi}^{r} + \varepsilon_{o\ell}$ to a type-i woman and $\zeta_{mj}^{r} + \eta_{o\ell'}$ to a type-j man. Because the partners on the other side are outside the model, their supply is not subject to a market-clearing condition.

The pools define who is available to whom. Let Φ_{ij}^{r} denote the systematic joint surplus created by a type-i woman and a type-j man in region r. For a given type pairing, this surplus is the same in either pool:

$$U_{ij}^{rh} + V_{ij}^{rh} = \Phi_{ij}^{r}. \tag{5}$$

Pool composition can nevertheless change how that surplus is divided. Thus U_{ij}^{rh} and V_{ij}^{rh} may differ across pools even though their sum is fixed. Appendix A.1 derives these type-level payoffs

[5] The i.i.d. Type-I extreme-value specification yields multinomial-logit choice shares and therefore imposes independence of irrelevant alternatives across partner types, singlehood, and the outside partnership. IIA can be particularly restrictive when the choice set contains many close substitutes. This concern is less acute here because the modeled partner alternatives are broad categories defined only by party and education (Train, 2009; Galichon and Salanié, 2022).

from individual stability.

Let μ_{ij}^{rh} denote the mass of (i, j) couples in the pool, μ_{i0}^{rh} the mass of single type-i women, and μ_{0j}^{rh} the mass of single type-j men. Let μ_{io}^{rh} and μ_{oj}^{rh} denote the masses of type-i women and type-j men who form outside partnerships. Under the extreme-value assumption, these payoff comparisons imply logit choice shares (McFadden, 1974; Choo and Siow, 2006). For each type pairing, the mass of couples relative to the mass of single women or single men is

$$\frac{\mu_{ij}^{rh}}{\mu_{i0}^{rh}} = \exp\Big(U_{ij}^{rh}\Big), \qquad \frac{\mu_{ij}^{rh}}{\mu_{0j}^{rh}} = \exp\Big(V_{ij}^{rh}\Big). \tag{6}$$

The first ratio compares (i, j) couples with single type-i women; the second compares the same couples with single type-j men. A higher systematic payoff therefore makes that pairing more common relative to singlehood. The same logit choice rule determines the outside-partnership shares:

$$\frac{\mu_{io}^{rh}}{\mu_{i0}^{rh}} = \exp(\zeta_{fi}^{r}), \qquad \frac{\mu_{oj}^{rh}}{\mu_{0j}^{rh}} = \exp(\zeta_{mj}^{r}). \tag{7}$$

Because ζ_{gt}^{r} is common across pools, these odds also hold for regional totals. We therefore calibrate the payoff from the baseline outside-partnership-to-single ratio,

$$\zeta_{gt}^{r} = \log\left(\frac{O_{gt}^{r}}{L_{gt}^{r}}\right), \tag{8}$$

and hold it fixed in counterfactuals while allowing the mass choosing an outside partnership to respond. This fixed ratio follows from the imposed logit structure and the assumption of a constant outside-partnership payoff.

Multiplying the two ratios in equation (6) and using equation (5) gives the Choo–Siow gravity equation,

$$\mu_{ij}^{rh} = \sqrt{\mu_{i0}^{rh}\mu_{0j}^{rh}}\ \exp\left(\frac{\Phi_{ij}^{r}}{2}\right). \tag{9}$$

Matches are more common when the two types create more surplus and when more potential partners of either type are available in the pool. The outside alternative leaves this gravity equation unchanged because the pairwise odds in equation (6) remain relative to singlehood. The gravity equation, the outside-partnership odds, and the feasibility conditions, which count each person as single, in a modeled partnership, or in an outside partnership, jointly determine a unique equilibrium within each pool:

$$\mu_{i0}^{rh} + \mu_{io}^{rh} + \sum_{j} \mu_{ij}^{rh} = n_{i}^{rh}, \qquad \mu_{0j}^{rh} + \mu_{oj}^{rh} + \sum_{i} \mu_{ij}^{rh} = m_{j}^{rh}. \tag{10}$$

Here, n_i^{rh} and m_j^{rh} are the masses assigned to the pool by the education-based rule above. The transfers that divide the surplus cancel from the gravity equation, so they do not have to be observed. Appendix A.2 derives equation (9) from the type-level choices. Appendix A.6 establishes uniqueness and describes how we compute the equilibrium.

Regional outcomes sum the two pool-specific equilibria:

$$\mu^r_{ij} = \sum_{h=1}^{2} \mu^{rh}_{ij}, \qquad \mu^r_{i0} = \sum_{h=1}^{2} \mu^{rh}_{i0}, \qquad \mu^r_{0j} = \sum_{h=1}^{2} \mu^{rh}_{0j}.$$

Outside partnerships aggregate in the same way: $\mu^r_{io} = \sum_h \mu^{rh}_{io}$ and $\mu^r_{oj} = \sum_h \mu^{rh}_{oj}$. When $\pi = 1/2$, each pool contains half of every type and is a scaled copy of the full regional market. Appendix A.3 proves that the two pool equilibria then sum to the standard one-pool equilibrium. Values above one half concentrate each education group in its education-centered pool and allow contact to affect matching.

2.3 Surplus and Model Outputs

We build systematic joint surplus from type-specific intercepts and bilateral party and education components. The region-specific intercepts α^r_i and γ^r_j collect factors that shift the value of every modeled match involving a given type. They primarily govern whether a type forms a modeled partnership rather than remaining single or forming an outside partnership. Let S^p_{ij} and S^e_{ij} denote the party and education components of bilateral surplus, and let $S_{ij} = S^p_{ij} + S^e_{ij}$ denote total bilateral surplus. Higher values of S_{ij} raise joint surplus. Section 5 constructs these components from the experimental estimates. A common conversion factor λ maps S_{ij} from the experiment into the units of joint surplus used in the matching model. We parameterize surplus in region r as

$$\Phi^r_{ij} = \alpha^r_i + \gamma^r_j + \lambda S_{ij}. \tag{11}$$

The sixteen values of Φ^r_{ij}—one for each female–male type pairing—form the region's 4×4 surplus matrix. The bilateral component captures how party and education change surplus across pairings. Its components and the conversion factor are common across regions, while the type intercepts vary by region. This restriction applies the experimental party and education effects uniformly across regions while allowing the baseline value of partnership for each type to differ across regions. Section 7 estimates the conversion factor λ, the contact structure, and the region-specific intercepts α^r_i and γ^r_j.

The surplus specification restricts (1) which components enter joint surplus and (2) how they are combined. The female- and male-type intercepts enter joint surplus additively, and bilateral surplus contains only the party and education components measured in the experiment. A single conversion factor scales both components, so the relative magnitudes estimated in the experiment are preserved. The party and education components also enter additively, without an additional effect for their combination. Appendix A.4 explains these restrictions and their role in identification.

Given the surplus matrix and the type masses in each pool, the equilibrium determines the number of couples in every type pairing, the singlehood rate and outside-partnership share of every type, and the systematic payoffs U^{rh}_{ij} and V^{rh}_{ij}. It also determines ex-ante welfare, or expected utility before idiosyncratic tastes are realized. Under the Type-I extreme-value assumption, the change in a type-i woman's expected utility between two equilibria equals the change in the following inclusive value (McFadden, 1978; Galichon and Salanié, 2022):

$$\mathcal{W}_i^{rh} = \log\left(1 + \exp(\zeta_{fi}^r) + \sum_j \exp(U_{ij}^{rh})\right). \tag{12}$$

Inside the logarithm, 1 represents singlehood, whose systematic payoff is normalized to zero; $\exp(\zeta_{fi}^r)$ represents an outside partnership; and the sum includes every male partner type available in the pool. For a type-j man, the analogous expression replaces U_{ij}^{rh} with V_{ij}^{rh} and ζ_{fi}^r with ζ_{mj}^r, and sums over female types i. A type's regional welfare averages its two pool-specific values using its pool-assignment probabilities. Appendix A.5 derives this inclusive-value expression from the individual choice problem. Welfare is measured in normalized utility units. Section 5 maps welfare changes into partner-income equivalents.

We report welfare changes as an average among adults included in the equilibrium and as an average among all Democratic and Republican adults aged 25–45. Let $N = \sum_{r,g,t} N_{gt}^r$ and $N^{\text{eq}} = \sum_{r,g,t} N_{gt}^{\text{eq},r}$, and write $\Delta\mathcal{W}_{gt}^r$ for a type's regional welfare change. Holding population masses constant,

$$\Delta\overline{\mathcal{W}}_{\text{eq}} = \frac{\sum_{r,g,t} N_{gt}^{\text{eq},r}\Delta\mathcal{W}_{gt}^r}{N^{\text{eq}}}, \qquad \Delta\overline{\mathcal{W}}_{\text{all}} = \frac{N^{\text{eq}}}{N}\Delta\overline{\mathcal{W}}_{\text{eq}}. \tag{13}$$

The first expression averages welfare changes among the adults included in the equilibrium. The second spreads the same aggregate change across all Democratic and Republican adults by multiplying the first average by the share N^{eq}/N. Singles seeking only same-sex partners and adults in current same-sex partnerships remain in the overall denominator but do not contribute to welfare change. Reported singlehood similarly adds $L_{gt}^{\text{SS},r}$ to the singles selected in equilibrium and divides by N.

The counterfactuals change preferences, regional type masses, and contact. After each change, we resolve the equilibrium in every meeting pool and compare the resulting sorting, singlehood, and welfare with the baseline.

3 Experimental Design

The matching framework in Section 2 requires estimates of how people value a potential partner's party and education. We identify these preferences using a discrete choice experiment. Respondents choose between two profiles that vary in party identification, education, income, attractiveness, religiosity, and family plans. From these choices, we estimate how partisan distance, crossing the party line, and crossing the BA threshold affect the utility respondents assign to a profile, holding the other displayed traits fixed. Variation in partner income places these effects on a common scale. Because income enters the choice model in logs, we express the tradeoffs as percentage changes in partner income and, where useful, as dollar amounts evaluated at the experiment's median profile

Figure 1: A paired-profile choice task

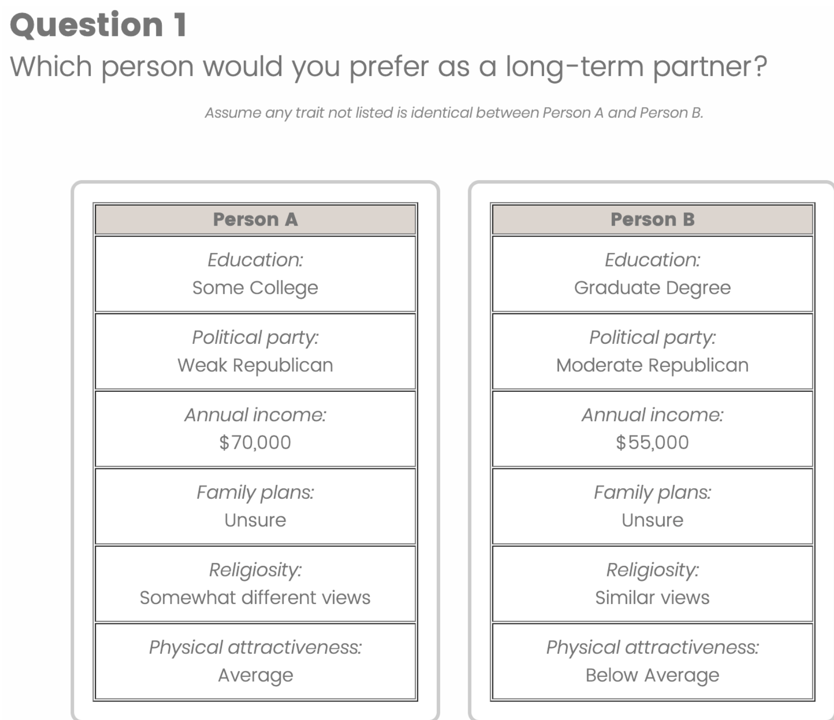

Notes. Screenshot from the survey instrument. Profile gender is not displayed. Reported partner-gender preference determines the earnings grid; for respondents open to both women and men, the survey selects one of the two grids at random. The order of the six attributes is randomized across respondents and held fixed across a respondent's tasks.

income. The analysis sample contains 423 Democrats and Republicans aged 25–45 recruited through Prolific.

3.1 Choice Tasks and Profiles

Each respondent completes 12 paired-profile tasks, choosing the person they would prefer as a long-term partner from two hypothetical profiles. Respondents are asked to answer as if they were single and to treat each profile as someone of a gender they would consider. The instructions state that every trait not listed is the same between the two people, including age. This isolates tradeoffs among the displayed traits from beliefs about their unlisted correlates. Among otherwise eligible respondents, 79 percent correctly identify before the tasks that unlisted traits should be held constant and 93 percent correctly identify that they should answer as if single; respondents who answer incorrectly see the correct instruction before continuing. In a post-task check, 80 percent describe the unlisted traits as similar across the two profiles (Appendix B.3). Figure 1 shows the task respondents faced.

The two profiles list six characteristics, shown in Table 1. The design allows party to vary independently of the other displayed traits. Income is drawn from education- and partner-gender-specific grids, so profiles at a given education level can have different earnings while implausible schooling–income combinations are excluded. Attractiveness, religiosity, and family plans are also shown because respondents might otherwise infer them from a partner's party or education.

Table 1: Attributes shown on partner profiles

Attribute	Displayed levels	Number of levels
Party	Strong Democrat through strong Republican	6
Education	High school, some college, BA, graduate degree	4
Income	Education- and partner-gender-specific grid	7
Attractiveness	Below average, average, above average	3
Religiosity	Similar, somewhat different, very different	3
Family plans	Wants children, unsure, does not want	3

Party, education, religiosity, and family plans are coded as differences between the respondent and the profile. Attractiveness and income use the profile's displayed level. Appendix B.1 lists the exact coding of every displayed level.

3.2 Choice Model

We analyze the choices using a standard random-utility model for paired alternatives (Manski, 1977; McFadden, 1974; Train, 2009). Respondent i chooses profile A when it provides more utility than profile B. We write utility from profile k as

$$\widetilde{u}_i(k) = u_i(k) + \varepsilon_i(k), \tag{14}$$

where $u_i(k)$ is the component determined by the profile's displayed attributes and $\varepsilon_i(k)$ captures profile-specific factors not represented by those attributes.

Party and education are expressed relative to the respondent. Party enters through two terms: absolute distance on the six-level partisan scale and an indicator for whether the respondent and profile have different party affiliations. The first captures how preferences change with partisan distance; the second allows an additional change when a profile crosses the party line. Education enters through an indicator for whether the respondent and profile fall on opposite sides of the BA threshold. Using the same education distinction that defines types in the matching model lets the experimental estimate enter the structural model directly. For family plans, separate indicators distinguish a one-category difference from disagreement between opposite ends of the scale, allowing the latter to carry a different penalty.

To write these terms formally, let p_i and p_k denote the respondent's and profile's partisan positions, and let P_i and P_k denote their party affiliations. Let BA_i and BA_k indicate that the respondent and profile, respectively, have at least a bachelor's degree. Let y_k denote the profile's annual income in dollars and define $g(y_k) = c\log(y_k/65{,}000)$, with $c = 6.5$. Only differences in profile utility affect choices, so the income contribution to a comparison between profiles A and B is $c\log(y_A/y_B)$. Let x_{ik} contain controls for attractiveness, religious distance, the two family-plan differences, and whether the profile earns more than the respondent, and let $\boldsymbol{\beta}_i^{\text{ctrl}}$ denote the corresponding coefficients. We parameterize $u_i(k)$ as

$$u_i(k) = -\beta_i^{\text{dist}}|p_i - p_k| - \beta_i^{\text{party}}\mathbf{1}[P_i \neq P_k] - \beta_i^{\text{edu}}\mathbf{1}[\text{BA}_i \neq \text{BA}_k] + x_{ik}'\boldsymbol{\beta}_i^{\text{ctrl}} + g(y_k). \tag{15}$$

The minus signs orient the three focal coefficients as penalties. The distance coefficient β_i^{dist} is the penalty from moving one step farther apart while holding fixed whether the profile crosses the party line. The party coefficient β_i^{party} is the additional penalty from crossing that line while holding partisan distance fixed. We call this additional penalty the *discrete party-line shift*. These two estimates combine into the binary cross-party penalty used in the matching model. The education coefficient β_i^{edu} is the signed penalty from crossing the BA threshold: a positive value favors the respondent's own side, while a negative value favors the other side.

3.3 Estimation

We assume that the shocks $\varepsilon_i(k)$ are independent Type-I extreme value draws with scale σ. Under this standard specification, the choice probability conditional on a respondent's coefficients takes the conditional-logit form:

$$\Pr_i(A \succ B) = \left[1 + \exp\left(-\frac{u_i(A) - u_i(B)}{\sigma}\right)\right]^{-1}. \tag{16}$$

A larger σ means that the same difference in displayed attributes produces a smaller change in choice probability.

Choice probabilities identify utility only up to scale, so one coefficient must be normalized. We fix the coefficient on $g(y_k)$ at one and estimate σ (Train and Weeks, 2005). The log form allows the response to a given dollar increase to decline as income rises. It predicts held-out choices better than linear income in all five respondent-level folds, so log income is the main specification and linear income is a robustness check (Appendix C.2).

Each respondent has a coefficient vector that remains fixed across the 12 tasks. Coefficients may differ across respondents, including respondents in the same observable type. Twelve choices are too few to estimate every coefficient independently for each person, so a hierarchical Bayesian mixed logit combines each respondent's choices with information from other respondents while retaining person-specific coefficients (Train, 2009). The coefficients are normal draws around group means. The means of β_i^{dist}, β_i^{party}, and β_i^{edu} vary across eight respondent types: women and men crossed with Democratic or Republican identification and BA or non-BA education. The mean coefficient on whether the profile earns more than the respondent varies by sex and party. The means for attractiveness, religious distance, and family plans, along with the dispersion of every coefficient, are common across types. The hierarchy therefore combines type-level means with individual preference variation. Appendix B.2 states the full hierarchy.

We jointly estimate one error scale σ, common to all respondents, with the preference distributions. The posterior draws of the type-specific partisan-distance, party-line-shift, and education means provide the experimental inputs to bilateral surplus in Section 5. The attractiveness, religiosity, family-plan, and higher-earner coefficients remain in the choice model as controls but do not enter bilateral surplus. Appendix B.2 provides the parameter estimates and convergence diagnostics.

3.4 Task Assignment and Data Collection

Profile pairs were assigned using either a randomized or an adaptive design. Random assignment provides broad coverage of the design space. For adaptive assignment, we use the Bayesian Adaptive Choice Experiment of Drake et al. (2024). The procedure begins with a distribution of possible preference coefficients; choices from the pilot and randomized-main batches supplied this starting distribution. After each choice, it updates the distribution for that respondent and selects the next profile pair expected to reduce uncertainty about the coefficients the most. Both assignment designs used the same 12-task format and profile attributes. Appendix B.4 explains the adaptive procedure and compares how much information the adaptive and randomized tasks provide. Appendix B.6 uses simulations to compare recovery under the implemented design with an equal-budget randomized benchmark.

Beginning with the randomized-main batch, respondents were told that their choices would generate a personalized summary comparing their own preferences with those of other respondents. Respondents who wanted an accurate summary therefore had a reason to choose carefully. The design does not separately identify whether the feedback changed behavior. Appendix B.5 describes the feedback and the conditions under which it can encourage sincere responses.

The final choice model pools all paired-profile choices across the four batches and both assignment designs. Appendix B.3 maps the assignment designs to the four data-collection batches.

3.5 Analysis Sample

The analysis uses 423 respondents recruited through Prolific, an online participant-recruitment platform, who satisfy the sample criteria. Respondents are Democrats or Republicans aged 25–45 who report a binary gender and are open to a partner of the other sex. Of the 423 respondents, 363 report interest only in partners of the other sex and 60 report interest in both women and men. Eligibility therefore depends on reported partner-gender interest, not the gender of a respondent's current partner. These criteria place each respondent in one of the eight types defined above. Every respondent in the analysis sample completed all 12 choice tasks and passed the preregistered attention check. As described in Section 2, the structural counterfactuals hold current same-sex partnerships fixed.

Appendix B.3 gives the batch structure and full sample cascade, compares the sample with the population used by the matching model, and reports the attention and comprehension checks.

Figure 2: Raw party and education choices across income gaps

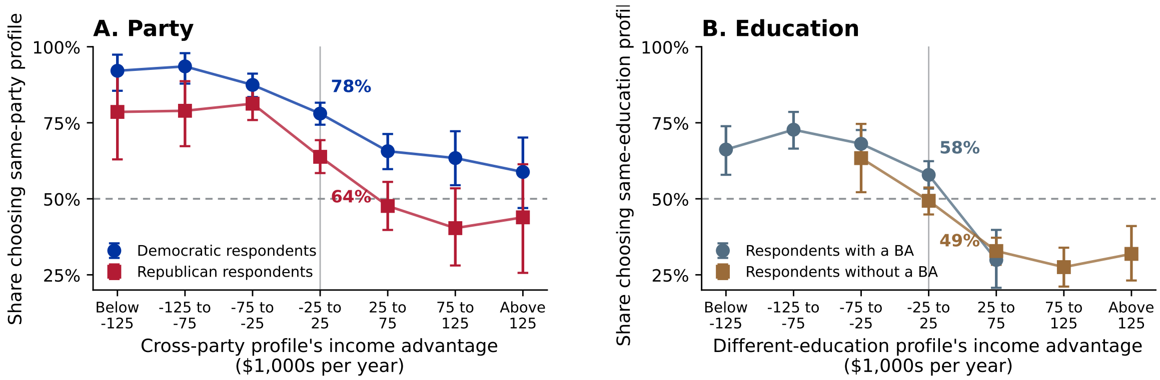

Notes. Each marker is the observed choice share within an income-gap bin; thin segments connect adjacent bins. The center bin contains gaps from -$25,000 to $25,000; the four interior bins span $50,000 each; the endpoints collect gaps beyond ±$125,000. Panel A uses all 2,562 tasks with one same-party and one cross-party profile. Panel B uses all 2,590 tasks with one profile on each side of the respondent's BA threshold. BA profiles draw from higher income levels, so the income-gap range differs for respondents with and without a BA; bins with fewer than 20 tasks for a group are omitted. Appendix Table 13 reports every group–bin count, including omitted cells. Vertical lines show 95 percent confidence intervals that account for repeated choices by the same respondent. The dashed line marks 50 percent.

4 Experimental Evidence

The experiment finds that party matters far more than education in partner choice. After accounting for the other displayed traits, the choice model predicts an 81 percent same-party choice share between otherwise comparable profiles, compared with 54 percent for a profile on the respondent's side of the BA threshold. The party effect is comparable in magnitude to the effects of family plans, attractiveness, and religion, while education matching lies well below the other displayed traits. Party preferences are stronger among Democratic respondent types but appear in both parties, and measured party preferences predict the political proximity of respondents' actual partners.

4.1 Party, Education, and Income in the Raw Choices

Figure 2 plots raw choice rates against the income advantage of the cross-party profile in Panel A and the profile across the respondent's BA threshold in Panel B. Positive values mean that this profile earns more.

Respondents choose the same-party profile in 73 percent of choice tasks when the two profiles' incomes differ by no more than $25,000. The share is 78 percent among Democrats and 64 percent among Republicans. It declines as the cross-party profile's income advantage grows, but remains 53 percent when that profile earns more than $125,000 extra. The raw choices therefore show a strong same-party preference even when the cross-party profile has a large income advantage.

Respondents display little overall preference for education matching in the raw choices. When the profiles' incomes differ by no more than $25,000, respondents with a BA choose the profile on their side of the BA threshold in 58 percent of tasks, while respondents without a BA do so in

49 percent. The latter estimate implies no systematic preference for a non-BA profile; a general preference for a BA partner would instead push their same-side share below 50 percent. Outside the center bin, both groups tend to favor whichever profile earns more. Some income-gap bins do not appear for both groups because the education-specific income grids make those comparisons rare; the figure omits group–bin cells with fewer than 20 tasks.

Figure 3 shows the response to income across all tasks with unequal profile incomes. Respondents choose the higher-income profile 56 percent of the time when the gap is no more than \$25,000, 66 percent when the gap is between \$30,000 and \$100,000, and 69 percent when it exceeds \$100,000. Most of the increase occurs before the largest income gaps, which is consistent with a diminishing response to a given dollar difference. The raw bins alone do not distinguish log from linear income because the other profile attributes continue to vary and choice probabilities are nonlinear. Even at the largest gaps, nearly one-third of choices favor the lower-income profile. Family-plan agreement is the least responsive to a large income advantage for the other profile (Appendix Figure 12).

Figure 3: Raw response to differences in partner income

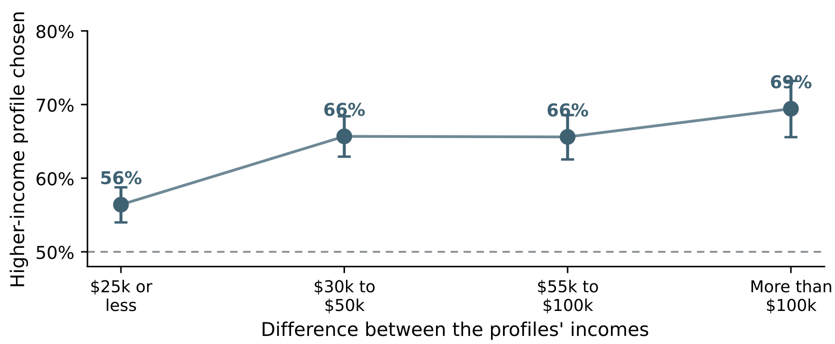

Notes. The figure uses all 4,757 tasks in which the profiles have different incomes and groups them by the absolute income difference. Markers are raw choice shares; vertical lines are 95 percent confidence intervals that account for repeated choices by the same respondent. Other displayed attributes continue to vary.

4.2 Choice-Model Estimates

The choice model described in Section 3.2 is fitted to all 5,076 assigned choices. We then use the fitted model to predict choices between standardized pairs of profiles that differ only in the attribute being compared. For each posterior draw, the calculation averages over the estimated distribution of preferences within each respondent type and then gives the eight types equal weight. A share of 50 percent indicates no average preference between the two profiles; movement away from 50 percent indicates a stronger preference for the trait named first.

Figure 4 places the headline party comparison on the same predicted-choice scale as the other displayed traits. Moving from a cross-party to a same-party profile changes two components of utility: partisan distance falls by 2.11 steps on average under the experimental design, and the profile crosses onto the respondent's side of the party line.[6] Holding income and the other displayed traits fixed, this combined comparison gives a predicted same-party choice share of 81 percent, with a 95 percent credible interval of [78, 84] percent. A one-step reduction in distance alone gives 65 percent [62, 67]. Changing only the party affiliation at a fixed distance gives 67 percent [62, 72]. For example, moving from a moderate Democrat to a weak Democrat changes distance only, whereas moving from a weak Democrat to a weak Republican changes distance by one step and activates the separate party-line shift. The combined comparison supplies the party component of bilateral surplus in Equation (17).

At the median displayed partner income of $65,000, a one-step reduction in partisan distance is equivalent to about $23,700 in annual partner income, or 36 percent (95 percent credible interval: $19,500–$28,000, or 30–43 percent). The log specification also permits comparisons that do not depend on a reference income. If the cross-party profile earns twice as much as the otherwise comparable same-party profile, the predicted same-party choice share is 70 percent [66, 74]. If it earns three times as much, the share is 61 percent [57, 66]. These comparisons occur in the experiment: the cross-party profile earns at least twice as much in 22 percent of the 2,562 party-comparison tasks (571 tasks) and at least three times as much in 10 percent (252 tasks).

Log income predicts held-out choices better than linear income in all five prespecified respondent-level folds. We therefore use log income in the main specification and retain linear income as a robustness check (Appendix C.2).

The remaining rows of Figure 4 apply the same measure to the other displayed attributes. The predicted shares are 77 percent for shared rather than opposite family plans, 77 percent for above- rather than below-average attractiveness, and 74 percent for similar rather than very different religious views. The corresponding share is 54 percent for a profile on rather than across the respondent's BA threshold. Partisanship is therefore among the leading displayed attributes, while education matching lies well below the others. A separate post-task rating produces the same broad ranking: respondents place family plans first, with political identity and attractiveness close behind (Appendix C.4). This ordering also accords with prior experiments and online-dating evidence that party labels affect romantic evaluations and partner choice (Huber and Malhotra, 2017; Nicholson et al., 2016; Easton and Holbein, 2021).

[6] Giving equal weight to each of the three partisan positions on each side for both respondents and profiles, mean distance is 3 steps for cross-party profiles and 8/9 of a step for same-party profiles. The difference is $3 - 8/9 = 19/9 \simeq 2.11$. Appendix D.3 derives the factor and compares alternative weights.

Figure 4: Predicted choices across partner attributes

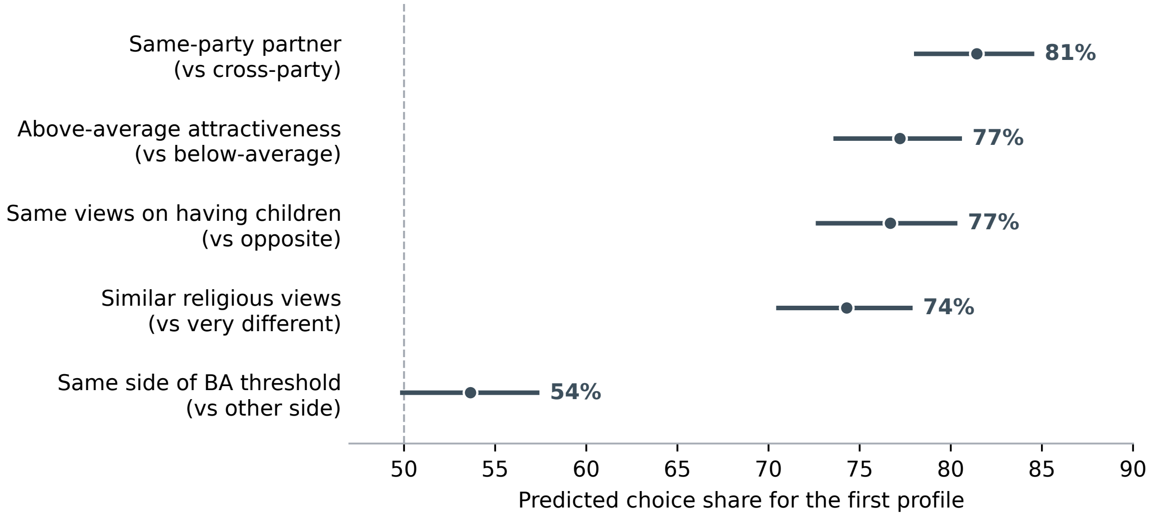

Notes. Posterior medians with 95% credible intervals from the hierarchical choice model. Each row gives the predicted probability of choosing the first profile over an otherwise identical second profile. The calculation integrates over estimated preference heterogeneity within each respondent type and then gives the eight types equal weight. The party row combines the average 2.11-step difference in partisan distance with the shift from crossing the party line. The education row compares a profile on the respondent's side of the BA threshold with one on the other side. The party and education rows provide the preference inputs used by the matching model; the other rows are experimental benchmarks. Attractiveness and religiosity span two steps on their three-level scales. The dashed line marks 50 percent.

The 2020 American Perspectives Survey provides a second ordinal benchmark. Among U.S. adults, the shares rating a characteristic one of the most important or very important when deciding whom to date were 59 percent for shared views on having children, 39 percent for a shared religious background, 31 percent for shared political views, and 16 percent for a prospective date having at least a four-year degree (Cox et al., 2020). Family plans, religion, and politics again rank above education.

4.3 Heterogeneity across Respondent Types

The largest source of heterogeneity is the respondent's own party. Figure 5 reports predicted choices separately for the eight respondent types. Averaging across gender and education, the predicted same-party choice share is 91 percent [89, 93] among Democratic respondent types and 72 percent [66, 77] among Republican respondent types. The Democratic share is larger in every posterior draw, and the same ordering appears within gender and education categories. The middle panel isolates the party-line shift while holding partisan distance fixed. Its predicted same-side shares are 77 percent among Democrats and 57 percent among Republicans; the Republican interval includes 50 percent. Republicans still prefer politically closer profiles on average, but partisan distance accounts for most of their combined same-party preference.[7]

The model also allows respondents with the same observed sex, party, and education to have different coefficients. Averaged equally across the eight types, the estimated within-type distributions imply that 89 percent of respondents prefer a same-party to an otherwise comparable cross-party profile (Appendix C.3).

Education preferences are modest and heterogeneous. Panel C orients each estimate toward a profile on the respondent's own side of the BA threshold, matching the coefficient used by the structural model. Averaging across gender and party, the predicted same-side share is 58 percent [52, 63] for BA respondent types and 49 percent [44, 55] for non-BA respondent types. Expressed instead as a preference for a BA profile, three of the four BA types favor BA, while two of the four non-BA types favor BA and two favor non-BA. The experiment therefore provides much weaker evidence for education preferences than for partisan alignment, despite strong education sorting in observed couples. The structural model attributes the remaining education-sorting gap to the degree of separation between education-centered meeting pools (Section 7.3).

Figure 5: Predicted choices by respondent type

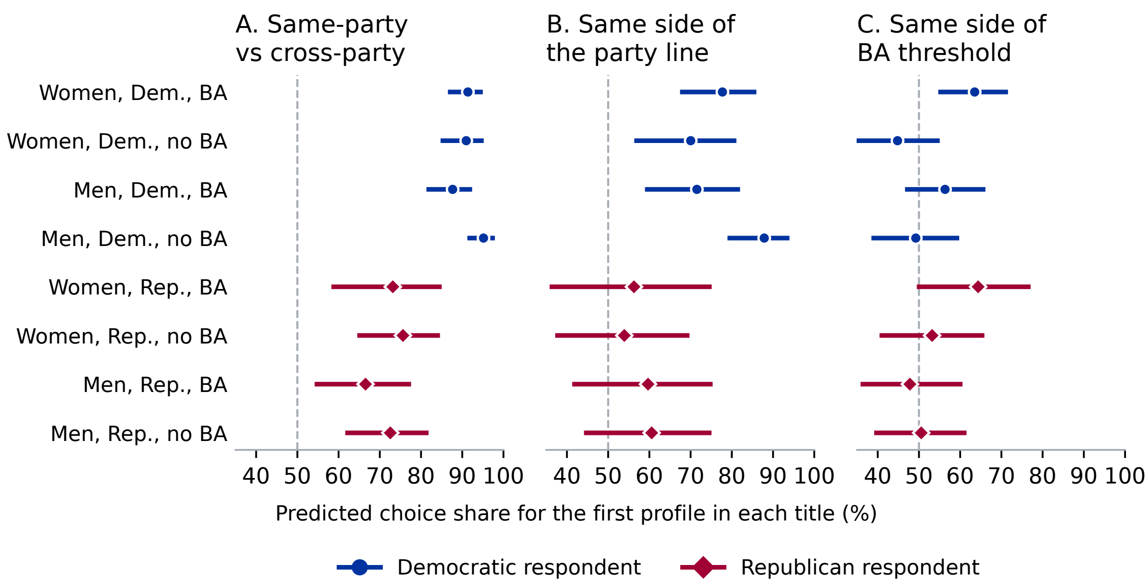

Notes. Posterior medians with 95% credible intervals by respondent type. Each estimate is the predicted probability of choosing the first profile in the panel title over an otherwise identical second profile, integrating over estimated preference heterogeneity within the type. Panel A combines the average 2.11-step difference in partisan distance with the party-line shift. Panel B holds partisan distance fixed. Panel C compares a profile on the respondent's side of the BA threshold with one on the other side. The dashed line marks 50 percent.

[7] This Democratic–Republican asymmetry differs from Iyengar et al. (2018), where Republicans more often object to a child marrying across party lines. Their respondents evaluate a child's partner; respondents here choose for themselves.

4.4 What a Party Label Captures

Party labels can influence choices through both the political positions and the broader social meanings they convey. To distinguish these channels, we fielded a mechanism module with 120 respondents. The module followed the main choice experiment. In the module, respondents chose between two profiles that independently varied party, income, and positions on abortion, same-sex marriage, immigration, and gun purchases.[8] Appendix C.5 gives the design, sample construction, and full results. Democratic respondents prefer the more liberal position on all four issues. Republican respondents' point estimates favor the more conservative position on all four, although their estimates are less precise in the smaller Republican sample.

The mechanism module estimates the residual party-label effect: the change in profile utility when only the party label switches from the respondent's party to the opposing party, holding income and the four positions fixed. A hierarchical choice model allows the label and issue-position preferences to vary across respondents and estimates separate population averages for Democratic and Republican respondents. The predicted same-party-label choice share is 79 percent, with a 95 percent credible interval of [71, 86] percent. This residual is economically large, but the 79 and 81 percent shares are not a decomposition because they come from experiments with different profile designs and choice scales. The within-module decomposition attributes 27 percent of the combined penalty to the residual label among Democratic respondents and 67 percent among Republican respondents (Appendix E.10).

A separate follow-up asks what an opposing-party label brings to mind. Specific issues are selected by 75 percent of respondents, moral values or worldview by 68 percent, political leaders or style by 65 percent, family roles by 59 percent, and how a partner might treat them or people like them by 51 percent. Only 14 percent select ambition, work ethic, or money.

The large party penalty may also reflect beliefs that Democrats and Republicans are far apart on policy. For each issue, respondents estimated how many out of 100 Democrats and Republicans would support the liberal position displayed in the module: abortion as a personal choice, legal same-sex marriage, a path to citizenship for unauthorized immigrants, and stricter gun-purchase rules. Because all four positions are liberal, the perceived partisan gap is predicted Democratic support minus predicted Republican support. The benchmark is the survey-weighted gap in support for the same positions among Democratic and Republican identifiers and leaners aged 25–45 in the 2024 ANES (ANES, 2025).

Democratic respondents perceive wider partisan gaps than Republican respondents on each issue. Both groups overstate the immigration gap, although their predictions differ. Democratic respondents predict 71 percent support among Democrats and 23 percent among Republicans, while Republican respondents predict 74 and 41 percent; the corresponding ANES shares are 58 and 40 percent. The perceived gaps are therefore 48 and 33 percentage points, respectively, compared with

[8] Issue selection combined evidence from the American Perspectives Survey (APS; Cox et al., 2020) on which disagreements make dating difficult with standard items showing within-party variation and cross-party overlap in the 2024 American National Election Studies (ANES, 2025). Within the APS's broader LGBT-rights domain, we chose same-sex marriage for its relevance to partnership and concise wording.

18 points in the ANES. On same-sex marriage, Democratic and Republican respondents perceive gaps of 45 and 38 points, compared with 35 points in the ANES. Averaged across all four issues, the gaps are 50 and 36 points ($p = 0.046$). Panel A of Figure 16 in Appendix C.5 shows the group-specific predictions and ANES benchmarks for all four issues.

Respondents who perceive wider partisan gaps also express larger partisan-distance penalties in their choices. We summarize each respondent's party preference with the personalized-feedback score calculated from their 12 main-experiment choices. Among the 100 Democratic and Republican respondents, the Spearman rank correlation between the perceived partisan gap and this experimental preference score is 0.25 ($p = 0.012$). Comparing respondents within party gives a partial rank correlation of 0.20 ($p = 0.043$). The mechanism evidence therefore connects the party penalty to perceived policy distance and to the broader social meanings respondents associate with party labels (Appendix Figure 16).

4.5 Validation with Actual Partnerships

Respondents' actual partnerships provide two validation checks. First, same-party partnering in the experimental sample closely matches national data: 86 percent of partnered respondents whose partners identify with or lean toward either major party are in same-party couples. The 14 percent cross-party share is close to the 16 percent share in the HCMST calibration sample (Section 6) and resembles the same-party partnering that Huber and Malhotra (2017) document in online dating.

Second, party preferences measured in the experiment predict actual-partner proximity. We place the respondent and partner on the experiment's six levels of party identification, from strong Democrat to strong Republican. Partner proximity ranges from zero at opposite endpoints to five at the same position.

The rank correlation between the experimental party-preference score and actual-partner political proximity is 0.18 ($p = 0.004$). After the choice tasks, a separate follow-up asks respondents to rate the importance of a partner's political identity on a one-to-five scale. That rating has a similar correlation with partner proximity, 0.17 ($p = 0.004$), although it is not plotted. Panel A of Figure 6 plots mean partner proximity by bins of the experimental score.

Education provides a different comparison. To make the measure comparable across respondents with and without a BA, the education score measures how strongly each respondent favors profiles on their own side of the BA threshold. Among 293 partnered respondents who report a partner's education, 76 percent have a same-education partner. The education score is only weakly related to whether a respondent has a same-education partner: the rank correlation is 0.07 ($p = 0.25$). A separate post-task rating of education importance has a similar correlation of 0.06 ($p = 0.32$). The individual-level data therefore combine substantial education matching with little relationship between that matching and measured education preferences. The matching model accounts for this pattern through education-centered meeting pools.

Figure 6: Experimental preferences and actual-partner matching

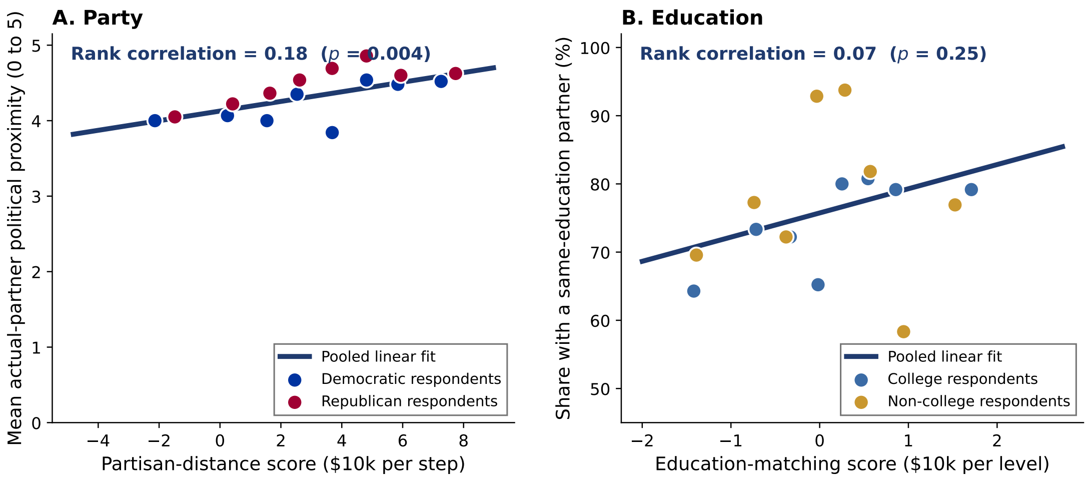

Notes. This is a descriptive binned scatterplot. Respondents are assigned to eight quantile bins based on the experimental score. Markers show group-specific means within bins; lines are pooled least-squares fits to individual observations. Panel A reports actual-partner political proximity, defined as five minus distance on the six-position scale. Panel B reports same-education shares. The party score is the estimated penalty for a one-step increase in partisan distance. The education score is oriented so that a larger value means a stronger preference for profiles on the respondent's own side of the BA threshold. Both scores come from the seven-coefficient model used to produce personalized feedback; their scale does not affect the ranks, bins, or correlations. Partnered subsamples contain 261 respondents in Panel A and 293 in Panel B.

5 From Experimental Preferences to Match Surplus

The experiment provides two inputs to the matching model: how a potential partner's party and education affect utility, measured relative to log partner income. This section combines women's and men's estimates into bilateral surplus and explains the conversion factor that determines how strongly these experimentally measured differences affect observed matching.

5.1 From Experimental Estimates to Bilateral Surplus

Attractiveness, religiosity, family plans, and whether the profile earns more than the respondent serve as controls in the choice model. Because these traits vary independently of party and education, they allow the party and education coefficients to compare otherwise similar profiles. The controls do not enter the matching model, whose partner types are defined by party and education.

Party. For respondents of gender g and party-by-education type t, let $\bar{\beta}_t^{\text{dist},g}$ and $\bar{\beta}_t^{\text{party},g}$ denote the type-level mean penalties from moving one step farther apart on the six-level partisan scale and from crossing the party line, respectively. The matching model distinguishes same-party from cross-party matches, so we combine these estimates into the cross-party penalty

$$\psi_t^g = \bar{\beta}_t^{\text{party},g} + \bar{\beta}_t^{\text{dist},g}\bar{d}_p, \tag{17}$$

where $\bar{d}_p = 2.11$ is the average increase in partisan distance from a same-party to a cross-party profile under the experimental design (Section 4.2; Appendix D.3).

Education. The education term is the type-level mean penalty from crossing the respondent's BA threshold:

$$\tau_t^g = \bar{\beta}_t^{\text{edu},g}. \tag{18}$$

A positive τ_t^g favors the respondent's own side of the threshold, while a negative value favors the other side.

Bilateral surplus. The individual penalties use the choice model's utility units, in which partner income enters as $c \log(y/65{,}000)$ with $c = 6.5$. In the woman–man equilibrium defined in Section 2, a match combines the preferences of both partners, so each bilateral component sums the woman's and man's penalties. For a woman of type i and a man of type j,

$$S_{ij}^p = -\left(\psi_i^f + \psi_j^m\right) \mathbf{1}[P(i) \neq P(j)], \qquad S_{ij}^e = -\left(\tau_i^f + \tau_j^m\right) \mathbf{1}[E(i) \neq E(j)]. \tag{19}$$

where $P(t)$ and $E(t)$ denote type t's party and education. The minus signs convert penalties into surplus: a pairing receives a lower value when either partner dislikes the relevant difference. We repeat this mapping for every posterior draw, carrying experimental uncertainty into the structural estimates and counterfactuals.

5.2 Income Interpretation and the Conversion Factor

The log-income specification gives the experimental estimates a direct interpretation. If $a > 0$ is an individual's penalty for a partner characteristic, avoiding that characteristic is equivalent to forgoing the share

$$w(a) = 1 - \exp(-a/c) \tag{20}$$

of a potential partner's income. At partner income y, the corresponding dollar amount is $y\,w(a)$. The local dollar equivalent, reported in Section 4.2, is ya/c. These tradeoffs do not require the matching model or its conversion factor.

The conversion factor λ instead links the experimentally measured surplus differences to behavior in the marriage market. As specified in equation (11), the matching model includes the bilateral component as λS_{ij}. A larger λ makes the experimental party and education differences generate larger differences in matching surplus and therefore stronger sorting across partner types. The scale of unobserved tastes in the experiment need not equal the scale of unobserved determinants of real partnerships, so λ is estimated from observed couple pairings rather than imposed from the experimental normalization.

The same bridge gives a partner-income interpretation to counterfactual welfare. For a welfare change $\Delta\mathcal{W}_t = \mathcal{W}_t^{CF} - \mathcal{W}_t^0$, define its log-partner-income equivalent as

$$q_t = \frac{\Delta\mathcal{W}_t}{c\lambda}. \tag{21}$$

We report q_t as an approximate percentage change, $100q_t$ percent, and as a local dollar equivalent, $y^\star q_t$, at a stated reference level of annual partner income $y^\star$. The percentage measure does not depend on the reference income; the dollar measure does.

5.3 What the Experiment Does and Does Not Supply

The experiment determines the relative party and education components of bilateral surplus. It does not determine overall partnership rates, the payoff from an outside partnership, or how people are distributed across meeting pools. Section 6 constructs the regional population inputs, including the groups held fixed outside the woman–man equilibrium and the baseline masses choosing singlehood, modeled partnerships, and outside partnerships. Section 7 then estimates the conversion factor λ, the own-pool share π, and the region-specific type intercepts α_i^r and γ_j^r from observed singlehood and couple pairings. The outside-partnership payoff is calibrated from the baseline outside-partnership-to-single ratio as described in equation (8).

6 Data

Four data sources provide the population and couple inputs to the matching model. The Cooperative Election Study (CES, Kuriwaki, 2022) supplies regional population shares and singlehood rates by sex, party, and education. The How Couples Meet and Stay Together study (HCMST, Rosenfeld et al., 2023) supplies regional couple counts by female and male type and, among adults classified as single, identifies those who report attraction only to people of the same sex. The American Community Survey (ACS) supplies linked partners' age, sex, and education and identifies existing same-sex partnerships, while the American National Election Studies (ANES) help interpret uncertain partner-party reports in HCMST. We use the four census regions as markets because finer geographic divisions would leave too few HCMST couples in the party-by-education pairing cells. Regional population composition and meeting-pool assignment jointly determine which partner types are available. Baseline singlehood and partnership shares discipline the value of the three equilibrium outcomes, while the couple distribution informs sorting across partner types.

6.1 The Cooperative Election Study

We use the CES to measure regional type masses and singlehood rates. This nationally representative survey interviews more than 60,000 respondents annually. We pool the 2017–2024 waves and retain respondents aged 25–45 who identify with or lean toward the Democratic or Republican Party. These age and party restrictions match the experimental sample and the matching model's type space. Non-leaning independents, about 21 percent of CES respondents in this age range, are excluded. The final sample contains 85,146 respondents.

Within each region, the CES supplies the weighted population share and singlehood rate for each sex-and-type group. The matching model normalizes the regional population to one, so the population shares are its type masses and preserve the region's sex ratio and party-by-education composition. A group's singlehood rate follows the CES marital-status categories. Respondents who select "Married" or "Domestic Partnership" are counted as partnered; those who select "Separated," "Divorced," "Widowed," or "Single/Never Married" are counted as single. In the analysis sample, 44.6 percent are married and 6.8 percent report a domestic partnership. The CES item misses some cohabiting couples, leaving its partnered share closer to the married share in household data than to the married-or-cohabiting share. Responses with missing marital status are excluded. Every regional group contains more than 1,000 respondents. Appendix E.1 compares the CES measure with ACS marital-status and co-residence measures. For the historical comparison, we fit the matching model using ACS marriage and cohabitation rates while retaining CES information on party (Appendix F).

Table 2 reports the national versions of these inputs. Democratic women with a BA make up 16.3 percent of the population, compared with 12.2 percent for Democratic men with a BA. The pattern reverses among Republicans without a BA: men account for 13.8 percent of the population and women for 11.3 percent. Singlehood rates are highest among Democratic men without a BA, at 66.6 percent, and lowest among Republican women with a BA, at 31.1 percent.

Table 2: Population shares and singlehood rates by sex, party, and education (%, CES 2017–2024, ages 25–45)

	Population share		Singlehood rate	
Type	Men	Women	Men	Women
D, BA	12.2	16.3	42.1	44.9
D, NoBA	16.3	15.7	66.6	55.7
R, BA	6.8	7.6	32.9	31.1
R, NoBA	13.8	11.3	54.7	38.7

Notes. Group shares give each group's share of the full population; singlehood rates are calculated within groups. Estimates are survey weighted, giving each year equal weight. D/R include party identifiers and leaners; BA/NoBA indicate respondents with and without a bachelor's degree. N = 85,146; non-leaning independents are excluded.

6.2 How Couples Meet and Stay Together

We use HCMST to measure regional couple counts by female and male type. The study is nationally representative of U.S. adults and records relationship histories and partner characteristics. The 2017 baseline surveyed 3,510 respondents. The 2020 and 2022 follow-ups re-interview the same respondents without remeasuring party identification, so we construct the couple distribution from the baseline.

The estimation sample contains 485 married or cohabiting opposite-sex couples whose respondent belongs to HCMST's 25–44 age group and whose party and education can be coded for both partners. HCMST reports the partner's age in ten-year intervals, some of which cross the model's 25–45 boundary. We therefore use exact partner ages from the ACS to estimate the probability that partners in those intervals fall within the model's age range. Appendix E.4 gives the sample construction and weighting procedure.

Uncertain partner-party reports are the main measurement concern in HCMST. Respondents classify 36.6 percent of partners as undecided, independent, or other, and the baseline couple table excludes those observations. We assess that choice with two ANES-based checks: reallocating the uncertain HCMST reports and replacing HCMST with a couple table in which both partners report their own party (ANES, 2025). Cross-party couples account for 19.8 percent after the reallocation and 14.4 percent in the ANES couple table, compared with 16.0 percent in the baseline HCMST table. We retain the classified HCMST sample as the baseline because the reallocation assumes that uncertain reports have the same composition across the two data sources. HCMST also provides a larger usable sample, 485 couples rather than 222, while the ANES partner study requires both partners to participate.

Table 3 reports the full national couple distribution. 74.4 percent of couples have the same education, and the four same-party, same-education cells account for 62.3 percent of all couples.

Table 3: Couples by female and male type (%, HCMST 2017)

	Female type				
Male type	D, BA	D, NoBA	R, BA	R, NoBA	Total
D, BA	17.1	3.8	2.1	0.0	23.1
D, NoBA	8.6	13.0	0.7	2.5	24.8
R, BA	5.3	0.5	13.7	4.9	24.5
R, NoBA	1.8	3.1	5.3	17.5	27.7
Total	32.9	20.5	21.7	24.9	100.0

Notes. Cells give shares using the HCMST survey weight multiplied by the ACS-based probability that the partner is aged 25–45. D/R include party identifiers and leaners; BA/NoBA indicate partners with and without a bachelor's degree. The sample contains $N = 485$ married or cohabiting opposite-sex couples whose respondent belongs to HCMST's 25–44 age group.

6.3 Constructing the Equilibrium Population

The CES observes whether a respondent is partnered but not the partner's age, sex, or party. We therefore combine all four sources to divide each region-by-sex-by-type cell into the five states in equation (1): single and seeking only same-sex partners; single and open to a woman–man partnership; in an existing same-sex partnership; in a modeled woman–man partnership between Democrats or Republicans aged 25–45; or in an opposite-sex partnership outside that market.

Among adults classified as single, HCMST asks whether they are attracted to men, women, or both. We use the responses to distinguish people who report attraction only to the same sex from those who are open to a woman–man partnership. Applying the CES age, party, and partnership-status restrictions, we estimate one share for men and one for women. Among 142 men and 168 women with valid attraction reports, 27 men and 17 women report attraction only to the same sex. The survey-weighted shares are 6.4 percent for men and 3.1 percent for women. The detailed cells are small, so we apply these sex-specific shares uniformly across party, education, and region. Singles open to both women and men remain eligible for the woman–man equilibrium. A sensitivity analysis uses HCMST sexual identity instead of attraction.

ACS partner links separate existing same-sex partnerships from opposite-sex partnerships. The 2017–2021 ACS identifies both members of each co-residential married or unmarried partnership, so we observe each person's sex, exact age, and education. Among linked partnered adults aged 25–45, 2.04 percent have a same-sex partner. We estimate this share separately by region, sex, and education. Because the ACS does not observe party, we apply each rate equally to Democrats and Republicans within the same demographic cell.

We next determine which remaining opposite-sex partnerships belong to the modeled market. The ACS supplies the probability that the partner is also aged 25–45. HCMST supplies the partner's party and education, and ANES supplies the classification of uncertain partner-party reports described above. Leaning Independents are classified as Democrats or Republicans throughout. Adults who do not lean toward either party are not part of the model population. A Democrat or Republican partnered with a non-leaning Independent enters the outside-partnership group. We reconcile the resulting female and male totals within each region, preserving the estimated type composition within sex. The remaining partnered mass forms the outside-partnership group, which includes opposite-sex partnerships with non-leaning Independents and with people outside the age range. The modeled partnership mass never exceeds the available opposite-sex partnered mass in any cell.

Only singles open to a woman–man partnership, adults in modeled partnerships, and adults in opposite-sex partnerships outside the modeled D/R market enter the equilibrium. They may move among these three outcomes in the counterfactuals. Singles seeking only same-sex partners and adults in existing same-sex partnerships remain fixed. Nationally, the equilibrium contains 96.6 percent of Democratic and Republican adults aged 25–45; the two fixed groups account for 2.36 and 1.07 percent, respectively. Reported full-population singlehood adds the fixed single group after solving the equilibrium. Full-population welfare assigns zero change to both fixed groups, while cross-party and same-education shares remain conditional on modeled partnerships.

7 Identification, Estimation, and Validation

The matching model combines the experimental preference estimates with observed partnership patterns. The experiment determines how party and education enter match surplus. Regional population shares determine which partner types are available, and the meeting pools determine whom people can encounter. Within the woman–man equilibrium, adults may remain single, form a partnership within the modeled D/R market, or form an opposite-sex partnership outside it—for example, with a non-leaning Independent or someone outside ages 25–45. The regional HCMST couple tables identify the scale of the experimental surplus and the degree of education-based separation in meeting opportunities. CES, ACS, HCMST, and ANES together determine the baseline population in each relationship state.

7.1 Data, Parameters, and Estimation

Three sets of parameters govern different parts of matching. The conversion factor λ determines how strongly the party and education differences measured in the experiment affect who partners with whom. The common own-pool share π determines how much more often people encounter potential partners with the same education. Section 7.3 relaxes this restriction by giving each party-by-education type its own assignment probability. The regional type intercepts α_i^r for women and γ_j^r for men shift the value of every potential match involving a type and therefore affect how often that type forms a modeled partnership.

For each region and type, we set the payoff from an opposite-sex partnership outside the modeled D/R market so that the baseline model reproduces its observed frequency relative to singlehood. Under the model's normalization, this payoff equals the log ratio of the two population shares. We hold the payoff fixed in the counterfactuals, while the shares choosing singlehood, an outside partnership, and a modeled partnership can change.

We estimate the remaining parameters in a nested procedure. At each candidate pair (λ, π), we choose the regional type intercepts to reproduce the baseline singlehood rates among people who enter the woman–man equilibrium. Given the calibrated outside payoff, this also reproduces each type's baseline outside-partnership rate. We then solve the regional equilibria and compare the resulting pattern of partner types with HCMST.

The comparison uses the association between female and male types rather than differences in the marginal number of each type observed in CES and HCMST. Female- and male-type fixed effects in every region align the fitted couple margins with HCMST. Let C^r be the matrix of observed HCMST couple counts in region r, $\widetilde{C}^r(\lambda, \pi)$ the fitted matrix after aligning those margins, and D their Poisson deviance. Following Galichon and Salanié (2024), we estimate

$$(\hat{\lambda}, \hat{\pi}) = \arg\min_{\lambda,\pi} \sum_r D\Big(C^r, \widetilde{C}^r(\lambda, \pi)\Big)\,. \tag{22}$$

Appendix E.5 gives the fixed-effect calculation and the complete objective. Appendix E.7 derives the one-pool benchmark at $\pi = 1/2$.

Uncertainty intervals repeat the full procedure on resampled inputs. Each replication draws the experimental preferences and choice-error scale jointly from the log-income posterior and resamples CES respondents, linked ACS couples, HCMST respondents, and ANES dyads. The same HCMST respondent weight is used wherever that respondent contributes to the attraction measure or the couple table. All 1,000 replications completed successfully. Appendix E.6 describes the resampling procedure and numerical checks.

7.2 Parameter Estimates and Model Fit

The conversion factor is $\hat{\lambda} = 0.161$, with a 95 percent interval of $[0.131, 0.207]$. The own-pool share is $\hat{\pi} = 0.917$ $[0.893, 0.939]$, so an estimated 91.7 percent of each education group is assigned to its education-centered pool. Party sorting primarily identifies λ: the experiment finds large party differences in surplus, and the observed cross-party share determines how strongly those differences must enter the marriage market. Education sorting primarily identifies π: the experiment finds much weaker education preferences than the degree of education sorting observed among couples.

Table 4: Matching-model estimates and couple sorting

Panel A. Parameter estimates		
Parameter	Estimate	95% interval
Conversion factor λ	0.161	$[0.131, 0.207]$
Own-pool share π	0.917	$[0.893, 0.939]$
Panel B. National couple sorting (%)		
	Observed	Fitted model
Cross-party	16.0	16.2
Same education	74.4	74.4

Notes. Panel A reports the joint estimates using the 423-respondent log-income experimental posterior, regional population inputs, and 485 HCMST couples. Intervals use 1,000 joint bootstrap replications. Panel B compares HCMST with the fitted matrices $\widetilde{C}^r$ in equation (22). The fitted shares use the HCMST female- and male-type margins and therefore evaluate sorting across partner types rather than the marginal number of each type who partners.

Figure 7: Joint uncertainty in the conversion factor and meeting-pool parameter

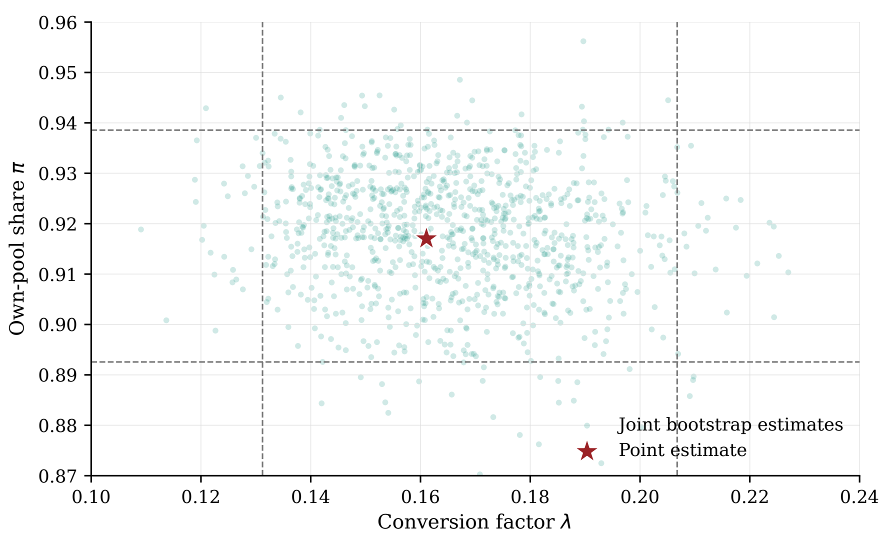

Notes. The star marks the joint point estimate. Dots are estimates from the 1,000 joint bootstrap replications; dashed lines mark the parameter-specific 95 percent intervals. Each replication redraws the experimental posterior and the CES, ACS, HCMST, and ANES inputs and re-estimates both parameters.

The model reproduces the two national sorting moments and fits the full regional couple tables

(Figure 8). The observed deviance is 40.514. In a fitted-reference check, we repeatedly draw regional couple tables from the fitted model, re-estimate λ and π, and compare each minimized deviance with the observed value. The simulated deviance is at least as large as the observed deviance with probability $p = 0.589$. This check finds no evidence of a general failure to reproduce the regional association patterns.

Figure 8: Model fit across census regions

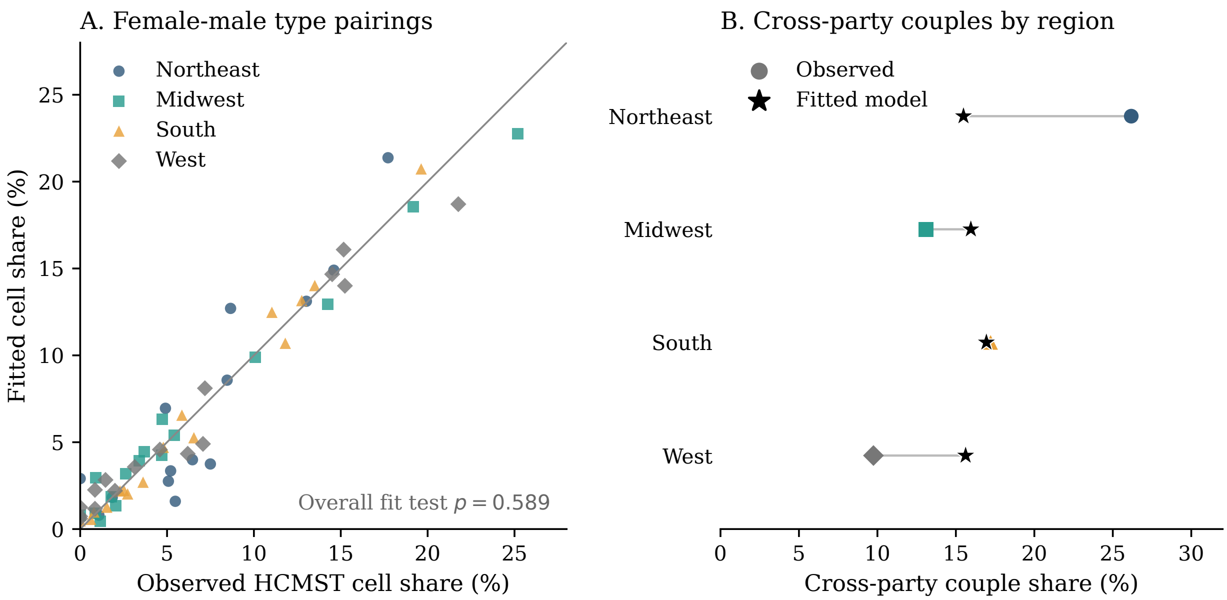

Notes. Panel A compares observed HCMST shares with the fitted female–male type-pairing shares in each region. Panel B compares observed and fitted cross-party shares. The fitted tables use the HCMST female- and male-type margins. The overall fitted-reference test re-estimates λ and π in each simulated sample.

Because baseline singlehood and outside-partnership shares determine the type intercepts and outside payoffs, the fit test concerns association across partner types rather than those relationship-state shares.

7.3 Contact and Surplus Restrictions

The headline specification uses one common π rather than four type-specific assignment probabilities. Relaxing that restriction recovers almost the same two-pool structure: assignment differs by 79.8 percentage points across education groups and only 0.6 point across parties. The conversion factor remains 0.161, and the three additional assignment parameters lower deviance by only 0.701, from 40.514 to 39.813. The unrestricted model's AIC is 5.30 points higher. We therefore use the one-parameter education-centered assignment as the headline and the four-parameter model as a test of that restriction.

A third unrestricted pool improves deviance by 3.441 and lowers the conversion factor to 0.128. Because the additional pool is not identified under the two-pool model, we use a fitted-reference bootstrap rather than a conventional chi-squared comparison. The resulting value, $p = 0.239$, does not reject two pools in favor of three. The three-pool specification nevertheless provides a consequential sensitivity: under complete removal of partisan aversion, predicted singlehood is 42.8 rather than 39.9 percent, and the local partner-income equivalent at a $62,500 profile income is $10.9 thousand rather than $13.6 thousand per adult.

Raw deviance continues to fall as more pools are added, but each additional pool consumes four residual degrees of freedom, and AIC favors the education-based two-pool specification (Table 5). With four pools, the assignment probabilities can make contact nearly type-specific and drive the conversion factor to its zero boundary, allowing contact to absorb almost all of the observed sorting. We use the four-pool fit only to illustrate how flexible contact can account for sorting. Its boundary estimate cannot establish that measured preferences are unimportant.

Table 5: Fit and parsimony across meeting-pool specifications

Model	Assignment parameters	Deviance	Residual df	ΔAIC
One pool	0	139.93	35	97.41
Education-based two pools	1	40.514	34	0.00
Unrestricted two pools	4	39.813	31	5.30
Unrestricted three pools	8	36.373	27	9.86
Unrestricted four pools	12	31.73	23	13.22

Notes. Each specification is fitted to 64 female–male type cells across four regional HCMST tables. Residual degrees of freedom subtract 28 regional female- and male-type fixed effects, one conversion factor, and the reported assignment parameters. ΔAIC is relative to the education-based two-pool model. The fitted-reference tests reported in the text are preferred to the asymptotic chi-squared approximation. The four-pool fit illustrates the flexibility of contact. The search does not establish a global optimum.

The data reject a common conversion factor across regions. Allowing separate factors yields 0.096 in the Northeast, 0.181 in the Midwest, 0.157 in the South, and 0.219 in the West. Deviance falls by 10.819, and a fitted-reference test of the common-factor restriction gives $p = 0.023$. These factors rescale the same nationally estimated preference matrix and therefore do not identify regional preferences. The regional fit does not tell us whether the choice-task error scale or the scale of unobserved match values differs across regions. Without separate evidence on either scale, we use one conversion factor to keep the national model simple and report the regional estimates as a sensitivity.

7.4 External Evidence

The external evidence assesses different parts of the model. The North Carolina voter file provides a held-out comparison of party sorting, while residential and workplace composition provide separate evidence on the education-based meeting structure.

We use North Carolina because its public statewide voter file records party registration and residential address for individuals, and both major parties are well represented among registrants (North Carolina State Board of Elections, 2026). The file is not used to estimate the model. Using the national preference, conversion-factor, and meeting-pool estimates with North Carolina population inputs, the model predicts that 15.39 percent of couples are cross-party, compared with

17.52 percent in the voter file. The correlation across the four party-pairing cells is 0.990, with an RMSE of 2.55 percentage points. The voter file does not record education, includes adults outside the model's age range, and infers partnerships from shared addresses. North Carolina singlehood is a local calibration input, so the held-out comparison concerns party sorting among couples rather than the partnership rate.

The ACS provides separate evidence on education-based meeting opportunities. Education sorting is stronger where adults with and without a BA are more residentially separated. We measure sorting as the observed same-education couple share minus the share expected under random matching given each place's education composition. A one-standard-deviation increase in residential separation is associated with a 1.7 percentage point increase in education sorting across 51 states and a 3.1 percentage point increase across 660 commuting zones; the corresponding t-statistics are 7.9 and 24.5. These associations are consistent with education-segmented meeting, although they may also reflect differences in preferences or other local characteristics.

Workplace composition points in the same direction. In 2020, workers without a college degree were employed at establishments that were 76.7 percent non-college, on average, whereas college graduates were employed at establishments that were about 38.0 percent non-college—a 38.7 percentage point difference in exposure to non-college coworkers (Dillon et al., 2025). The analogous partisan difference in coworker composition is 11.7 percentage points (Frake et al., 2026). The studies use different worker samples, and workplaces are only one setting in which potential partners meet. Their relative magnitudes therefore support the education-centered structure without identifying the model's meeting pools.

7.5 Counterfactual Design

The historical counterfactual asks how today's marriage market would differ if partisan aversion returned to its mid-1990s level. We also vary party composition to assess its separate contribution. Both exercises retain the estimated meeting structure. Adults in the woman–man equilibrium may remain single, partner with a Democratic or Republican adult of the opposite sex aged 25–45, or form an opposite-sex partnership outside that group. The payoff from the last option remains fixed, so the share choosing it may change. Existing same-sex partnerships and people seeking only same-sex partners remain fixed.

We use the pooled 1992–2000 ANES as the mid-1990s benchmark and the pooled 2016–2024 ANES as the current benchmark (ANES, 2026). Pooling several studies avoids relying on noisy single-year sex-by-education cells. The aversion counterfactual reduces every type's cross-party penalty by 40 percent, the rounded proportional reduction that returns the current ANES in-party minus out-party feeling-thermometer gap to its mid-1990s level. We assume that partner-specific partisan aversion changes in proportion to political affect. The exercise does not estimate the effect of an intervention.

For composition, we calculate the historical-minus-current change in Democratic log odds within each sex-by-education group and apply that national change to every CES region. The transformation preserves each region's population and sex-by-education totals. Each type also carries forward its baseline shares in singlehood, another opposite-sex partnership, an existing same-sex partnership, a same-party modeled partnership, and a cross-party modeled partnership before the new equilibrium is solved.

The four equilibria are the current baseline, mid-1990s aversion only, mid-1990s composition only, and both changes together. Comparing them separates the preference and composition channels, and their interaction, under a common meeting structure. Holding that structure fixed is the maintained restriction that permits the historical accounting exercise. The restriction is particularly plausible for the political dimension: when pool assignment is allowed to differ by party, the fitted Democratic–Republican difference within each education group is less than one percentage point, and the counterfactual responses remain close to those in the baseline model (Appendix E.8). This does not require schools, workplaces, or dating technology to have remained unchanged. It requires changes in those settings not to materially alter the estimated partnership response to partisan aversion and composition.

A separate benchmark lowers the estimated own-pool share π to $1/2$, making the two pools equivalent to one common meeting pool. This exercise shows the role of education-segmented meeting in the model but is not assigned a historical date. We also report the endpoint that removes the cross-party penalty completely.

Singlehood and welfare are reported for the full Democratic and Republican population aged 25–45, assigning zero modeled change to the two fixed same-sex groups. Cross-party and same-education shares remain conditional on modeled woman–man partnerships. Welfare is reported in the log-partner-income units defined in Section 2.3 and as a local partner-income equivalent at a stated reference income.

8 Counterfactual Effects

Returning partisan aversion to its mid-1990s benchmark increases cross-party matching and partnership formation. We also change party composition to its historical benchmark; it has little effect on the overall partnership rate. Both exercises hold today's meeting opportunities and other determinants of partnership fixed.

8.1 Mid-1990s Partisan Aversion and Composition

We use 1992–2000 as the historical window and refer to it as the mid-1990s benchmark. This period precedes much of the subsequent rise in out-party animosity (Iyengar et al., 2019). Pooling the five ANES studies centered on 1996 makes the sex-by-education composition estimates less sensitive to sampling variation than a single-year benchmark. The comparison window pools the 2016, 2020, and 2024 ANES studies (ANES, 2026).

The aversion change reduces each type's estimated cross-party penalty by 40 percent. The pooled ANES in-party minus out-party feeling-thermometer gap is 27.1 in the mid-1990s window and 44.9 in the current window. Returning the current gap to its earlier level implies a 39.6 percent proportional reduction, which we round to 40 percent. The mapping asks how the current marriage market would differ if partner-specific partisan aversion changed in the same proportion as this historical measure of political affect. It does not interpret the feeling-thermometer gap itself as a causal treatment.

The composition change applies the historical change in Democratic log odds within each sex-by-education group to every CES region. It preserves regional population and sex-by-education totals while changing the Democratic and Republican shares within those cells. Panel B of Figure 9 shows the resulting national shares. The counterfactual values are not raw historical ANES shares; they are current CES shares after applying the corresponding historical ANES log-odds shift.

Figure 9: Historical Inputs to the Counterfactual

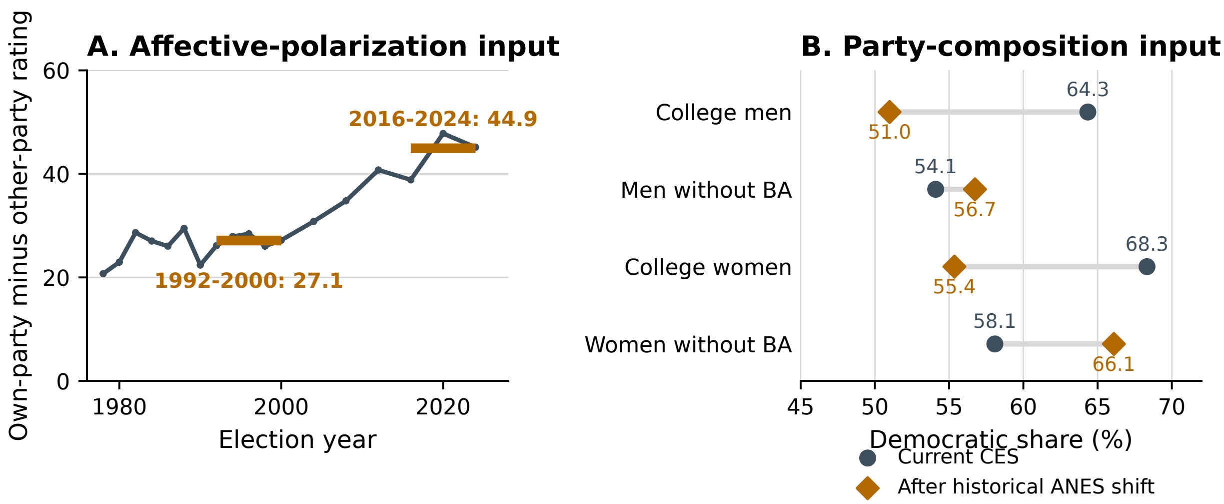

Notes. Panel A shows the weighted ANES in-party minus out-party feeling-thermometer gap among Democratic and Republican identifiers and leaners aged 25–45. Horizontal segments show the pooled means for 1992–2000 and for 2016, 2020, and 2024. Panel B shows current national CES Democratic shares and the counterfactual shares obtained by applying the pooled ANES historical-minus-current change in Democratic log odds within each sex-by-education group to every CES region. Non-leaning Independents are excluded. Sources: CES and ANES (Kuriwaki, 2022; ANES, 2026).

We hold the education-based meeting structure fixed in the historical counterfactual. The data support this restriction on the political dimension: when pool assignment can vary by party, fitted Democratic–Republican differences within education groups are less than one percentage point, and the response to lower partisan aversion changes little (Appendix E.8). Thus, the couple data do not require partisan meeting pools to explain party sorting. Schools, workplaces, neighborhoods, and dating technology may have changed over time; we assume those changes do not materially alter the estimated effects of aversion and party composition.

Table 6: Mid-1990s Partisan Aversion and Composition in the CES Partnership Measure

	Cross-party couples (%)	Same-education couples (%)	Singlehood (%)	Welfare equivalent (%)
Current baseline	15.41	75.77	48.58	—
Mid-1990s aversion	26.11	76.29	46.33	5.3
Mid-1990s composition	15.92	75.59	48.56	-0.4
Both changes	26.83	76.16	46.25	5.0

Notes. Entries use CES reports of marriage or domestic partnership; the CES item misses some cohabiting couples. The aversion exercise removes 40 percent of every type's cross-party penalty. The composition exercise applies the national 1992–2000 minus 2016–2024 ANES change in Democratic log odds within sex-by-education groups to each CES region. All fitted parameters, non-party surplus, regional population and sex-by-education totals, meeting opportunities, and outside-partnership payoffs remain fixed. Existing same-sex partnerships and people seeking only same-sex partners receive zero modeled change. The last column reports the welfare change in log-partner-income units as an approximate percentage, averaged over the full Democratic and Republican population. The historical partnership-rate comparison below also counts cohabitation, using ACS rates (Appendix F).

8.2 Preference and Composition Channels

When both inputs are set to their mid-1990s benchmarks, the cross-party share of modeled couples rises from 15.41 to 26.83 percent and singlehood falls from 48.58 to 46.25 percent. The welfare change is approximately 5.0 percent of partner income. At a reference partner income of $62,500, the corresponding local dollar equivalent is $3.12 thousand per adult.

Within the model, the 2.33-point increase in partnership separates into 2.25 percentage points from lower aversion, 0.02 point from the change in party composition, and 0.06 point from their interaction. Lower aversion therefore accounts for 96.5 percent of the modeled increase. Mid-1990s composition slightly raises cross-party matching but has almost no effect on aggregate partnership formation.

For Table 6, 1,000 bootstrap replications give a median aversion-only reduction in singlehood of 2.235 points (95 percent interval [2.107, 2.376]). This interval excludes uncertainty in the historical composition adjustment; Appendix Figure 19 shows the full response curve.

For comparison with the historical decline in marriage and cohabitation, we fit the model using ACS partnership rates that include cohabiting couples (Appendix F). A 40 percent reduction in partisan aversion then raises the partnered share by 2.04 points among Democrats and Republicans. Expressed over all adults using the Democratic and Republican share of the CES population, the increase is 1.51 points. The 95 percent interval for this all-adult effect is [1.41, 1.59] points, conditional on the aversion change. In the CPS ASEC, the married-or-cohabiting share of adults aged 25–45 falls from 66.75 percent in 1995 to 61.67 percent in 2025, a 5.08-point decline (Flood et al., 2023). The modeled aversion response equals 29.8 percent of that decline. This calculation holds current composition and meeting opportunities fixed; it does not establish how much of the historical decline aversion caused. The CPS relationship measure misses some cohabiting couples in both periods.

8.3 Model Benchmarks

Three additional benchmarks use the CES partnership measure in Table 6 and are not assigned a historical date. Removing the cross-party penalty completely raises the cross-party share to 48.87 percent and lowers singlehood to 40.03 percent. This endpoint measures the full contribution of the estimated penalty; the historically grounded exercise above is the paper's main counterfactual.

Combining the two education-centered meeting pools leaves party sorting nearly unchanged but sharply reduces education sorting. The cross-party share moves from 15.41 to 15.44 percent, while the same-education share falls from 75.77 to 53.64 percent and singlehood falls to 42.47 percent. This benchmark shows that education-segmented meeting is consequential for education sorting and partnership formation but does not provide evidence of independently partisan meeting segmentation.

Equalizing college composition provides a separate benchmark for the education imbalance between women and men. Within each region, we move a common fraction of Democratic and Republican men without a BA into the corresponding BA types until the number of college-educated men equals the number of college-educated women. The required fraction ranges from 12.3 to 19.4 percent across regions. Holding partner preferences, meeting segmentation, type-specific outside payoffs, and the remaining structural parameters fixed, singlehood falls from 48.58 to 47.50 percent, a 1.08-point reduction with a 95 percent interval of [0.93, 1.22]. Cross-party and same-education couple shares change by only 0.26 and 0.18 percentage point. The exercise changes type masses and assigns the added mass the calibrated characteristics of existing college-educated men. It does not estimate how schooling would change an individual's preferences or partnership prospects (Appendix E.3).

9 Conclusion

Observed sorting in the marriage market can reflect preferences for similar partners, differences in who is available, or limited opportunities to meet. These forces operate differently across dimensions. In our model, experimentally measured partisan aversion accounts for party sorting, whereas education-centered meeting opportunities account for the education sorting left unexplained by weak education preferences. Differences in the relative numbers of Democrats and Republicans have little effect on partnership formation.

Partisan aversion also affects whether partnerships form. Returning aversion to its mid-1990s benchmark raises the married-or-cohabiting share of adults aged 25–45 by 1.51 percentage points, equivalent to 29.8 percent of its measured 1995–2025 decline. This comparison holds today's meeting opportunities fixed and leaves other causes of the historical decline outside the exercise.

Preferences and meeting opportunities may not evolve independently. Cross-party contact can reduce partisan hostility, yet aversion to cross-party interaction may limit the encounters that could weaken it (Braghieri et al., 2024). Partner choice may also carry political divisions across generations: parent–child partisan agreement is stronger when parents share a party (Iyengar et al., 2018). By shaping the political environment within households, assortative partnership may influence how strongly children identify with a party and how they view the opposing party. Studying these feedbacks would extend the analysis from the private costs of partisan aversion to its longer-run political consequences.

References

Alford, John R., Peter K. Hatemi, John R. Hibbing, Nicholas G. Martin, and Lindon J. Eaves, "The Politics of Mate Choice," *The Journal of Politics*, April 2011, *73* (2), 362–379.

Andrew, Alison and Abi Adams, "Revealed Beliefs and the Marriage Market Return to Education," *The Quarterly Journal of Economics*, July 2025, *140* (3), 2107–2162.

ANES, "ANES 2024 Time Series Study [dataset and documentation]," Stanford University and the University of Michigan 2025. August 8, 2025 version, including the Spouse and Partner Study module. Available at https://electionstudies.org.

_ , "ANES Time Series Cumulative Data File (1948-2024) [dataset and documentation]," Stanford University and the University of Michigan 2026. February 5, 2026 version. Available at https://electionstudies.org.

Becker, Gary S., "A Theory of Marriage: Part I," *Journal of Political Economy*, July 1973, *81* (4), 813–846.

Belot, Michèle and Marco Francesconi, "Dating Preferences and Meeting Opportunities in Mate Choice Decisions," *Journal of Human Resources*, 2013, *48* (2), 474–508.

Bertrand, Marianne, Emir Kamenica, and Jessica Pan, "Gender Identity and Relative Income within Households," *The Quarterly Journal of Economics*, May 2015, *130* (2), 571–614.

Braghieri, Luca, Peter Schwardmann, and Egon Tripodi, "Talking across the Aisle," Discussion Paper 19753, Centre for Economic Policy Research 2024. Revised June 2026.

Chiappori, Pierre-André, Bernard Salanié, and Yoram Weiss, "Partner Choice, Investment in Children, and the Marital College Premium," *American Economic Review*, August 2017, *107* (8), 2109–2167.

_ **, Monica Costa Dias, Costas Meghir, and Hanzhe Zhang**, "Changes in Marital Sorting: Theory and Evidence from the United States," *Journal of Political Economy*, October 2025, *133* (10), 3045–3076.

Choo, Eugene and Aloysius Siow, "Who Marries Whom and Why," *Journal of Political Economy*, February 2006, *114* (1), 175–201.

Ciscato, Edoardo, "Assessing Racial and Educational Segmentation in Large Marriage Markets," *The Review of Economic Studies*, 2025, *92* (6), 3788–3839.

Cox, Daniel A., Jacqueline Clemence, and Eleanor O'Neil, "Partisan Attachment: How Politics Is Changing Dating and Relationships in the Trump Era," Technical Report, Survey Center on American Life, American Enterprise Institute 2020. Findings from the January 2020 American Perspectives Survey.

Currarini, Sergio, Matthew O. Jackson, and Paolo Pin, "An Economic Model of Friendship: Homophily, Minorities, and Segregation," *Econometrica*, 2009, *77* (4), 1003–1045.

Dillon, Francis M., Edward L. Glaeser, and William R. Kerr, "Workplace Segregation between College and Non-college Workers," NBER Working Paper 33653, National Bureau of Economic Research April 2025.

Drake, Marshall, Fernando Payró, Neil Thakral, and Linh T. Tô, "Bayesian Adaptive

Choice Experiments," 2024. Working paper.

Duncan, Otis Dudley and Beverly Duncan, "A Methodological Analysis of Segregation Indexes," *American Sociological Review*, 1955, *20* (2), 210–217.

Easton, Matthew J. and John B. Holbein, "The Democracy of Dating: How Political Affiliations Shape Relationship Formation," *Journal of Experimental Political Science*, 2021, *8* (3), 260–272.

Eika, Lasse, Magne Mogstad, and Basit Zafar, "Educational Assortative Mating and Household Income Inequality," *Journal of Political Economy*, December 2019, *127* (6), 2795–2835.

Fernández, Raquel, Nezih Guner, and John Knowles, "Love and Money: A Theoretical and Empirical Analysis of Household Sorting and Inequality," *The Quarterly Journal of Economics*, 2005, *120* (1), 273–344.

Finkel, Eli J., Christopher A. Bail, Mina Cikara, Peter H. Ditto, Shanto Iyengar, Samara Klar, Lilliana Mason, Mary C. McGrath, Brendan Nyhan, David G. Rand, Linda J. Skitka, Joshua A. Tucker, Jay J. Van Bavel, Cynthia S. Wang, and James N. Druckman, "Political sectarianism in America," *Science*, October 2020, *370* (6516), 533–536.

Fisman, R., S. S. Iyengar, E. Kamenica, and I. Simonson, "Gender Differences in Mate Selection: Evidence From a Speed Dating Experiment," *The Quarterly Journal of Economics*, May 2006, *121* (2), 673–697.

Fisman, Raymond, Sheena S. Iyengar, Emir Kamenica, and Itamar Simonson, "Racial Preferences in Dating," *Review of Economic Studies*, January 2008, *75* (1), 117–132.

Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, J. Robert Warren, Daniel Warren, Annie Chen, Grace Cooper, Stephanie Richards, Megan Schouweiler, and Michael Westberry, "IPUMS, Current Population Survey: Version 11.0," 2023.

Frake, Justin, Reuben Hurst, and Max Kagan, "Political Segregation in the US Workplace," *Nature Human Behaviour*, 2026, *10* (8), 1455–1473.

Galichon, Alfred and Bernard Salanié, "Cupid's Invisible Hand: Social Surplus and Identification in Matching Models," *The Review of Economic Studies*, September 2022, *89* (5), 2600–2629.

_ **and** _ , "Estimating separable matching models," *Journal of Applied Econometrics*, September 2024, *39* (6), 1021–1044.

Goldman, Benjamin, Jamie Gracie, and Sonya Porter, "Who Marries Whom? The Role of Segregation by Race and Class," *American Economic Review*, 2026. Forthcoming.

Greenwood, Jeremy, Nezih Guner, Georgi Kocharkov, and Cezar Santos, "Marry Your Like: Assortative Mating and Income Inequality," *American Economic Review*, May 2014, *104* (5), 348–353.

Gualdani, Cristina and Shruti Sinha, "Partial Identification in Matching Models for the Marriage Market," *Journal of Political Economy*, May 2023, *131* (5), 1109–1171.

Hersh, Eitan and Yair Ghitza, "Mixed partisan households and electoral participation in the United States," *PLOS ONE*, October 2018, *13* (10), e0203997.

Hitsch, Günter J, Alı Hortaçsu, and Dan Ariely, "Matching and Sorting in Online Dating," *American Economic Review*, March 2010, *100* (1), 130–163.

Hoffman, Matthew D. and Andrew Gelman, "The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo," *Journal of Machine Learning Research*, 2014, *15*, 1593–1623.

Huber, Gregory A. and Neil Malhotra, "Political Homophily in Social Relationships: Evidence from Online Dating Behavior," *The Journal of Politics*, January 2017, *79* (1), 269–283.

Hudde, Ansgar and Daniela Grunow, "Why Do Partners Often Prefer the Same Political Parties? Evidence from Couples in Germany," *Social Forces*, 2025, *103* (4), 1581–1603.

Iyengar, Shanto, Tobias Konitzer, and Kent Tedin, "The Home as a Political Fortress: Family Agreement in an Era of Polarization," *The Journal of Politics*, October 2018, *80* (4), 1326–1338.

_ , Yphtach Lelkes, Matthew Levendusky, Neil Malhotra, and Sean J. Westwood, "The Origins and Consequences of Affective Polarization in the United States," *Annual Review of Political Science*, May 2019, *22* (1), 129–146.

Kalmijn, Matthijs, "Intermarriage and Homogamy: Causes, Patterns, Trends," *Annual Review of Sociology*, August 1998, *24* (1), 395–421.

Kuriwaki, Shiro, "Cumulative CES Common Content," 2022. Harvard Dataverse, doi:10.7910/DVN/II2DB6.

Levendusky, Matthew S. and Neil Malhotra, "(Mis)perceptions of Partisan Polarization in the American Public," *Public Opinion Quarterly*, 2016, *80* (S1), 378–391.

Low, Corinne, "Pricing the Biological Clock: The Marriage Market Costs of Aging to Women," *Journal of Labor Economics*, April 2024, *42* (2), 395–426.

Manski, Charles F., "The Structure of Random Utility Models," *Theory and Decision*, 1977, *8* (3), 229–254.

Mare, Robert D., "Five Decades of Educational Assortative Mating," *American Sociological Review*, February 1991, *56* (1), 15.

McFadden, Daniel, "Conditional Logit Analysis of Qualitative Choice Behavior," in Paul Zarembka, ed., *Frontiers in Econometrics*, New York: Academic Press, 1974, pp. 105–142.

_ , "Modelling the Choice of Residential Location," in Anders Karlqvist, Lars Lundqvist, Folke Snickars, and Jörgen W. Weibull, eds., *Spatial Interaction Theory and Planning Models*, Amsterdam: North-Holland, 1978, pp. 75–96.

Nicholson, Stephen P., Chelsea M. Coe, Jason Emory, and Anna V. Song, "The Politics of Beauty: The Effects of Partisan Bias on Physical Attractiveness," *Political Behavior*, 2016, *38*, 883–898.

North Carolina State Board of Elections, "Voter Registration Data," 2026.

Phan, Du, Neeraj Pradhan, and Martin Jankowiak, "Composable Effects for Flexible and Accelerated Probabilistic Programming in NumPyro," *arXiv preprint arXiv:1912.11554*, 2019.

Rosenfeld, Michael J., Reuben J. Thomas, and Sonia Hausen, "How Couples Meet and Stay Together, Waves 1–7," Stanford University Libraries, Stanford Digital Repository 2023. Data set.

Ruggles, Steven, Sarah Flood, Matthew Sobek, Daniel Backman, Annie Chen, Grace

Cooper, Stephanie Richards, Renae Rodgers, and Megan Schouweiler, "IPUMS USA: Version 15.0," 2024.

Schwartz, Christine R. and Robert D. Mare, "Trends in educational assortative marriage from 1940 to 2003," *Demography*, November 2005, *42* (4), 621–646.

Shapley, L S and M Shubik, "The assignment game I: The core," *International Journal of Game Theory*, 1971, *Volume 1*, 111–130.

Sleiman, Yara, Georgios Melios, and Paul Dolan, ""Sleeping with the Enemy": Partisanship and Tolerance in Online Dating," *Political Science Research and Methods*, 2025, pp. 1–23.

Taflinger, Shannon and Ansgar Hudde, "Why Do Young US Americans Avoid Cross-Partisan Dating? A Closer Look at Mediators and Variation by Gender and Party," *European Sociological Review*, 2026, p. jcag020.

Train, Kenneth E., *Discrete Choice Methods with Simulation*, 2nd ed., Cambridge University Press, 2009.

_ **and Melvyn Weeks**, "Discrete Choice Models in Preference Space and Willingness-to-Pay Space," in Riccardo Scarpa and Anna Alberini, eds., *Applications of Simulation Methods in Environmental and Resource Economics*, Dordrecht: Springer, 2005, pp. 1–16.

Understanding America Study, "UAS 306: Election Tracking Poll Baseline and Consent 2020 [dataset and documentation]," Center for Economic and Social Research, University of Southern California 2020. Available at https://uasdata.usc.edu.

Appendix

A Matching Model Details

This appendix provides the derivations and computational details for Section 2. It derives the gravity equation from individual stability, shows how the meeting-pool model nests the standard one-pool model, explains the surplus parameterization and welfare measure, and describes computation (Choo and Siow, 2006; Galichon and Salanié, 2022).

A.1 From individual matching to partner-type choices

Section 2.2 represents matching as a choice among partner types. This subsection shows why the stable matching problem between individuals can be written that way. The derivation begins with individual match surplus and the conditions for stability. Transferable utility expresses the payoff available from each potential partner, and separability allows those payoffs to be grouped by partner type. The resulting systematic payoffs are U_{ij} for women and V_{ij} for men. Individual tastes determine the choice among these type-level payoffs, singlehood, and outside partnership.

For this derivation, fix a region r and a meeting pool h. We suppress these superscripts in this subsection and the next: write Φ_{ij} for Φ_{ij}^r, and write U_{ij} and V_{ij} for U_{ij}^{rh} and V_{ij}^{rh}.

At the individual level, joint surplus consists of systematic surplus for the type pairing and the two partners' idiosyncratic tastes. Index a woman by ℓ and a man by ℓ', and write $i(\ell)$ and $j(\ell')$ for their types. A match between a type-i woman and a type-j man produces

$$\widetilde{\Phi}_{\ell\ell'} = \Phi_{ij} + \varepsilon_{j\ell} + \eta_{i\ell'}, \tag{23}$$

where Φ_{ij} is systematic joint surplus, $\varepsilon_{j\ell}$ is the woman's taste for a type-j partner, and $\eta_{i\ell'}$ is the man's taste for a type-i partner. Singlehood pays $\varepsilon_{0\ell}$ to the woman and $\eta_{0\ell'}$ to the man. Equation (23) imposes separability: unobserved surplus equals the woman's taste for the man's type plus the man's taste for the woman's type. People of the same observed type may value partner types differently, but there is no additional shock unique to the pair.

A stable equilibrium rules out gains from choosing singlehood or outside partnership and from forming a new modeled pair. Let $\widetilde{u}_\ell$ and $\widetilde{v}_{\ell'}$ denote equilibrium payoffs, including transfers:

$$\widetilde{u}_\ell \geq \varepsilon_{0\ell}, \qquad \text{with equality if } \ell \text{ is single,} \tag{24}$$

$$\widetilde{v}_{\ell'} \geq \eta_{0\ell'}, \qquad \text{with equality if } \ell' \text{ is single,} \tag{25}$$

$$\widetilde{u}_\ell + \widetilde{v}_{\ell'} \geq \widetilde{\Phi}_{\ell\ell'}, \qquad \text{with equality if } \ell \text{ and } \ell' \text{ are matched.} \tag{26}$$

The first two inequalities compare each person's equilibrium payoff with singlehood; equality holds for people who are single. Stability also requires $\widetilde{u}_\ell \geq \zeta_{fi} + \varepsilon_{o\ell}$ and $\widetilde{v}_{\ell'} \geq \zeta_{mj} + \eta_{o\ell'}$, with equality for people choosing an outside partnership. The third displayed inequality is the no-blocking condition for a modeled pair. A potential pair cannot create more surplus than its members currently receive in total, because they could otherwise match with each other and divide the gain. Equality holds for an actual modeled match because the partners' payoffs exhaust its surplus.

These stability conditions express each person's equilibrium payoff as the best payoff available from singlehood, outside partnership, or a modeled match. Consider a woman ℓ of type i and hold the men's equilibrium payoffs fixed. A match with man ℓ' creates surplus $\widetilde{\Phi}_{\ell\ell'}$. To make him as well off as in his current equilibrium outcome, she must leave him at least $\widetilde{v}_{\ell'}$. Transferable utility lets her keep the remainder, so the largest payoff she can obtain from this match is $\widetilde{\Phi}_{\ell\ell'} - \widetilde{v}_{\ell'}$. Her equilibrium payoff is therefore

$$\widetilde{u}_\ell = \max\Big\{\varepsilon_{0\ell},\ \zeta_{fi} + \varepsilon_{o\ell},\ \max_{\ell'}\big(\widetilde{\Phi}_{\ell\ell'} - \widetilde{v}_{\ell'}\big)\Big\}. \tag{27}$$

We next group the potential men in equation (27) by type. For type-j men, substituting equation (23) gives

$$\max_{\ell' : j(\ell')=j}\big(\widetilde{\Phi}_{\ell\ell'} - \widetilde{v}_{\ell'}\big) = \Phi_{ij} + \varepsilon_{j\ell} + \max_{\ell' : j(\ell')=j}\big(\eta_{i\ell'} - \widetilde{v}_{\ell'}\big). \tag{28}$$

The last term in equation (28) determines how the systematic surplus is divided. For a type-j man ℓ', $\widetilde{v}_{\ell'} - \eta_{i\ell'}$ is the systematic payoff he must receive from a type-i woman to remain as well off as in his current equilibrium outcome. The woman can keep the most surplus by choosing the type-j man with the lowest such requirement. This minimum depends on the two types i and j, but not on the woman's identity ℓ, so it is common to all type-i women. Denote it by V_{ij}:

$$V_{ij} = \min_{\ell' : j(\ell')=j}\big(\widetilde{v}_{\ell'} - \eta_{i\ell'}\big) = -\max_{\ell' : j(\ell')=j}\big(\eta_{i\ell'} - \widetilde{v}_{\ell'}\big). \tag{29}$$

Subtracting this required payoff from systematic joint surplus leaves the woman's systematic payoff, U_{ij}:

$$U_{ij} = \Phi_{ij} - V_{ij}. \tag{30}$$

The best payoff available to woman ℓ among type-j men is therefore $U_{ij} + \varepsilon_{j\ell}$. By construction, $U_{ij} + V_{ij} = \Phi_{ij}$. The systematic payoff U_{ij} is common to all type-i women, while the taste draw allows those women to value type-j men differently.

The woman's maximization over individual men therefore becomes a maximization over partner types. Let $\mathcal{I}$ and $\mathcal{J}$ denote the sets of women's and men's types. Then

$$\widetilde{u}_\ell = \max_{j \in \mathcal{J} \cup \{0,o\}} \big(U_{ij} + \varepsilon_{j\ell}\big), \qquad U_{i0} = 0, \quad U_{io} = \zeta_{fi}. \tag{31}$$

The same argument gives the man's choice over female partner types:

$$\widetilde{v}_{\ell'} = \max_{i \in \mathcal{I} \cup \{0,o\}} \big(V_{ij} + \eta_{i\ell'}\big), \qquad V_{0j} = 0, \quad V_{oj} = \zeta_{mj}. \tag{32}$$

Equations (31) and (32) deliver the type-level choice problems used in equations (3) and (4). Singlehood has a taste draw but no systematic component, which gives the normalizations $U_{i0} = V_{0j} = 0$. In an (i,j) match, the partners' realized payoffs are $U_{ij} + \varepsilon_{j\ell}$ and $V_{ij} + \eta_{i\ell'}$. The next subsection converts these choices into matching shares and derives the gravity equation.

A.2 From type choices to the gravity equation

Under the Type-I extreme-value assumption in Section 2.2, equation (31) is a multinomial-logit choice problem. Let $\Pr_i(j)$ denote the probability that a type-i woman matches with type j, and let $\Pr_i(0)$ denote the probability that she remains single. With the taste scale normalized to one,

$$\Pr_i(j) = \frac{\exp(U_{ij})}{1 + \exp(\zeta_{fi}) + \sum_{k \in \mathcal{J}} \exp(U_{ik})}, \qquad \Pr_i(0) = \frac{1}{1 + \exp(\zeta_{fi}) + \sum_{k \in \mathcal{J}} \exp(U_{ik})}. \tag{33}$$

Let n_i and m_j denote the masses of type-i women and type-j men in the pool. Let μ_{ij} denote the mass of modeled (i,j) couples, μ_{i0} and μ_{0j} the single masses, and μ_{io} and μ_{oj} the masses in outside partnerships. The latter satisfy $\mu_{io} = e^{\zeta_{fi}} \mu_{i0}$ and $\mu_{oj} = e^{\zeta_{mj}} \mu_{0j}$. Feasibility requires

$$\mu_{i0} + \mu_{io} + \sum_j \mu_{ij} = n_i, \qquad \mu_{0j} + \mu_{oj} + \sum_i \mu_{ij} = m_j. \tag{34}$$

Because each type contains a continuum of people, choice probabilities equal population shares: $\mu_{ij} = n_i \Pr_i(j)$ and $\mu_{i0} = n_i \Pr_i(0)$. Taking their ratio gives

$$\frac{\mu_{ij}}{\mu_{i0}} = \exp(U_{ij}), \qquad \text{so} \qquad U_{ij} = \log\left(\frac{\mu_{ij}}{\mu_{i0}}\right). \tag{35}$$

The same argument for men yields

$$\frac{\mu_{ij}}{\mu_{0j}} = \exp(V_{ij}), \qquad \text{so} \qquad V_{ij} = \log\left(\frac{\mu_{ij}}{\mu_{0j}}\right). \tag{36}$$

Adding equations (35) and (36) and using $U_{ij} + V_{ij} = \Phi_{ij}$ gives

$$\Phi_{ij} = 2 \log \mu_{ij} - \log \mu_{i0} - \log \mu_{0j}. \tag{37}$$

This inversion recovers the systematic joint surplus of each type pairing from couples and singles within an observed pool, without observing the transfer between partners. Exponentiating gives

$$\mu_{ij} = \sqrt{\mu_{i0}\mu_{0j}}\ \exp\left(\frac{\Phi_{ij}}{2}\right). \tag{38}$$

The two pools share the same surplus Φ^r_{ij} but have separate equilibria. Restoring superscripts yields equation (9); U^{rh}_{ij}, V^{rh}_{ij}, and matching counts may differ between pools.

A.3 Meeting pools and nesting

Let pool 1 be college-centered and pool 2 non-college-centered. Let $q^h_t(\pi)$ denote the probability that a person of type t enters meeting pool h. For either women or men, the baseline assignment rule is

$$q^1_t(\pi) = \begin{cases} \pi, & E(t) = \text{BA}, \\ 1-\pi, & E(t) = \text{NoBA}, \end{cases} \qquad q^2_t(\pi) = 1 - q^1_t(\pi), \qquad \pi \in [1/2, 1]. \tag{39}$$

The function $E(t)$, defined in Section 2.1, gives type t's education component. Equation (39) assigns a share π of each education group to the pool centered on that group: college types enter pool 1 with probability π, and non-college types enter pool 2 with the same probability. The remaining share $1-\pi$ of each group enters the other pool. Because people form a continuum within each type, the assignment probabilities determine the pool masses exactly: $n^{rh}_i = q^h_i(\pi)n^r_i$ for women and $m^{rh}_j = q^h_j(\pi)m^r_j$ for men. Conditional on type, pool assignment is independent of the taste draws. Each pool solves equations (34) and (9) on its assigned masses using the common regional surplus matrix Φ^r. Regional couples and singles are the sums of the two pool-specific equilibria.

The matching system is homogeneous of degree one in the masses. If $(\mu_{ij}, \mu_{i0}, \mu_{0j})$ is the equilibrium for masses (n_i, m_j), then $(c\mu_{ij}, c\mu_{i0}, c\mu_{0j})$ is the equilibrium for masses (cn_i, cm_j) for any $c > 0$: the feasibility conditions and gravity equation both scale by c. Ratios of counts are unchanged, so match rates, singlehood rates, and the systematic payoffs U_{ij} and V_{ij} are unchanged as well.

Proposition 1 (Nesting)**.** *When $\pi = 1/2$, the sum of the two meeting-pool equilibria equals the equilibrium of a single unsegmented market containing the entire region.*

At $\pi = 1/2$, equation (39) assigns half of every female and male type to each pool. Both pools therefore have the region's type composition at half its scale. By degree-one homogeneity, each pool's equilibrium couples and singles are half of the one-pool equilibrium counts. Their sum is the one-pool equilibrium. When $\pi > 1/2$, each education group is more concentrated in its education-centered pool, so pool composition changes and the nesting argument no longer applies.

A.4 Surplus parameterization and identification

Because meeting-pool membership is unobserved, the data report regional totals rather than separate counts for each pool. Equation (37) recovers surplus from couples and singles within one pool, whereas the observed regional totals combine the two pool-specific equilibria:

$$\mu_{ij}^{r} = \mu_{ij}^{r1} + \mu_{ij}^{r2}, \qquad \mu_{i0}^{r} = \mu_{i0}^{r1} + \mu_{i0}^{r2}, \qquad \mu_{0j}^{r} = \mu_{0j}^{r1} + \mu_{0j}^{r2}.$$

When $\pi > 1/2$, the two pools have different type compositions, and each has its own matching equilibrium. Their couple and single masses therefore differ even though the same Φ^r governs both pools. Applying equation (37) to the regional totals would treat the entire region as a single pool. The resulting surplus estimate would attribute sorting caused by differences in pool composition to surplus. We use this estimate only as a one-pool benchmark.

The baseline model parameterizes the surplus matrix shared by the two pools as

$$\Phi_{ij}^{r} = \alpha_i^r + \gamma_j^r + \lambda S_{ij}, \qquad S_{ij} = S_{ij}^{p} + S_{ij}^{e}. \tag{40}$$

The region-specific intercepts shift the value of every match involving a given type. Only the sum $\alpha_i^r + \gamma_j^r$ enters joint surplus: $(\alpha_i^r + c, \gamma_j^r - c)$ produces the same surplus matrix for any c. We therefore normalize one intercept in each region to zero. This selects a unique set of female- and male-type intercepts without changing joint surplus. The bilateral component S_{ij} varies across pairings. The experiment supplies its party and education components, and the common conversion factor λ maps them into joint-surplus units. Given Φ^r and π, the calibration solves the equilibrium in each pool, adds the pool-specific couple and single masses to form regional predictions, and compares those predictions with the observed regional masses.

This parameterization makes the depolarization counterfactual well defined. The experiment identifies S_{ij}^{p} as the party component that is removed, while S_{ij}^{e} and the other primitives are held fixed. Matching data alone cannot provide that label: even within an observed pool, they identify total surplus rather than a unique decomposition into party, education, and other components.

The parameterization also imposes two restrictions. First, one conversion factor scales the party and education components together, preserving their relative magnitudes from the experiment. Second, the components enter additively. The value of matching on party or education can still vary across types, but there is no additional effect tied to the combination of the two match statuses. Thus, when a couple differs on both party and education, the bilateral effect is the sum of the separately measured party and education effects, with no extra penalty or premium for the double mismatch.

The restrictions also determine what the matching data can estimate. Adding an unrestricted bilateral residual B_{ij} alongside λS_{ij} would leave λ unidentified: for any scalar c, replacing λ by $\lambda + c$ and B_{ij} by $B_{ij} - cS_{ij}$ leaves Φ_{ij}^{r} unchanged. The common conversion factor preserves the relative party and education components measured in the experiment. Section 7.3 instead relaxes restrictions on meeting-pool assignment, and Section 7.3 allows the conversion factor to differ across regions.

The experimental estimates of S_{ij} themselves carry uncertainty. The calibration carries that uncertainty forward by rebuilding S_{ij} and re-estimating λ, π, and the intercepts across experimental draws. Appendix E compares the structured specification with the unrestricted one-pool benchmark and reports the contact diagnostics.

A.5 Welfare

The welfare measure in Section 2.3 follows from the same partner-type choice problem that generates matching shares. This subsection derives its inclusive-value form, relates it to equilibrium singlehood, and defines type-level welfare across meeting pools.

Return to the explicit region and pool superscripts. Before a type-i woman in pool h of region r observes her taste draws, her expected payoff is

$$\mathbb{E}_{\varepsilon}\left[\max\left\{\varepsilon_{0\ell},\ \zeta^{r}_{fi}+\varepsilon_{o\ell},\ \max_{j\in\mathcal{J}}\left(U^{rh}_{ij}+\varepsilon_{j\ell}\right)\right\}\right]=\log\left(1+\exp(\zeta^{r}_{fi})+\sum_{j\in\mathcal{J}}\exp(U^{rh}_{ij})\right)+C, \tag{41}$$

where the equality follows from the Type-I extreme-value distribution (McFadden, 1978; Galichon and Salanié, 2022). The additive constant C is the same in every equilibrium and therefore cancels from welfare changes. Dropping it gives the normalized welfare measure in equation (12).

The logit choice probabilities link this inclusive value directly to equilibrium singlehood. The odds relation in equation (35) and feasibility give

$$1+\exp(\zeta^{r}_{fi})+\sum_{j\in\mathcal{J}}\exp(U^{rh}_{ij})=1+\frac{\mu^{rh}_{io}}{\mu^{rh}_{i0}}+\sum_{j\in\mathcal{J}}\frac{\mu^{rh}_{ij}}{\mu^{rh}_{i0}}=\frac{n^{rh}_{i}}{\mu^{rh}_{i0}}. \tag{42}$$

Because μ^{rh}_{i0}/n^{rh}_{i} is the type's singlehood rate within the pool, normalized welfare is its negative log, $\log(n^{rh}_{i}/\mu^{rh}_{i0})$. The analogous expression for a type-j man is $\log(m^{rh}_{j}/\mu^{rh}_{0j})$.

Regional type welfare averages the two pool-specific values using the assignment probabilities in equation (39):

$$\mathcal{W}^{r}_{i}=\sum_{h=1}^{2}q^{h}_{i}(\pi)\mathcal{W}^{rh}_{i},\qquad \mathcal{W}^{r}_{j}=\sum_{h=1}^{2}q^{h}_{j}(\pi)\mathcal{W}^{rh}_{j}. \tag{43}$$

Section 5.2 converts changes in these type-level welfare measures from normalized utility units into partner-income equivalents.

A.6 Existence, uniqueness, and computation

The transferable-utility assignment game has stable outcomes, and stable matchings maximize total surplus (Shapley and Shubik, 1971). Stable matching in separable large markets can be characterized as a convex optimization problem. The multinomial-logit specification yields a unique equilibrium matching matrix for given surplus, type masses, and outside payoffs (Galichon and Salanié, 2022). To apply the usual computation with the additional outside partnership, combine singlehood and outside partnership into a non-modeled-match mass $\nu_{i0}=(1+e^{\zeta_{fi}})\mu_{i0}$ for women and analogously for men. The resulting system is the standard matching system with effective surplus $\Phi_{ij}-\log(1+$

$e^{\zeta_{fi}}) - \log(1 + e^{\zeta_{mj}})$. The fixed outside odds then separate each composite mass back into singles and outside partnerships.

Within a pool, substituting the gravity equation into feasibility leaves a system in the singles masses alone:

$$\left(1 + e^{\zeta_{fi}}\right)\mu_{i0} + \sum_j \sqrt{\mu_{i0}\mu_{0j}} \exp\left(\frac{\Phi_{ij}}{2}\right) = n_i, \tag{44}$$

$$\left(1 + e^{\zeta_{mj}}\right)\mu_{0j} + \sum_i \sqrt{\mu_{i0}\mu_{0j}} \exp\left(\frac{\Phi_{ij}}{2}\right) = m_j. \tag{45}$$

To solve the equilibrium in a pool, we take its type masses and surplus matrix as given. The iterative proportional fitting procedure begins with candidate single masses for women. Holding those masses fixed, it updates men's single masses so that the male feasibility conditions hold. It then holds the updated men's masses fixed and updates women's single masses so that the female conditions hold. The gravity equation maps the updated single masses into couple masses. The procedure repeats these steps until couples and singles exhaust every type's mass to a numerical tolerance.

This equilibrium calculation is nested within the model calibration. For any candidate values of the conversion factor, assignment parameter π, and type intercepts, the calibration constructs the surplus matrix and the masses assigned to each pool. The iterative procedure maps those inputs into predicted couples and singles. We solve the equilibrium in both meeting pools in each region and add the pool-specific outcomes to obtain the regional matching.

B DCE Design and Estimation Details

This appendix provides the implementation and estimation details for the discrete choice experiment in Section 3. Appendix B.1 defines the variables constructed from the displayed attributes, and Appendix B.2 describes the hierarchical estimator. Appendix B.3 documents the data-collection batches, analysis sample, and response-quality measures. Appendix B.4 describes adaptive task assignment and compares it with random assignment. Appendix B.5 describes the personalized feedback, and Appendix B.6 evaluates recovery under the implemented assignment design.

B.1 Partner attributes in the choice model

Before the choice tasks, respondents report their own gender, party identification, education, income, family plans, and preferred partner gender. Gender, party, and education place respondents in the eight types defined in Section 2.1. Party, education, income, and family plans also determine how a profile is compared with the respondent. For example, for a BA-educated Moderate Democrat who wants children and earns \$60,000, a Moderate Republican profile with some college, \$70,000 in income, and no plans for children is three steps away on the partisan scale, crosses both the party line and the BA threshold, takes the opposite position on family plans, and earns more than the respondent. Preferred partner gender determines the income grid from which profile earnings are

Table 7: Variables used in the choice model

Displayed attribute	Variable	Definition
Party	$\|p_i - p_k\|$; $\mathbf{1}[P_i \neq P_k]$	Distance on the six-level partisan scale; indicator for different party affiliations
Education	$\mathbf{1}[\text{BA}_i \neq \text{BA}_k]$	Indicator for opposite sides of the BA threshold
Income	$c\log(y_k/65{,}000)$; b_{ik}	Log annual partner income, with $c = 6.5$; indicator that the profile earns more than the respondent
Attractiveness	a_k	-1, 0, and 1 for below average, average, and above average
Religiosity	r_{ik}	0, 1, and 2 for similar, somewhat different, and very different views
Family plans	f_{ik}^{adj} and f_{ik}^{opp}	Indicators for adjacent and opposite positions on the three-level scale

drawn. For respondents open to either gender, the survey selects one grid at random.

Table 7 defines the variables for respondent i and profile k. The two profiles in a task always differ on at least one displayed attribute. For each choice, the model compares the profiles using the difference between their values of every variable in the table.

Party is displayed as Strong Democrat, Moderate Democrat, Weak Democrat, Weak Republican, Moderate Republican, or Strong Republican. The choice model includes both distance on this scale and an indicator for whether the respondent and profile have different party affiliations. Distance allows the value of a profile to change with partisan intensity; the indicator allows an additional change when a profile crosses the party line. The analysis excludes respondents who identify as Independent, and the profiles do not include an Independent level.

The matching model distinguishes Democratic from Republican types but does not retain partisan intensity. It therefore combines the distance coefficient and the discrete party-line shift into one binary cross-party penalty. Section 5.1 explains this mapping, and Appendix D.3 derives the design-based conversion of the distance coefficient. This appendix focuses on how the two party variables are coded in the choice model.

Education is displayed as high school, some college, a bachelor's degree, or a graduate degree. The baseline uses only whether the respondent and profile fall on opposite sides of the BA threshold. This specification matches the education partition in the matching model. It also makes explicit that, in the baseline specification, variation within the college and non-college groups does not receive a separate coefficient.

Partner income comes from seven-point grids calibrated to ACS earnings distributions for the profile's education and the respondent's preferred partner gender. The grids allow income to vary within an education category while excluding implausible schooling–income combinations. The fitted model uses $c\log(y_k/65{,}000)$ with its coefficient fixed at one. The constant $c = 6.5$ sets the numerical utility unit, and the denominator disappears from utility differences. Dividing all taste coefficients and the choice-error scale by c expresses the same model in unscaled log-income units.

The higher-earner indicator b_{ik} allows a discrete response when the profile's income crosses the respondent's own income, beyond the smooth log-income effect (Bertrand et al., 2015).

Attractiveness, religiosity, and family plans enter the same utility function as party, education, and income. Their inclusion means that the party and education coefficients compare profiles with the other displayed traits held fixed. Respondents are also instructed to treat unlisted traits, including age, race, and personality, as identical across the two profiles. The comprehension question reported in Appendix B.3 checks whether they understood this instruction.

B.2 Pooled hierarchical estimation

We use a hierarchical mixed logit to estimate the type-specific party and education means required by the matching model while allowing preferences to vary among respondents of the same type. Because each respondent makes only 12 choices, the model pools information across respondents rather than estimating each respondent's coefficients separately.

Let θ_i collect respondent i's eight taste coefficients. The corresponding element of μ_i is the mean for the respondent's group under the pooling rule in Table 8. We assume

$$\theta_i \mid \mu_i, \Sigma \sim \mathcal{N}(\mu_i, \Sigma), \qquad \Sigma = \operatorname{diag}(s_1^2, \ldots, s_8^2). \tag{46}$$

The coefficients in θ_i are fixed across a respondent's tasks. The diagonal covariance assumes that, conditional on the applicable population means, respondent-level deviations in the eight coefficients are uncorrelated. This restriction reduces the number of covariance parameters estimated from only 12 choices per respondent.

Table 8 states the pooling restrictions. The party-distance, party-line-shift, and education means vary across the eight gender-by-party-by-education types. The higher-earner mean varies by gender and party. The remaining means are common across respondents, and each of the eight dispersions is common across types.

Table 8: Pooling in the hierarchical choice model

Parameter	Population grouping	Number
Party-distance mean	Gender × party × education	8
Party-line-shift mean	Gender × party × education	8
BA-threshold mean	Gender × party × education	8
Higher-earner mean	Gender × party	4
Attractiveness, religiosity, and family-plan means	Common across respondents	4
Coefficient dispersions	One for each taste coefficient, common across types	8
Choice-error scale	Common across respondents	1

For respondent i, let y_{is} denote the profile chosen in task s, let x_{is} contain the variables for the two profiles in that task, and let S_i denote the number of tasks. The marginal likelihood is

$$L = \prod_i \int \prod_{s=1}^{S_i} P(y_{is} \mid x_{is}, \theta_i, \sigma)\; f(\theta_i \mid \mu_i, \Sigma)\; d\theta_i, \tag{47}$$

where $P(\cdot)$ is the conditional-logit probability in equation (16). The product inside the integral uses the same coefficient vector for all of a respondent's choices. The integral averages how well those choices fit across the possible preference vectors the respondent could have drawn from the population distribution in equation (46). For each candidate set of population parameters, we approximate this average with 1,600 fixed standard-normal draws transformed by the candidate means and dispersions.

In the numerical units defined above, each population mean has an independent normal prior centered at zero with standard deviation ten. The within-type standard deviations s_q have half Student-t priors with three degrees of freedom and scale two, and the common choice-error scale σ has a half Student-t prior with three degrees of freedom and scale ten. In unscaled log-income units these three prior scales are $10/c$, $2/c$, and $10/c$. Variation in displayed income identifies σ because the income coefficient is fixed. Sampling the scale jointly with the taste distributions propagates its uncertainty into the preference estimates.

We combine the resulting likelihood with the priors to obtain the posterior distribution of the population parameters. We sample from this distribution using the No-U-Turn Sampler implemented in NumPyro (Hoffman and Gelman, 2014; Phan et al., 2019). Six chains each use 500 warm-up iterations followed by 1,000 retained draws, yielding 6,000 posterior draws. Across all estimated parameters, the maximum $\widehat{R}$ is 1.004 and the minimum effective sample size is 1,366; there are no divergent transitions. The final fit uses 1,600 integration draws per respondent, following numerical checks of smaller integration rules. This is distinct from the number of retained posterior draws.

B.3 Data collection, analysis sample, and response quality

Profile pairs were assigned using either the randomized or adaptive design described in Section 3.4. Table 9 maps the two assignment procedures to the four data-collection batches. The hierarchical estimator pools all retained paired-profile choices across these batches. Recruitment balanced respondents by gender and targeted the national education distribution across Prolific's education categories. It imposed no party quota, so the eight gender-by-party-by-education cells have different sample sizes.

Profile pairs in the first choice task were assigned using the same rule in every batch, with the aim of creating variation in whether a profile earns more than the respondent. When the income grids contain values on both sides of the respondent's income, one profile is assigned an income above that amount and the other an income at or below it. Otherwise, both incomes are drawn from the available grids. The remaining attributes are drawn randomly, and each profile's income is drawn from the grid for its education and the respondent's preferred partner gender. From the second task onward, the pilot, randomized main batch, and randomized top-up draw profile pairs randomly; the adaptive main batch uses the respondent's preceding choices to select each new pair.

Table 9: Data-collection batches and task assignment

Batch	Fielding dates	N	Tasks 2–12	Feedback
Pilot	May 4–8	85	Random	No
Randomized main	May 15–26	146	Random	Yes
Adaptive main	May 27–June 7	86	Adaptive	Yes
Randomized top-up	June 28–29	106	Random	Yes
Total		423		

Notes. Dates are the fielding dates for the four analysis batches in 2026. N is the post-restriction estimation sample. Every batch uses the common first-task construction described in the text. The pilot supplies the initial reference distribution for personalized feedback. The randomized top-up also includes a descriptive party-signal battery after the 12 choice tasks. The separate 120-person sample analyzed in Appendix C.5 completed the same 12 core tasks followed by an eight-task mechanism module and is outside this estimation sample.

The analysis sample contains respondents who can be assigned to one of the eight respondent types: Democrats or Republicans aged 25–45 who report a binary gender and an interest in opposite- or either-sex partners. Respondents open to either-sex partners remain because they evaluated the same opposite-sex profiles. The final requirements are completion of all 12 tasks and a pass on the preregistered attention check.

Table 10 begins with the 799 Prolific-linked survey records from the four data-collection batches used in the analysis. Four of those records stopped before the respondent answered the survey items needed to determine eligibility and construct the variables used to compare profiles. The remaining rows apply the substantive sample restrictions one at a time.

Table 10: From Prolific survey records to the estimation sample

Restriction	Responses
Prolific-linked survey records in the four analysis batches	799
Answers the survey items required for sample assignment	795
Democrat or Republican, reported binary gender, and interested in an opposite-sex or either-sex partner	579
Completes all 12 choice tasks	573
Prolific-verified age 25–45	426
Passes the attention check	423
Estimation sample	423

Notes. Each row applies the stated restriction to the row above. The four records removed in the second row ended before the respondent answered gender, party, education, income, family plans, or preferred partner gender. Party is measured on a seven-point self-identification scale; leaners count as partisans and pure Independents are excluded. Nine otherwise eligible respondents have no Prolific age record and are removed with the age restriction.

The sample resembles the target population in gender and party but is more educated. Table 11 compares it with Democratic and Republican adults aged 25–45 in the Cooperative Election Study (CES). The DCE college share is about nine percentage points higher. This imbalance does not give the college types greater weight in the headline preference averages: Section 4 gives the eight type estimates equal weight, and the matching model later combines them using population type shares from the CES rather than sample shares from Prolific.

Table 11: Estimation sample and CES population

Characteristic	Prolific DCE	CES (partisans, 25–45)
Mean age	35.4	34.7
% Female	50.1	50.9
% College (BA+)	51.8	42.9
% Democrat	66.0	60.5
Median income (bracket)	\$50–75k	\$50–60k

Notes. The DCE column is the unweighted 423-respondent estimation sample. CES figures are survey-weighted over Democratic and Republican adults aged 25–45 in the 2017-and-later cumulative file, with leaners counted as partisans. The income row compares the respondent's personal income in the DCE with household income in the CES and is therefore descriptive rather than a like-for-like benchmark. The DCE questionnaire does not collect race.

The preregistered attention question is the only response-quality measure used to exclude respondents. Asked after the choice tasks, it requires respondents to identify which trait did not appear in the profiles. Of the 426 respondents who otherwise meet the inclusion criteria, 423, or 99 percent, answer correctly.

Two questions before the choice tasks instead check whether respondents understand the instructions. Among the same 426 respondents, 79 percent correctly answer on the first attempt that attributes not shown in the profiles should be held constant; 93 percent correctly answer that they should make the hypothetical choices as if they were single. An incorrect answer sends the respondent to a correction page before the tasks and does not trigger exclusion. A separate question after the tasks asks how respondents actually pictured the unlisted traits. Among the final sample, 80 percent select one of the two responses that describe those traits as similar across the profiles.

B.4 Adaptive task assignment

For adaptive task assignment, we use the Bayesian Adaptive Choice Experiment of Drake et al. (2024). After each choice, the procedure updates a posterior distribution over the respondent's preference coefficients and selects the next profile pair expected to provide the most information about those coefficients. This concentrates the remaining tasks on informative comparisons without increasing the 12-task budget.

For tasks 2 through 12, the procedure tracks seven coefficients:

$$\theta_i^{\text{assign}} = (\beta_i^{\text{dist}}, \beta_i^{\text{edu}}, \beta_i^{\text{attr}}, \beta_i^{\text{relig}}, \beta_i^{\text{fam,adj}}, \beta_i^{\text{fam,opp}}, \beta_i^{\text{earn}}).$$

The first two coefficients are the penalties associated with partisan distance and crossing the BA threshold. The remaining coefficients measure preferences over attractiveness, religious distance, adjacent and opposite differences in family plans, and whether the profile earns more than the respondent. The assignment model used linear income and represented party through partisan distance alone. These were the rules used to collect the data. The final estimator uses log income and adds a separate penalty for crossing the party line; it conditions on the profiles actually assigned.

The prior used to choose tasks has type-specific means and standard deviations estimated from the pooled pilot and randomized main data. We double each standard deviation before using the prior, leaving its mean unchanged. This reduces how strongly the earlier data determine the tasks shown to a new respondent.[9]

For each candidate pair, the procedure uses the logit model in equation (16) to predict the respondent's choice. It holds fixed gender-specific choice-error scales estimated from the pilot and randomized main data. These scales are used only to choose tasks; the final estimator uses all batches and estimates the common scale described in Appendix B.2.

Figure 10 compares the information about preferences contained in the profile pairs actually shown in the adaptive and randomized main batches.

Figure 10: Information from adaptive and randomized task assignment

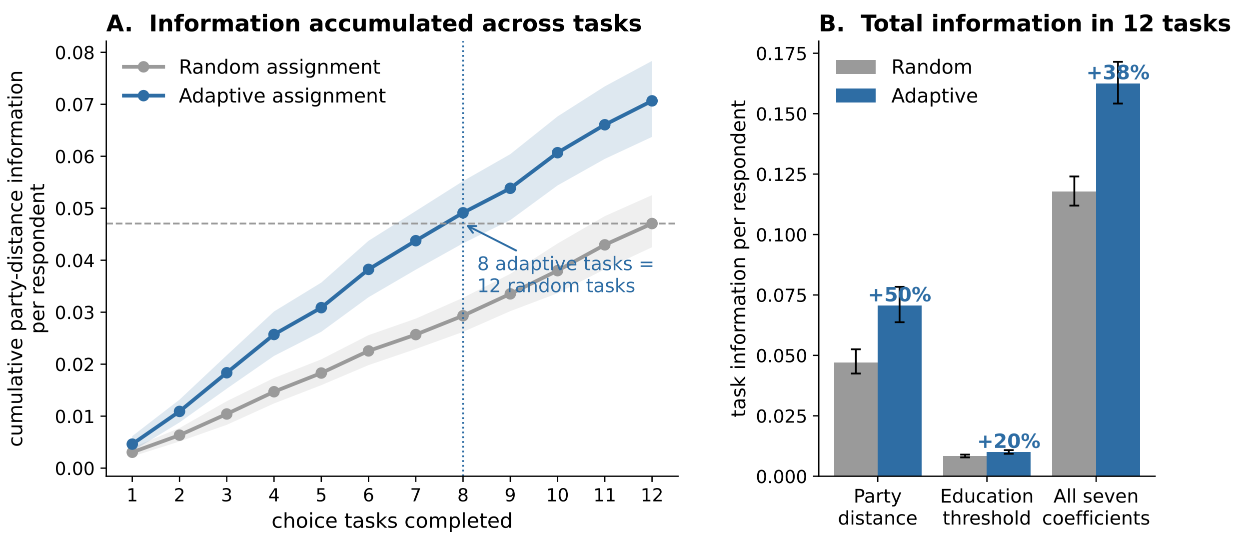

Notes. The figure uses profile pairs actually shown in the adaptive and randomized main batches. For every task, information is calculated using the population-mean coefficients applicable to the respondent's type and the gender-specific choice-error scale, both estimated from the pilot and randomized main data. The same coefficient and scale estimates are used for both batches. The measures first average respondents within type and then give the eight types equal weight. Shaded bands and capped error bars are pointwise 95 percent respondent-bootstrap confidence intervals, obtained by resampling each respondent's complete task sequence within type while holding these estimates fixed. Panel B's "all seven" measure includes partisan distance, the education threshold, attractiveness, religiosity, two family-plan terms, and the higher-earner indicator.

To compare the two assignment designs, we calculate Fisher information for every profile pair actually shown, using the same type-specific mean coefficients and gender-specific choice-error scales

[9] The preregistration specified the randomized main batch alone as the source of the adaptive prior. The implemented prior pooled the pilot and randomized main responses. Reconstructing the prior from the randomized main batch alone produces similar means and type-level intervals that are about 20 percent wider on average.

for both batches. Holding these estimates fixed ensures that differences in calculated information reflect the profile pairs produced by the two assignment designs. For a binary logit, Fisher information measures how much an observed choice can reveal about a coefficient (Train, 2009). The information that task s provides about coefficient q for respondent i is

$$I_{isq} = p_{is}(1 - p_{is}) \left(\frac{\Delta x_{isq}}{\sigma_{g_i}^{\text{assign}}} \right)^2, \tag{48}$$

where p_{is} is the predicted probability of choosing Profile B, Δx_{isq} is Profile B's value minus Profile A's value for the variable associated with coefficient q, g_i is respondent i's gender, and $\sigma_{g_i}^{\text{assign}}$ is the corresponding gender-specific scale. The expression is larger when the profiles differ more along the relevant dimension and when neither choice is close to certain.

Panel A uses only the party-distance coefficient. For each task number, we average I_{isq} across respondents within each type and then average the eight type-specific means. The curve adds these type-balanced averages through the indicated task, so a higher curve means that the task sequence has accumulated more information about the party-distance coefficient. Panel B first sums information across all 12 tasks for each respondent and then applies the same type-balanced averaging. Its "all seven" bar also sums across the seven coefficients used for adaptive assignment.

Panel A shows that the adaptive sequence accumulates as much information about party distance in eight tasks as the randomized sequence does in 12. Panel B shows that, over all 12 tasks, adaptive assignment provides about 50 percent more information about partisan distance, 20 percent more about the education threshold, and 38 percent more across all seven coefficients.

B.5 Personalized feedback and response incentives

Because the choices are hypothetical, the survey gives respondents a personal reason to consider them carefully. In the post-pilot batches, respondents were told before the tasks that their choices would generate a personalized summary of their partner preferences. Respondents who want an accurate summary therefore have a reason to choose profiles that reflect their preferences. This subsection describes how the summary is constructed and the condition under which it encourages sincere responses.

After the tasks, the summary reported the respondent's estimated willingness to pay for party, education, attractiveness, religiosity, and family plans. Figure 11 shows the feedback page. The blue marker and line show the respondent's posterior median and uncertainty interval. The gray dots show the distribution of estimates among respondents from earlier batches. For example, the adaptive batches use estimates pooled from the pilot and randomized main batch. The education panel measures the value of moving a profile up one displayed education level, rather than the BA-threshold mismatch used in the paper's main estimator.

The family-plans panel depends on the respondent's reported plans. For respondents who want children or do not want children, the blue marker measures the value of an opposite family plan.

Figure 11: Personalized preference feedback

Notes. Screenshot from the survey instrument. The blue marker and line show the respondent's posterior median and uncertainty interval. Gray dots show posterior medians for earlier respondents. In the family-plans panel, the gray dots show opposite-plan estimates; for respondents who are unsure about children, the blue marker instead shows their adjacent-plan estimate.

For respondents who are unsure, it instead measures the value of an adjacent family plan. The gray dots always show earlier respondents' opposite-plan estimates. They are therefore a general comparison group, rather than estimates of the identical coefficient, for respondents who are unsure.

The estimates in the summary come from a respondent-specific Bayesian choice model. For each respondent, population Monte Carlo approximates the posterior distribution of seven preference coefficients by combining normal priors centered at zero with the likelihood of the respondent's 12 choices. Coefficient values that better predict those choices receive more posterior weight. The posterior median and interval determine the blue marker and line shown in the summary.

The prior standard deviation for each coefficient is based on the estimated standard deviation of the corresponding coefficient across pilot respondents, with two adjustments.[10] The logit choice-error scale is fixed using estimates from earlier respondents rather than estimated separately from only 12 choices. For every respondent, the feedback posterior is calculated only after the 12th choice. In the adaptive batches, it is separate from the posterior used to select tasks. Neither posterior enters the final hierarchical estimator.

The feedback mechanism can encourage sincere choices only if respondents value an accurate summary. We formalize this condition as follows:

A1. Diagnostic preferences. A respondent values a summary that accurately describes their preferences and how they compare with those of previously collected respondents.

Under A1 and the maintained choice model, choices that reflect the respondent's preferences lead the posterior to place more weight on coefficients that describe those preferences. The design also makes targeted manipulation difficult: respondents were not shown the priors, calculation, or comparison distribution, and the summary appeared only after their answers were locked. These features do not eliminate strategic responding.

The mechanism has two limitations. First, a respondent who values flattering feedback, or feedback that confirms their political identity, more than accuracy may choose profiles to obtain a desired summary rather than to report their preferences. Second, with only 12 choices, the individual posterior remains sensitive to the zero-centered priors described above. Estimates of large preferences may therefore be pulled toward zero even when respondents answer sincerely. Under A1, the feedback gives respondents a reason to answer carefully and sincerely. Hypothetical-choice bias can still affect their answers.

B.6 Assignment-design recovery under linear income

Figure 10 shows that the profile pairs assigned adaptively contain more information about the coefficients used to select them. A simulation conducted under the linear-income specification compares the implemented mix of adaptive and randomized tasks with an equal-budget randomized benchmark. The simulation evaluates assignment and recovery with linear income. Its results do not establish recovery for the log-income estimator used in the main text.

Each of 20 simulation replications contains 423 respondents and matches the analysis sample's composition across the eight respondent types and assignment procedures. Population means, coefficient standard deviations, and the choice-error scale are calibrated to the eight-coefficient linear-income posterior. The implemented-design arm reproduces the survey assignment rules: the common first task in every batch, random assignment for tasks 2–12 in the three randomized batches, and the seven-coefficient adaptive procedure in the adaptive batch. The benchmark holds the simulated respondents, choice shocks, and first task fixed but assigns tasks 2–12 randomly in every

[10] Because the pilot hierarchical model estimated a BA-threshold mismatch rather than the education-level coefficient shown in the feedback, the latter uses the estimated standard deviation of the pilot mismatch coefficient. For opposite family plans, the feedback calculation uses six rather than the imprecisely estimated pilot value of 13.01.

batch. Each respondent completes 12 choices, and both arms are estimated with the eight-coefficient linear-income hierarchical model and a jointly estimated common choice-error scale. The simulation preserves the adaptive procedure used to collect the data while evaluating the additional party-line term in the estimation model.

Table 12: Linear-income recovery under the implemented and randomized designs

	RMSE			90% interval coverage	
Target	Random	Implemented	Difference (MC s.e.)	Random	Implemented
Combined cross-party penalty	3.400	3.320	-0.080 (0.151)	0.863	0.900
BA-threshold coefficient	2.003	1.986	-0.017 (0.125)	0.863	0.881

Notes. Each entry is based on 20 replications of the 423-respondent design. Within a replication, RMSE is calculated across the eight respondent-type means; the table reports the average RMSE across replications. The difference is implemented minus the randomized benchmark, with its Monte Carlo standard error in parentheses. Coverage is the share of the 160 type-by-replication 90 percent credible intervals that contain the corresponding population mean. Coefficients are measured in units of $10,000 of annual partner income. The combined cross-party component adds the party-line shift to the average cross-party change in partisan distance. Lower RMSE indicates better recovery; nominal coverage is 0.90.

The implemented design and randomized benchmark recover the linear-income inputs similarly. For the combined cross-party component, average RMSE is 3.32 under the implemented design and 3.40 under the benchmark. For the BA-threshold component, the corresponding values are 1.99 and 2.00. The paired differences are small relative to their Monte Carlo standard errors. Coverage under the implemented design is 0.90 for the party component and 0.88 for the education component, compared with 0.86 for both under the benchmark. All 40 fits converged and had no divergent transitions.

At the study's sample size and task budget, this linear-income simulation finds no material loss in recovery from using adaptive assignment for one batch, and no large efficiency gain over the randomized benchmark.

C Additional DCE Evidence and Interpretation

This appendix provides supporting evidence for the results in Section 4. It reports additional raw choice patterns, explains how the income specification determines the dollar interpretation, gives exact estimates by respondent type, compares the choice estimates with direct importance ratings, and examines the policy and social meanings attached to a displayed party label.

C.1 Additional raw choice patterns

Panel B of Figure 2 separates respondents by education because their available income gaps differ sharply. Respondents with a BA supply most tasks in which the same-education profile earns more, while respondents without a BA supply most tasks in which the different-education profile earns more. The center bin is balanced, with 484 tasks from respondents with a BA and 460 from

Figure 12: Raw choices on control attributes across income gaps

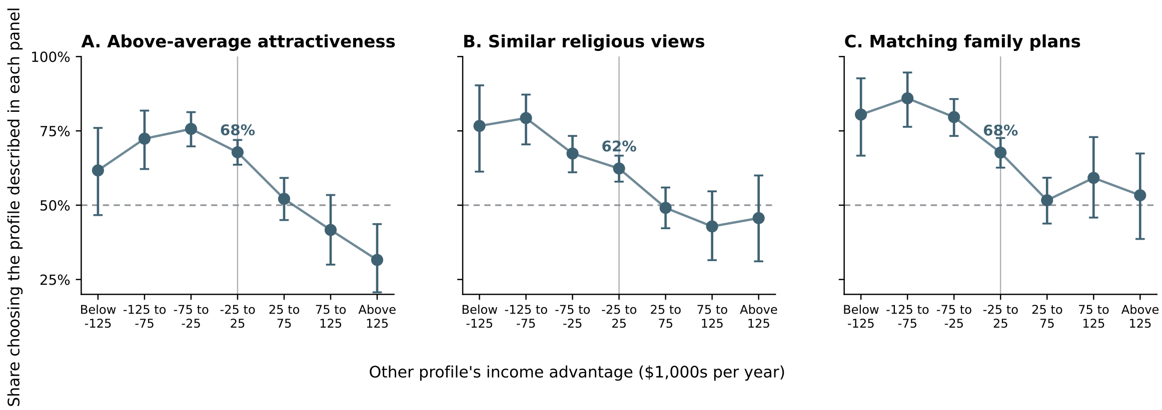

Notes. Panels A–C use tasks spanning the full range of the indicated attribute: above- versus below-average attractiveness, similar versus very different religious views, and family plans that match versus oppose the respondent's. The income gap equals the other profile's income minus the income of the profile described in each panel, so positive values mean that the other profile earns more. The seven bins match Figure 2. Markers are raw choice shares, and vertical lines are 95 percent confidence intervals that account for repeated choices by the same respondent. Other displayed attributes continue to vary, so the shares are descriptive rather than adjusted estimates.

respondents without one. Estimating education preferences separately by respondent type allows the choice model to use this variation while holding the other displayed attributes fixed.

Table 13 reports the group-specific support behind Panel B. The main figure omits two nonempty cells with fewer than 20 tasks: one task for respondents with a BA in the $75,000–$125,000 bin and four tasks for respondents without a BA in the -$125,000 to -$75,000 bin. The two zero cells cannot arise from the education-specific income grids. All tasks remain in the choice-model estimation.

Table 13: Task counts behind the raw education-choice figure

	Income advantage of the profile across the respondent's BA threshold ($1,000s)						
Respondent group	< -125	$[-125, -75)$	$[-75, -25)$	$[-25, 25]$	$(25, 75]$	$(75, 125]$	> 125
BA	145	209	420	484	87	1	0
NoBA	0	4	71	460	393	178	138

Notes. Counts include every task with one profile on each side of the BA threshold. Positive values mean that the profile across the respondent's BA threshold earns more. The main figure omits group–bin cells with fewer than 20 tasks; the choice model uses all tasks.

Figure 12 reports the raw choices for the three control attributes.

When the profiles span the full range of an attribute, respondents choose the above-average rather than below-average attractiveness profile 63 percent of the time, the profile with similar rather than very different religious views 61 percent of the time, and the profile with matching rather than opposite family plans 67 percent of the time. When the other profile earns over $75,000 more, the corresponding shares fall to 37, 44, and 56 percent. Respondents are therefore more willing to trade attractiveness or religious similarity for a large income advantage than agreement about family plans. These are raw shares; the choice model estimates each attribute's effect while holding the remaining attributes fixed.

Figure 2 uses all 2,562 tasks with one same-party and one cross-party profile and all 2,590 tasks with one profile on each side of the respondent's BA threshold. It groups each set into seven bins according to the income advantage of the cross-party or different-education profile. The center bin covers gaps from -\$25,000 to \$25,000, the four interior bins have widths of \$50,000, and the endpoints collect gaps beyond ±\$125,000. Bin counts range from 104 to 1,033 tasks in the party panel and from 138 to 944 in the education panel.

The choice model uses the full range of income differences in all 5,076 tasks. Among the 2,562 tasks pairing one same-party profile with one cross-party profile, the cross-party profile earns at least \$100,000 more in 197 tasks and at least \$200,000 more in 14. Across all tasks, the joint variation in income and the other displayed attributes identifies how much each attribute moves choices relative to income.

The main specification uses log income, so an income ratio rather than a dollar gap determines the smooth income contribution to utility. Large income advantages occur in the experiment, but the predicted choices also depend on the other displayed attributes. The standardized comparisons in Section 4.2 hold those attributes fixed. The next subsection compares the predictive performance of log and linear income and distinguishes local dollar equivalents from finite compensation.

C.2 How the income specification affects dollar equivalents

The baseline uses log income. Five-fold respondent-level cross-validation compares otherwise identical log- and linear-income specifications, keeping all 12 choices from a respondent in the same fold. Log income has a higher predictive log score and a lower mean squared prediction error in every fold (Figure 13). Across held-out choices, mean squared error falls from 0.1710 to 0.1678. The improvement is modest but consistent across this partition. The folds overlap in their training data, and adaptive-task predictions condition on the realized profile sequences; the exercise is not a new prospective experiment.

With concave income utility, a local dollar equivalent and a finite compensating amount are different objects. To make the distinction explicit, suppose income enters as $c \log y$, where $c > 0$ scales the income term, and let $A > 0$ be an attribute penalty in the same utility units. At reference income y, division by marginal utility gives the local equivalent

$$m(y) = \frac{A}{c/y} = \frac{Ay}{c}. \tag{49}$$

The local equivalent is not the finite amount associated with a discrete change. The income a respondent would forgo to obtain the preferred attribute is

$$\mathrm{WTP}(y) = y\left(1 - e^{-A/c}\right), \tag{50}$$

Figure 13: Held-out fit of log and linear partner income

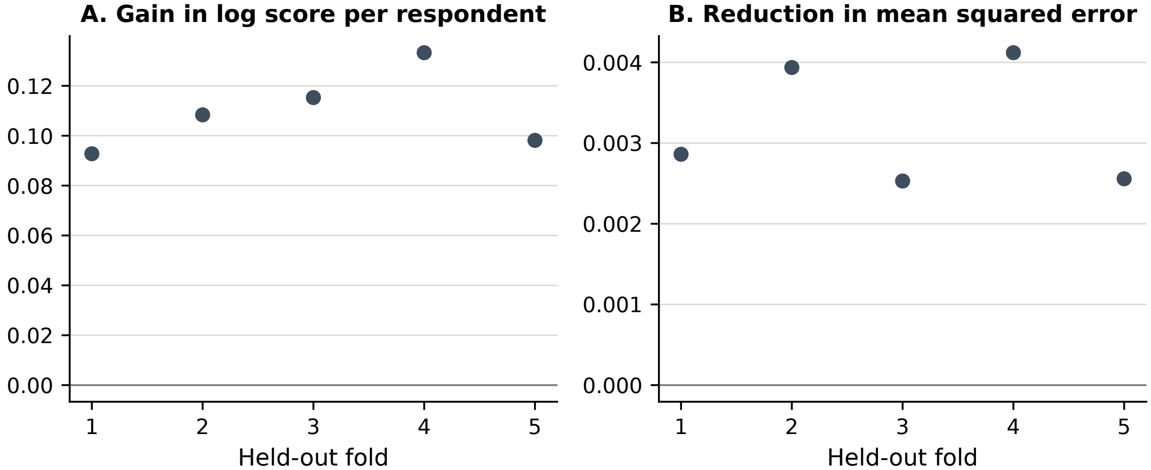

Notes. Each marker is one respondent-level validation fold. Panel A reports the gain in log predictive score per respondent under log rather than linear income. Panel B reports the reduction in mean squared prediction error. Positive values favor log income in both panels. The choice models are re-estimated on the remaining respondents; all tasks from a respondent stay in the same fold.

whereas the extra income required to compensate for its absence is

$$\mathrm{Comp}(y) = y\left(e^{A/c} - 1\right). \tag{51}$$

The two finite amounts coincide under linear income but diverge under log income, especially when A is large. With diminishing marginal utility, income forgone is smaller than the local equivalent because marginal utility rises as income falls; extra compensation is larger because marginal utility falls as income rises. Here $c = 6.5$ is the numerical scale used in estimation (Appendix B.2). All three formulas hold the higher-earner indicator fixed. If compensation moves a profile across the respondent's own income, that discrete utility change must also be included. Large finite changes can also leave the displayed income range. The main figures therefore report standardized choice probabilities; the local dollar comparison for a one-step party-distance change uses the median displayed income of $65,000. Structural welfare uses a separately stated reference of $62,500.

C.3 Exact estimates by respondent type

Table 14 reports the probabilities underlying Figure 5, together with the sample size of each respondent type. Each posterior draw integrates choice probabilities over the estimated normal preference distribution within type. The table then reports posterior medians and 95 percent credible intervals.

Table 14: Predicted choices by respondent type (percent)

Type	N	Same party	Party line	Same BA side
Women, D, BA	86	91.4 [87.1, 94.4]	77.7 [68.0, 85.4]	63.5 [55.2, 71.0]
Women, D, NoBA	62	91.0 [85.3, 94.7]	70.0 [56.9, 80.6]	44.8 [35.2, 54.5]
Women, R, BA	24	73.2 [58.8, 84.5]	56.3 [36.4, 74.5]	64.4 [50.0, 76.5]
Women, R, NoBA	40	75.6 [65.1, 84.0]	53.9 [37.7, 69.2]	53.2 [41.0, 65.3]
Men, D, BA	73	87.7 [81.9, 91.8]	71.5 [59.5, 81.4]	56.3 [47.2, 65.6]
Men, D, NoBA	58	95.2 [91.8, 97.3]	87.9 [79.5, 93.4]	49.2 [39.0, 59.2]
Men, R, BA	36	66.5 [54.7, 77.0]	59.7 [41.8, 74.8]	47.8 [36.4, 60.0]
Men, R, NoBA	44	72.6 [62.2, 81.3]	60.6 [44.7, 74.5]	50.5 [39.7, 61.0]

Notes. Posterior medians and 95 percent credible intervals from the 423-respondent log-income model. Same party combines the 2.11-step average distance contrast and the party-line shift. Party line holds distance fixed. Same BA side compares a profile on the respondent's side of the threshold with one across it. Other traits and income are held fixed; 50 percent indicates no average preference. D and R denote respondent party. College Republican women, with 24 respondents, is the only type below the preregistered target of 30.

The headline average gives the eight respondent types equal weight rather than weighting them by their realized sample shares. The matching model instead uses CES population masses. These are different aggregations of the same type-specific preference estimates.

Preference heterogeneity is distinct from uncertainty about the type mean. With independent normal deviations in the distance and party-line coefficients, the combined penalty within a type is normal with mean ψ_t and standard deviation $\{(19/9)^2 s_{\text{dist}}^2 + s_{\text{party}}^2\}^{1/2}$. The normal probability that this penalty is positive gives the model-implied share who prefer a same-party profile. Averaging those shares equally across types gives 89 percent at the posterior median. This preference share differs from the 81 percent predicted same-party choice share, which also incorporates choice noise.

Education preferences must be read relative to the respondent's own education. A same-side probability above 50 percent favors BA for a BA respondent and non-BA for a non-BA respondent. The point estimates favor BA for three of the four BA types and two of the four non-BA types; the intervals are wide. Appendix D reports the underlying signed coefficients.

C.4 Stated importance and choice-implied tradeoffs

Direct ratings provide a descriptive comparison with the profile choices. Respondents rate the importance of each partner attribute on a one-to-five scale. Family plans have the highest mean, followed by political identity and attractiveness. Table 15 places these ratings beside the standardized predicted-choice shares from the log-income model. The scales answer different questions, but both place political identity among the leading attributes and education below the others.

Table 15: Direct importance ratings and standardized predicted choices

Attribute	Mean rating	N	Predicted choice (%)
Political identity	3.46	420	81.4 [78.1, 84.5]
Children	3.67	413	76.7 [72.7, 80.2]
Attractiveness	3.43	420	77.2 [73.7, 80.5]
Religion	3.15	418	74.3 [70.5, 77.8]
Education	2.87	421	53.6 [49.9, 57.3]

Notes. Ratings run from one to five and use available responses in the 423-person analysis sample. Choice shares are posterior medians with 95 percent credible intervals. They compare same versus cross party, matching versus opposite family plans, above- versus below-average attractiveness, similar versus very different religion, and the same versus opposite side of the respondent's BA threshold. Choices integrate over within-type heterogeneity and weight the eight types equally.

C.5 Party labels, displayed issue positions, and partisan beliefs

The headline party coefficient measures the total response to a displayed party label while the other partner attributes vary independently. To ask how much of that response remains when profiles also display concrete political positions, we fielded a mechanism module with 120 additional respondents. The module followed the main choice experiment, in which respondents chose between two profiles. In the module, the profiles independently varied party, income, and positions on abortion, same-sex marriage, immigration, and gun purchases.

The prespecified attention rule retains 117 respondents overall. The analysis of the remaining party-label effect uses the 100 respondents who identify with or lean toward one of the two parties: 79 Democrats and 21 Republicans. Their eight choices each yield 800 choice tasks.

This module asks for the party-label effect that remains after the four issue positions are held fixed. Its reported model retains linear income, unlike the log-income specification for the main experiment. The hierarchical model includes an opposite-party indicator and the four issue positions. These five coefficients vary across respondents, and their population means differ by respondent party. Income enters in units of \$10,000 with coefficient one, and the common choice-error scale is estimated. A single opposite-party indicator captures the residual label effect because this module displays only Democratic and Republican labels. The reported average gives the two party means equal weight. Its dollar coefficients should not be subtracted from the main experiment's log-income coefficients.

The issue-position estimates have the expected partisan pattern. Democratic respondents prefer the more liberal position on abortion, same-sex marriage, immigration, and gun purchases; the 95 percent credible intervals exclude zero for all four. Republican respondents' point estimates favor the more conservative position on all four issues. Their intervals are wider because the analysis sample contains 21 Republicans, compared with 79 Democrats, and only the gun-purchase interval excludes zero.

The likelihood averages over 400 possible respondent-level preference vectors for each set of population parameters. We use the same priors and sampling procedure as the main choice model. Six independent runs produce 6,000 posterior draws. The maximum $\widehat{R}$ is 1.002, the minimum

Figure 14: Issue-position preferences in the party-label mechanism experiment

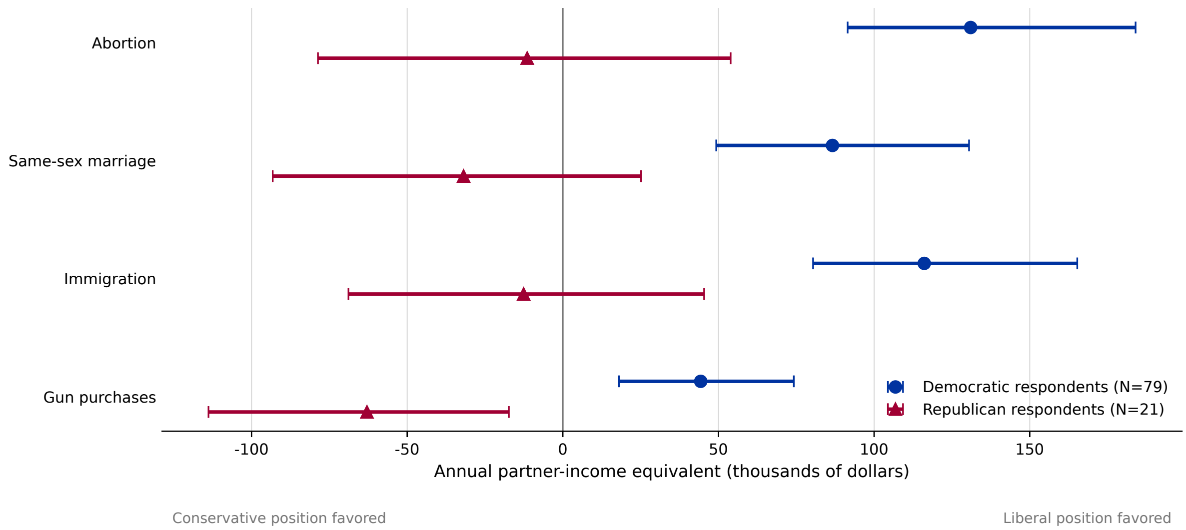

Notes. Markers are posterior medians from the hierarchical Bayesian choice model; lines are 95 percent credible intervals. Positive values indicate a preference for the liberal displayed position and negative values indicate a preference for the conservative displayed position. The horizontal axis is the annual partner-income equivalent in thousands of dollars. The income coefficient is fixed at one, and the common choice-error scale is estimated. The analysis contains 100 identifiers and leaners—79 Democrats and 21 Republicans—who completed eight mechanism choices each.

effective sample size is 2,397, and there are no divergent transitions.

After the four positions are held fixed, the opposite-party label carries an annual partner-income penalty of $81,000 [$56,000, $111,000] for Democratic respondents and $72,000 [$32,000, $118,000] for Republican respondents. Giving the two population means equal weight yields $76,000, with a 95 percent credible interval of [$51,000, $106,000]. Panel A of Figure 15 reports this equal-weight average.

The $76,000 measures the residual label effect within this module. Integrating over its estimated preference heterogeneity and choice noise gives the 79 percent same-party-label choice share reported in Section 4.4. For the supplementary counterfactual in Appendix E.10, we use the residual label's share of the module's combined label and belief-weighted issue penalty, not a dollar subtraction across experiments. Restricting that calculation to the model population leaves 88 respondents—69 Democrats and 19 Republicans. Applying their party-specific shares to the main experiment's penalty requires the additional assumption that the within-module proportions transfer across designs.

The checklist in Panel B helps describe that additional content. Seventy-five percent say that an opposing-party label brings specific political issues to mind, 68 percent select moral values or worldview, 65 percent select political leaders or style, and 59 percent select family life or gender roles. Fifty-one percent select treatment of themselves or people like them, while 14 percent select ambition, work ethic, or attitudes about money. Respondents could select multiple answers, so these categories describe overlapping associations rather than a decomposition of the label effect.

Figure 15: What remains in a party label after issue positions are shown

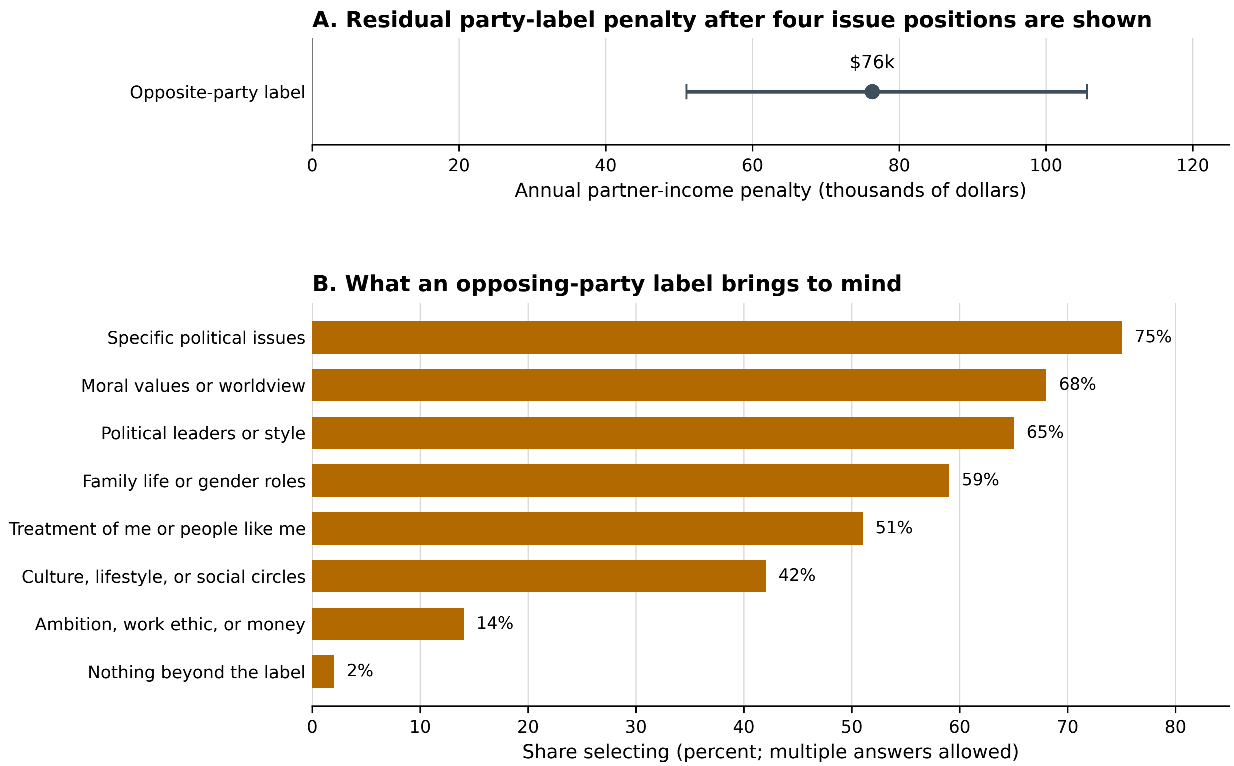

Notes. The supplementary sample contains 120 respondents, each of whom completed 12 core tasks before the eight-task mechanism module analyzed here. The prespecified attention rule leaves 117 respondents overall and 100 who identify with or lean toward one of the two parties for the analyses shown here. Panel A reports the equal-weight average of the Democratic and Republican population means from the hierarchical Bayesian choice model, after income and four independently randomized issue positions are held fixed. The marker is the posterior median; the line is the 95 percent credible interval. Panel B reports the share of the same 100 respondents selecting each association with an opposing-party label. Multiple selections were allowed.

All 117 respondents who pass the attention check answer two questions for each focal position: how many out of 100 Democrats would support it, and how many out of 100 Republicans would support it. The comparison by respondent party uses the 79 Democratic and 21 Republican identifiers and leaners in this sample. The four positions are treating abortion as a personal choice, allowing same-sex couples to marry, allowing unauthorized immigrants to remain and qualify for citizenship if they meet certain requirements, and making gun-purchase rules stricter. Each position corresponds to a named response category in the ANES. Following the standard measure of perceived policy polarization, we orient all four positions in the liberal direction and subtract predicted Republican support from predicted Democratic support (Levendusky and Malhotra, 2016). The ANES benchmark applies the same subtraction to survey-weighted support among Democratic and Republican identifiers and leaners aged 25–45 (ANES, 2025).

Panel A of Figure 16 shows who respondents believe occupies each party. On immigration, Democratic respondents predict 71 percent support among Democrats and 23 percent among Republicans, while Republican respondents predict 74 and 41 percent. The corresponding ANES shares are 58 and 40 percent. Democratic respondents therefore overestimate Democratic support and underestimate Republican support; Republican respondents mainly overestimate Democratic support. These beliefs imply gaps of 48 and 33 percentage points, respectively, compared with an 18-point ANES gap. On same-sex marriage, Democratic respondents predict 81 and 36 percent support, while Republican respondents predict 79 and 41 percent, compared with ANES shares of 88 and 54 percent. The resulting gaps are 45 and 38 points, compared with 35 points in the ANES. Democratic respondents also perceive wider gaps on abortion and gun purchases; the party-specific predictions and ANES benchmarks for all four issues appear in the figure.

Among the 100 respondents who identify with or lean toward a party, Democrats perceive an average gap of 50 percentage points across the four positions, compared with 36 points among Republicans. The difference is 14 points ($p = 0.046$), although the Republican subsample contains only 21 respondents. To relate these beliefs to the preceding choices, we use the respondent-level score from the seven-coefficient personalized-feedback model described in Appendix B.5. The score summarizes each respondent's 12 choices with one partisan-distance penalty. Because it is based on only 12 choices, it is a noisy summary of the respondent's preference. Its Spearman rank correlation with the mean perceived gap is 0.25 ($p = 0.012$). A partial rank correlation that accounts for the respondent's party is 0.20 ($p = 0.043$); the separate within-party correlations are 0.13 for Democratic respondents and 0.36 for Republican respondents. Panel B of Figure 16 plots the two relationships. These results connect beliefs about political distance to the party preferences expressed in the choices.

Figure 16: Perceived partisan differences and party preferences measured in the experiment

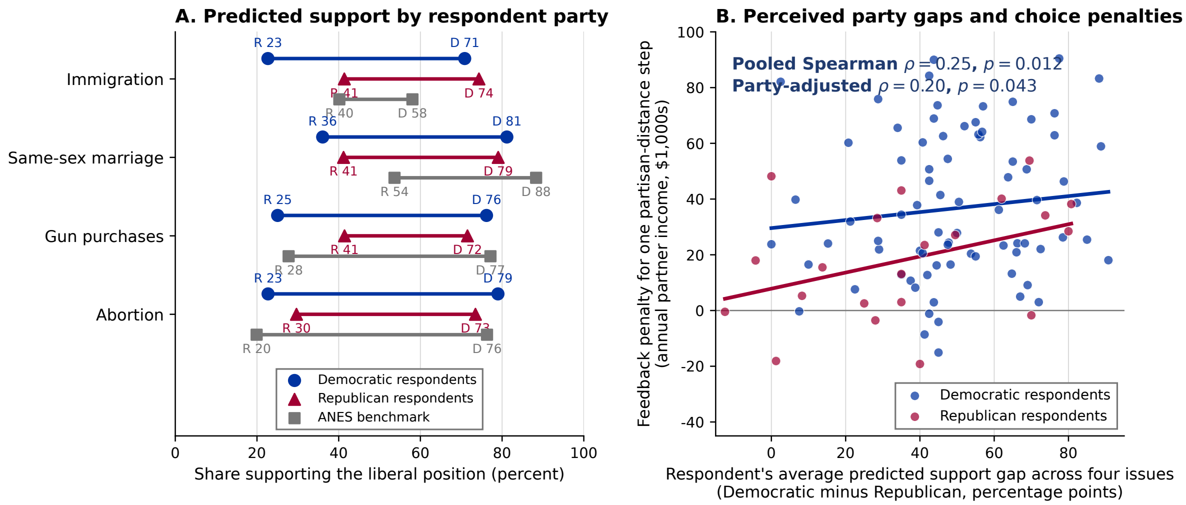

Notes. Panel A compares the mean predictions of Democratic and Republican respondents with the corresponding ANES shares. The letters D and R identify the party whose support is being predicted. All positions are oriented in the liberal direction. Both panels use the 100 respondents who identify with or lean toward one of the two parties. Panel B's horizontal axis averages each respondent's four predicted Democratic-minus-Republican differences. Each marker is one respondent. The vertical axis is the one-step party-distance penalty from the survey's personalized-feedback model, in thousands of partner-income dollars. This seven-coefficient, linear-income score is distinct from the current hierarchical estimate. The colored lines are separate fits for Democratic and Republican respondents.

D Experiment-to-Model Mapping Details

This appendix gives additional details on the mapping introduced in Section 5. It identifies the experimental objects carried into bilateral surplus, explains which estimated attributes remain outside the structural type space, documents the auxiliary reservation-value estimator, and derives the factor that converts the six-level party-distance coefficient into the binary cross-party contrast used by the matching model.

D.1 Objects carried to the matching model

Three sets of type-specific experimental means enter bilateral surplus: partisan distance, the discrete party-line shift, and crossing the BA threshold. Let $\bar{\beta}^{\text{dist},g}_{t,d}$, $\bar{\beta}^{\text{party},g}_{t,d}$, and $\bar{\beta}^{\text{edu},g}_{t,d}$ denote these means for gender g, party-by-education type t, and posterior draw d. Section 5.1 maps them into the matching-model primitives

$$\psi^{g}_{t,d} = \bar{\beta}^{\text{party},g}_{t,d} + \bar{d}_p\, \bar{\beta}^{\text{dist},g}_{t,d}, \qquad \tau^{g}_{t,d} = \bar{\beta}^{\text{edu},g}_{t,d}, \tag{52}$$

where $\bar{d}_p$ converts the party-distance penalty into the binary same-party versus cross-party penalty and $\bar{\beta}^{\text{party},g}_{t,d}$ supplies the discrete party-line penalty. Appendix D.3 derives the bridge factor. At the posterior medians, the combined cross-party penalty is positive for every respondent type. The education term is signed. Because its reference category is the respondent's own side of the BA threshold, a general preference for a college-educated partner appears with opposite signs for college and non-college respondents and is distinct from a preference for education matching.

The attractiveness, religiosity, family-plan, and higher-earner coefficients remain in the experimental utility function but not in the structural type space. They make the party and education estimates conditional on all displayed traits. Excluding them from bilateral surplus means that the matching counterfactuals change only the party and education components that the paper has chosen to model structurally. The accept-or-reject reservation values are excluded for a different reason: the matching model recovers type-specific outside-option components from observed singlehood rather than stated acceptance thresholds.

The higher-earner coefficient captures the discrete effect of a profile earning more than the respondent, beyond the smooth log-income effect. Retaining it in the choice model separates that relative-income response from income curvature, while excluding it from bilateral surplus keeps the structural type space focused on party and education.

Table 16: Type-level preference coefficients used in bilateral surplus

Type	Distance	Party line	BA threshold
Women, D, BA	3.86 [2.67, 5.23]	6.69 [4.07, 9.58]	2.76 [1.03, 4.70]
Women, D, NoBA	4.72 [3.33, 6.36]	4.57 [1.51, 7.66]	-1.05 [-3.20, 0.85]
Women, R, BA	2.48 [0.55, 4.46]	1.36 [-3.08, 5.83]	2.95 [0.01, 6.12]
Women, R, NoBA	3.11 [1.60, 4.75]	0.86 [-2.71, 4.44]	0.64 [-1.87, 3.16]
Men, D, BA	3.58 [2.33, 4.99]	4.97 [2.08, 8.03]	1.29 [-0.57, 3.36]
Men, D, NoBA	3.69 [2.31, 5.26]	10.49 [7.29, 14.14]	-0.15 [-2.34, 1.84]
Men, R, BA	1.17 [-0.33, 2.75]	2.13 [-1.83, 5.90]	-0.44 [-2.77, 2.09]
Men, R, NoBA	1.95 [0.37, 3.54]	2.33 [-1.16, 5.86]	0.11 [-2.15, 2.22]

Notes. Posterior medians and 95 percent credible intervals from the log-income model. Positive values are penalties for greater partisan distance, crossing the party line, or crossing the respondent's BA threshold. Coefficients use the fitted utility unit, in which income enters as $6.5\log(y/65{,}000)$, and are not dollar amounts. Divide by 6.5 to obtain unscaled log-income coefficients. The combined party penalty is constructed draw by draw before summarizing.

Figure 17 displays the bilateral surplus matrix used for the point estimates. Each cell adds the woman's and man's party and education terms according to equation (19). Cross-party pairings have substantially lower bilateral surplus; education differences make smaller changes.

The mapping is applied draw by draw. Each draw produces a new bilateral surplus component. The calibration then re-estimates the matching-model parameters and resolves the equilibrium. This carries experimental uncertainty, including its dependence on the jointly estimated choice scale, through the nonlinear matching and welfare calculations.

D.2 Auxiliary estimation of the accept-or-reject reservation value

The four accept-or-reject tasks ask a different question from the paired-profile tasks. A paired comparison identifies tradeoffs among partner attributes without requiring a stated reservation value. An accept-or-reject task locates a profile relative to remaining single. These stated reservation values are estimated in a separate auxiliary model and do not enter the matching-model calibration.

Figure 17: Experimental Bilateral Surplus by Partner Types

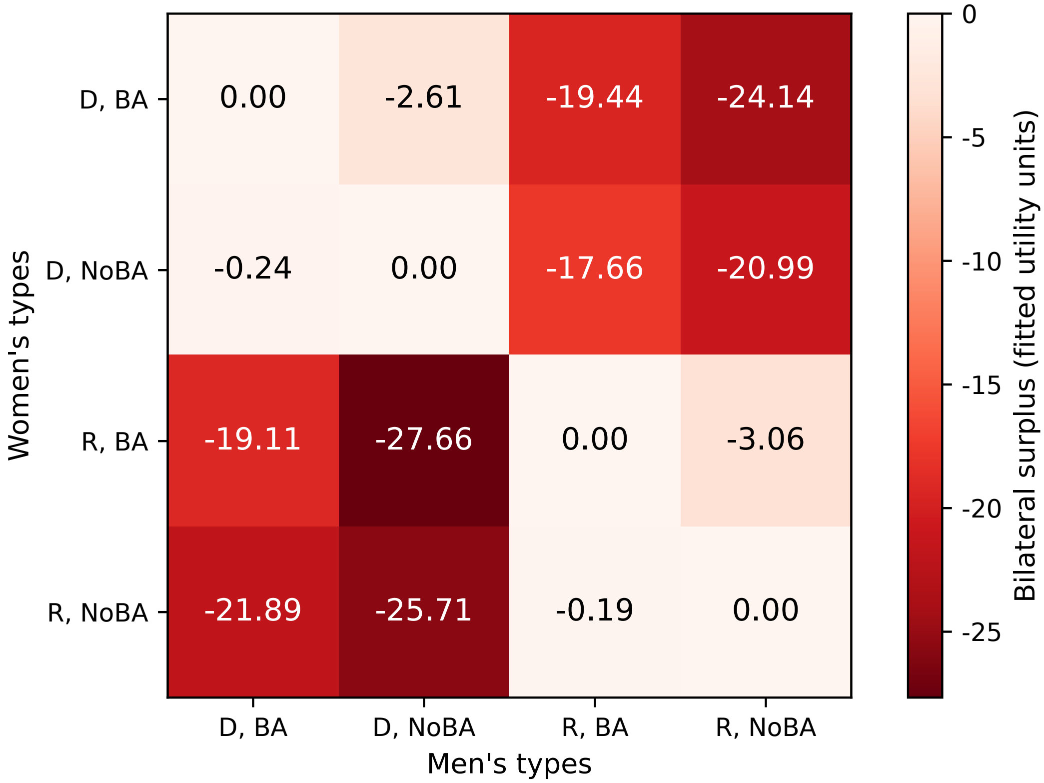

Notes. Cells show the bilateral component S_{ij} used by the point calibration, in the fitted log-income utility units defined above. Rows are women's types and columns are men's types. D and R indicate party; BA indicates a bachelor's degree or more. Darker red indicates lower surplus. Same-party, same-education cells are zero by construction. The conversion factor and regional type intercepts enter joint surplus separately.

For profile k, it writes

$$\Pr_i(\text{accept } k) = \left[1 + \exp\left(-\frac{u_i(k) - \delta_i}{\sigma^{\text{ar}}_{g_i}}\right)\right]^{-1}, \tag{53}$$

where δ_i is the respondent's stated reservation value for partnership, g_i is the respondent's gender, and $\sigma^{\text{ar}}_{g_i}$ is the corresponding accept-or-reject choice-error scale. Four binary observations do not support a separate δ_i for every respondent. The auxiliary estimator therefore recovers one reservation value δ^g_t for each gender-by-party-by-education type.

The analysis uses only randomly assigned accept-or-reject profiles. Wave 2 assigned these profiles adaptively to each respondent's evolving posterior and is excluded, so the type-level likelihood is evaluated on a common random-design basis. The accept-or-reject scale is profiled separately for women and men.

Let $\mathcal{I}^g_t$ denote the random-design respondents in type (g, t). For each paired-profile posterior draw d, we draw $R = 200$ coefficient vectors $\theta_{r,d}$ from that type's population distribution. A candidate δ is evaluated using

$$\log L^g_{t,d}(\delta) = \sum_{i \in \mathcal{I}^g_t} \log \left[\frac{1}{R} \sum_{r=1}^{R} \prod_{s=1}^{4} P\big(y_{is} \mid \theta_{r,d}, \delta, \sigma^{\mathrm{ar}}_g\big) \right]. \tag{54}$$

The product over s makes the four choices a panel governed by one coefficient vector; the average over r integrates over within-type preference heterogeneity. This auxiliary estimator uses seven taste coefficients, no separate party-line term, and linear income measured in units of \$10,000. It combines the likelihood with a $\mathcal{N}(-3, 3^2)$ prior for δ^g_t, evaluates the posterior on a 400-point grid from -10 to 2, and draws one δ^g_t for each paired-profile posterior draw.

The matching model sets the outside payoff from the ratio of outside partnerships to singles and estimates type intercepts from observed participation. Neither calculation uses these stated acceptance thresholds.

D.3 Party bridge factor

The experiment estimates a gradient on six partisan positions, whereas the matching model distinguishes only Democratic and Republican types. We code Strong Democrat as one, Moderate Democrat as two, Weak Democrat as three, Weak Republican as four, Moderate Republican as five, and Strong Republican as six. To construct a bridge that does not depend on the sample's distribution of partisan intensity, the baseline gives equal weight to the three partisan positions on each side for both respondents and profiles. Under this distribution, mean absolute distance is 8/9 for a same-party pairing and 3 for a cross-party pairing. Moving from the first to the second therefore increases expected distance by

$$\bar{d}_p = 3 - \frac{8}{9} = \frac{19}{9} \simeq 2.111. \tag{55}$$

Multiplying a cell's one-step distance penalty by $\bar{d}_p$ gives the distance component of this standardized cross-party contrast. Adding the cell's discrete party-line penalty gives the total penalty carried into the matching model.[11] Weighting both sides by the CES distribution of partisan intensity instead gives 2.587, a different, population-weighted contrast. The baseline uses the equal-weight factor in equation (55) for every respondent type. The calculation imposes equal respondent-intensity weights to standardize the contrast across types.

E Additional Identification and Estimation Details

This appendix documents the external inputs, the joint estimator, and the checks supporting Section 7. It distinguishes the model's equilibrium predictions from the adjusted couple tables used to estimate sorting and reports the supporting sensitivity analyses.

[11] Weighting respondents' partisan-identification categories by their frequencies in the retained experimental sample, while keeping profile positions equally weighted within party, gives $\bar{d}_p = 2.127$.

E.1 External Inputs and Market Normalization

The CES supplies regional population masses and singlehood rates, while HCMST supplies regional couple tables. Within each region, we normalize the full Democratic and Republican population to one, preserving the survey-weighted sex ratio and each sex's type composition. The woman–man equilibrium contains singles open to an opposite-sex partner, adults in modeled partnerships, and adults in outside opposite-sex partnerships. Same-sex-only singles and adults in existing same-sex partnerships remain fixed. Multiplying all masses by a common constant changes counts by that constant and leaves rates unchanged. CES regional population shares provide the weights for national outcomes.

The cumulative CES file contains 701,955 respondents. Restricting it to the 2017–2024 waves and ages 25–45 leaves 110,286 respondents. Requiring complete gender, education, party identification, partnership status, and a positive survey weight leaves 104,211. Excluding 19,065 non-leaning independents leaves the 85,146 respondents used to construct the model inputs. The smallest gender-by-party-by-education group within a census region contains 1,045 respondents.

The CES counts married respondents and those reporting a domestic partnership as partnered; separated, divorced, widowed, and never-married respondents are single. In the partisan analysis sample, 44.6 percent are married, 6.8 percent report a domestic partnership, and 48.6 percent are classified as single. Household data instead identify cohabitation from co-residence. The CES domestic-partnership category and the ACS unmarried-partner category are not equivalent, so the CES outcome should not be read as a comprehensive measure of being unpartnered. The CES classification supplies the participation rates used by the main estimates. The ACS partner links supply the conditional age and sex information used to construct partner eligibility.

HCMST provides a 4×4 couple table for each region. Rows index women's types i and columns index men's types j. Let C_{ij}^{r} be the survey-weighted count of couples in cell (i, j) of region r, including the ACS-based partner-age adjustment. Every regional cell enters the estimation objective. The estimator adjusts model tables to the observed margins separately in each region, so a discrepancy in a region's total weight or type composition does not count as an association discrepancy.

E.2 Constructing the Equilibrium Population

Population construction separates the five states in equation (1). Among adults classified as single, HCMST supplies the share attracted only to people of the same sex: 6.4 percent for men and 3.1 percent for women. The small samples motivate pooling within sex across party, education, and region. We apply these shares to CES single counts; singles attracted to both sexes remain in the woman–man equilibrium. Linked ACS partnerships identify the existing same-sex partnered share by region, sex, and education, applying the same rate to both parties within each cell. The resulting fixed groups contain 2.36 percent of the full population as singles and 1.07 percent as partnered adults.

For the remaining partnered population, the ACS determines whether the opposite-sex partner is also aged 25–45. Its household file links spouses and unmarried partners in both directions, observing exact age and sex. Eligibility probabilities condition on the respondent's sex, education, census region, age group, and marriage or cohabitation status. For a married person whose spouse lives elsewhere, we use the probability among otherwise similar married people whose spouse is observed. CES domestic partnerships receive the probability estimated from ACS unmarried partnerships. Combining this demographic adjustment with the party adjustment below assumes that partner party is independent of partner sex and age within the demographic groups.

HCMST and ANES determine which eligible partners are Democrats or Republicans. HCMST places many partners in a combined undecided, independent, or other category. ANES observes a person's own party and the partner's report of it. These reciprocal reports give the probability that someone placed in the combined category identifies as a Democrat, an Independent, or a Republican, conditional on the reporting partner's party. Applying these probabilities to HCMST produces a national table of partner types by party and education. We fit this table with female- and male-type effects, a separate association for each party pairing, and one common association for having the same education, without party-by-education interactions. Iterative proportional fitting applies the estimated association to each region while matching its CES margins. The resulting table gives each type the probability that its partner is a Democrat or Republican. Multiplying by the demographically eligible partnered population gives its raw modeled-partnership mass. This corrected table constructs populations; the baseline couple table used to estimate sorting retains classified D/R couples.

The final step equates the numbers of modeled partnered women and men within each region. The raw female total exceeds the raw male total in every region. We set their common total to the average and multiply each type within sex by the same factor. This preserves the type composition within sex. A feasibility cap prevents the modeled partnered mass from exceeding the available opposite-sex partnered mass in any cell. The cap does not bind at the point estimate and binds in 1.5 percent of joint-bootstrap replications. Nationally, reconciliation multiplies the raw male total by 1.058 and the raw female total by 0.948; regional factors range from 0.928 to 1.084, within the prespecified 10 percent limit.

Reconciliation changes the baseline split between modeled and outside partnerships, leaving each type's full population, observed singles, and fixed same-sex groups unchanged. The resulting equilibrium population is 96.6 percent of all D/R adults aged 25–45. For each type and region, the outside-partnership payoff is $\zeta^r_{gt} = \log(O^r_{gt}/L^r_{gt})$. Holding this payoff fixed does not hold the outside-partnership count fixed: singlehood, outside partnerships, and modeled partnerships all respond in the counterfactuals. Appendix E.11 reports sensitivities to partner-party classification and the construction of the fixed groups.

E.3 Equal-College-Composition Benchmark

The education-composition benchmark changes type masses rather than any individual's education. In each region, we move a common fraction of Democratic and Republican men without a BA into the corresponding BA types. We choose that fraction so that the total mass of college-educated men equals the observed mass of college-educated women. The regional fractions are 19.4, 17.7, 16.5, and 12.3 percent in the Northeast, Midwest, South, and West. This transformation preserves each region's population, its numbers of women and men, male party totals, the party composition of the men whose type changes, and every female type mass.

Within each resulting type, we preserve the baseline shares in singlehood, modeled partnerships, outside partnerships, and the fixed same-sex sector. The outside-partnership-to-single ratio therefore remains unchanged. We also hold fixed the experimental surplus matrix, the conversion factor, meeting-pool assignment, regional type intercepts, and outside-partnership payoffs, then re-solve the equilibrium. The exercise does not estimate how completing college would change a particular person's income, preferences, or partnership opportunities.

Table 17 reports the type-specific changes behind the aggregate result. Adding college-educated men erodes their scarcity advantage while improving the matching opportunities of college-educated women. The national singlehood reduction nevertheless remains positive because the composition change relaxes the aggregate imbalance between the two sides of the market.

Table 17: Singlehood Under Equal College Composition, by Type

Sex	Type	Current	Equal college	Change [95% interval]
Men	Democratic, BA	42.10	45.73	3.63 [3.04, 4.24]
	Democratic, no BA	66.57	64.71	−1.86 [−2.21, −1.54]
	Republican, BA	32.85	37.90	5.05 [4.14, 6.11]
	Republican, no BA	54.72	52.36	−2.35 [−2.85, −1.93]
Women	Democratic, BA	44.87	42.53	−2.33 [−2.79, −1.95]
	Democratic, no BA	55.70	56.27	0.57 [0.31, 0.92]
	Republican, BA	31.13	28.15	−2.98 [−3.65, −2.43]
	Republican, no BA	38.73	39.65	0.93 [0.61, 1.34]

Notes. Entries are full-population singlehood rates in percent. Changes are counterfactual minus current rates. Intervals are percentile intervals from 1,000 joint-bootstrap replications that resample the experimental posterior and the CES, ACS, HCMST, and ANES inputs and re-estimate the structural parameters. The counterfactual changes population composition while holding each type's calibrated characteristics fixed. It does not estimate individual responses to schooling.

E.4 The HCMST couple sample

HCMST has two roles. It provides the couple table used to estimate sorting and helps determine which CES partnerships fall inside the market. The sorting table uses couples in which both partners are classified as Democratic or Republican. For the population calculation, HCMST also retains Independent respondents and partners reported as undecided, Independent, or other. ANES shows how people identify when their partner gives one of those reports. This tells us what share of those HCMST partners are Democrats or Republicans.

Partner age requires a separate adjustment. HCMST reports partner age in broad ranges, and the ranges at either edge cross the 25–45 cutoff. We use exact partner ages from the ACS to estimate the share of each boundary range that falls inside the age window. That share multiplies

Table 18: HCMST Sample Construction and Party Reports Across Surveys

Panel A. HCMST sample construction

Sample restriction	Couples retained
2017 baseline respondents	3,510
Aged 25–44	1,205
Married or cohabiting	836
Opposite-sex couple	770
Partner-party question answered	765
Partner classified as Democratic or Republican	490
Respondent also classified as Democratic or Republican	485
Education observed for both partners	485

Panel B. Party reports in comparable partnered samples (%)

Survey	Report	Democratic	Independent or unclassified	Republican	*N*
HCMST	Respondent self-report	51.0	3.6	45.5	770
HCMST	Respondent's report of partner	32.1	36.6	31.3	765
ANES	Respondent self-report	46.7	4.5	48.8	559
ANES	Partner self-report	36.7	17.0	46.3	293
CES	Respondent self-report	45.3	17.8	36.9	53,475

Notes. Panel A reports unweighted counts after each restriction and before applying the ACS partner-age weights. Panel B reports survey-weighted shares among married or cohabiting respondents aged 25–45, except that HCMST's grouped age variable requires ages 25–44. The HCMST and ANES samples are limited to opposite-sex couples; CES does not report partner gender. Leaners are classified with a party. HCMST combines undecided, independent, and other in one response; ANES and CES identify a middle category through different question sequences. The HCMST partner row is reported by the respondent, while the ANES partner row is reported by the partner. The ANES partner study requires the partner to participate and does not provide a separate survey weight, so the main-respondent weight is used.

the HCMST survey weight. Couples in ranges wholly inside the window receive full weight, while couples in ranges wholly outside receive zero weight.

Panel B separates party composition from the way party is reported. HCMST respondents' own reports resemble the ANES respondent reports: 3.6 and 4.5 percent, respectively, occupy the middle category. The corresponding CES share is 17.8 percent, illustrating that this response is sensitive to the survey's question sequence. The distinctive pattern is the HCMST partner report. Respondents place 36.6 percent of partners in a category that combines uncertainty with independence and other responses, compared with 17.0 percent when ANES partners report their own party. The HCMST category therefore cannot be interpreted as a group of truly independent partners. Among reports classified as Democratic or Republican, the Democratic share is similar for HCMST respondents and partners, 52.9 and 50.7 percent.

Table 19 reports the share of unclassified partner reports within each observed group among the 741 couples with a Democratic or Republican respondent and a nonrefused partner-party response.

Unclassified partner reports vary little by respondent age or region. They are more common among Democratic and female respondents and substantially more common when either partner has less than a bachelor's degree. The baseline exclusion is therefore selective on observed characteristics, especially education.

The ANES partner study helps distinguish true independence from difficulty placing a partner. It observes the main respondent's own party and the partner's report of that respondent on the

Table 19: Observed Characteristics and Unclassified HCMST Partner Reports

Characteristic	Group	Couples	Unclassified	Weighted share (%)
Respondent party	Democratic	391	149	38.6
	Republican	350	107	31.4
Respondent gender	Women	397	156	39.7
	Men	344	100	30.2
Respondent education	BA	323	77	24.5
	No BA	418	179	42.3
Partner education	BA	336	75	21.8
	No BA	405	181	46.0
Respondent age	25–34	360	125	35.5
	35–44	381	131	34.9
Census region	Northeast	114	39	32.0
	Midwest	172	63	37.2
	South	276	93	34.4
	West	179	61	36.5

Notes. Unclassified means that the respondent places the partner in HCMST's combined undecided, independent, or other category. Couples with an unclassified respondent party or a refused partner-party response are excluded so that each comparison uses the 741 couples available for the reassignment exercises. Counts are unweighted; shares use the HCMST survey weight.

same seven-point grid. A self-reported partisan is placed in the middle category 12 percent of the time for Democrats and 16 percent for Republicans; a self-reported pure independent is placed there 73 percent of the time. Conditional on a middle-category proxy report, the person being classified self-reports as a partisan on the reporting partner's side 57 percent of the time when the reporting partner is a Democrat and 71 percent when the reporting partner is a Republican. Thus the HCMST partner category contains both true independents and partisans whom their partners cannot place precisely.

Two checks show that party sorting remains substantial when the baseline classification rule is changed. Reassigning unclassified HCMST partners using ANES self-reports conditional on the partner's report gives a cross-party share of 19.8 percent, compared with 16.0 percent in the baseline HCMST sample. Replacing HCMST with 222 ANES couples in which both partners report their own party gives 14.4 percent. The ANES partner study requires both partners to participate and supplies no weights correcting for partner participation. These comparisons therefore address the party-reporting concern with different sample limitations.

Re-estimating the model with the ANES-corrected HCMST table gives $\lambda = 0.133$ and $\pi = 0.931$. The correction increases the cross-party share to be fitted by assigning some unclassified partners to the party opposite the respondent. Table 23 reports the resulting complete-removal outcomes alongside the other input sensitivities.

The panel dimension of HCMST cannot refresh the couple table, because the 2020 and 2022 waves re-interview the same couples and do not re-ask party identification. It does provide descriptive evidence on relationship continuation. Among the 485 estimation couples, 269 were re-interviewed in 2020 and 225 in 2022. Dissolution is recorded for 2 of 45 cross-party couples at the 2020 follow-up and 2 of 37 at the 2022 follow-up, against 15 of 224 and 4 of 188 same-party couples, with no partner deaths in either group. Weighted by the attrition-adjusted panel weights, continuation rates are 93 and 95 percent for cross-party couples at the two follow-ups, against 92 and 98 percent for same-party couples. The relative rates differ across follow-ups. Small counts and selective follow-up response limit what this comparison can establish about differential survival.

E.5 Joint Estimation and Numerical Implementation

The estimator uses singlehood rates to choose the regional type intercepts and couple sorting to choose the conversion factor and own-pool share. At each value of (λ, π), it chooses the intercepts to reproduce each type's singlehood rate among people inside the market, solves the two-pool equilibrium, and compares the resulting couple table with HCMST.

Regional type intercepts. The regional type intercepts reproduce each type's singlehood rate among people inside the market. For fixed (λ, π) and an experimental surplus matrix S, the joint surplus in region r is

$$\Phi^r_{ij} = \alpha^r_i + \gamma^r_j + \lambda S_{ij}.$$

There are seven free intercepts per region. We set the Democratic-college male intercept $\gamma^r_{\mathrm{D,BA}}$ to zero and estimate all four female intercepts and the other three male intercepts. Adding a constant to every female intercept and subtracting it from every male intercept leaves Φ^r unchanged, so this normalization only fixes their location.

Let $s^{f,r}_i$ and $s^{m,r}_j$ denote the singlehood rates among people inside the market. Each rate equals the type's responsive single mass L^r_{gt} divided by its equilibrium population $N^{\mathrm{eq},r}_{gt}$. Same-sex-only singles are excluded from this rate and added back when reporting full-population singlehood. Write $\mu^r_{i0} = \sum_h \mu^{rh}_{i0}$ and $\mu^r_{0j} = \sum_h \mu^{rh}_{0j}$ for the model's regional single masses. We choose the intercepts to minimize the equally weighted squared differences between predicted and observed rates:

$$\min_{\alpha^r, \gamma^r : \gamma^r_{\mathrm{D,BA}}=0} \left\{ \sum_i \left(\frac{\mu^r_{i0}(\alpha^r, \gamma^r; \lambda, \pi)}{n^r_i} - s^{f,r}_i \right)^2 + \sum_j \left(\frac{\mu^r_{0j}(\alpha^r, \gamma^r; \lambda, \pi)}{m^r_j} - s^{m,r}_j \right)^2 \right\}. \tag{56}$$

Each evaluation solves the matching equilibrium at the candidate intercepts. The constructed baseline contains equal totals of modeled partnered women and men. Since outside-partnership odds are fixed, modeled partnerships for each type equal its population minus $(1 + e^{\zeta^r_{gt}})$ times its single mass. This accounting relation makes one of the eight singlehood targets redundant. Seven free intercepts reproduce all eight targets, and the calibrated outside payoffs also reproduce the baseline outside-partnership masses.

Matching equilibrium. Appendix A.6 describes the within-pool solution. We apply that solver to the type masses assigned to each pool and add the resulting counts:

$$\mu^r_{ij} = \mu^{r1}_{ij} + \mu^{r2}_{ij}.$$

The same Φ^r applies in both pools. Pool-specific partner availability and equilibrium payoffs determine how many couples each pool contributes.

Comparison at observed couple margins. We multiply each row and column of μ^r by positive factors A_i^r and B_j^r to obtain

$$\widetilde{C}_{ij}^r = A_i^r B_j^r \mu_{ij}^r, \qquad \sum_j \widetilde{C}_{ij}^r = \sum_j C_{ij}^r, \qquad \sum_i \widetilde{C}_{ij}^r = \sum_i C_{ij}^r. \tag{57}$$

Iterative row and column scaling enforces these constraints. The adjustment preserves cross-product ratios because the scaling factors cancel:

$$\frac{\widetilde{C}_{ij}^r \widetilde{C}_{i'j'}^r}{\widetilde{C}_{ij'}^r \widetilde{C}_{i'j}^r} = \frac{\mu_{ij}^r \mu_{i'j'}^r}{\mu_{ij'}^r \mu_{i'j}^r}.$$

These ratios describe association between partner types after accounting for the margins. The factors A_i^r, B_j^r only adjust the table used for this comparison; the intercepts α_i^r, γ_j^r determine participation and remain part of the structural model.

The discrepancy measure is the Poisson deviance,

$$D(C^r, \widetilde{C}^r) = 2 \sum_{i,j} \left[C_{ij}^r \log\left(\frac{C_{ij}^r}{\widetilde{C}_{ij}^r} \right) - C_{ij}^r + \widetilde{C}_{ij}^r \right], \tag{58}$$

with $0 \log(0/\widetilde{C}) = 0$. An empty observed cell requires no continuity correction. Since the adjustment matches the table total, the linear terms sum to zero. Equation (22) minimizes the sum of these regional deviances over λ and π.

Numerical checks. Every criterion evaluation starts the intercept calibration from the same values. The search over λ and π uses multiple starts and retains a converged solution. We check population balance in the equilibrium, agreement with the observed margins after scaling, and the intercept calibration at the reported optimum. The fitted singlehood rates must match their targets to numerical tolerance.

The joint objective. All seven converged starts in the joint two-parameter fit reach the same reported minimum, and a cold reevaluation reproduces its objective. Figure 18 shows a separately evaluated 25×25 grid around that estimate, recalibrating regional intercepts at each point. The grid checks the shape of the criterion near the selected estimate. Minima outside this neighborhood remain possible.

Figure 18: The joint estimation objective

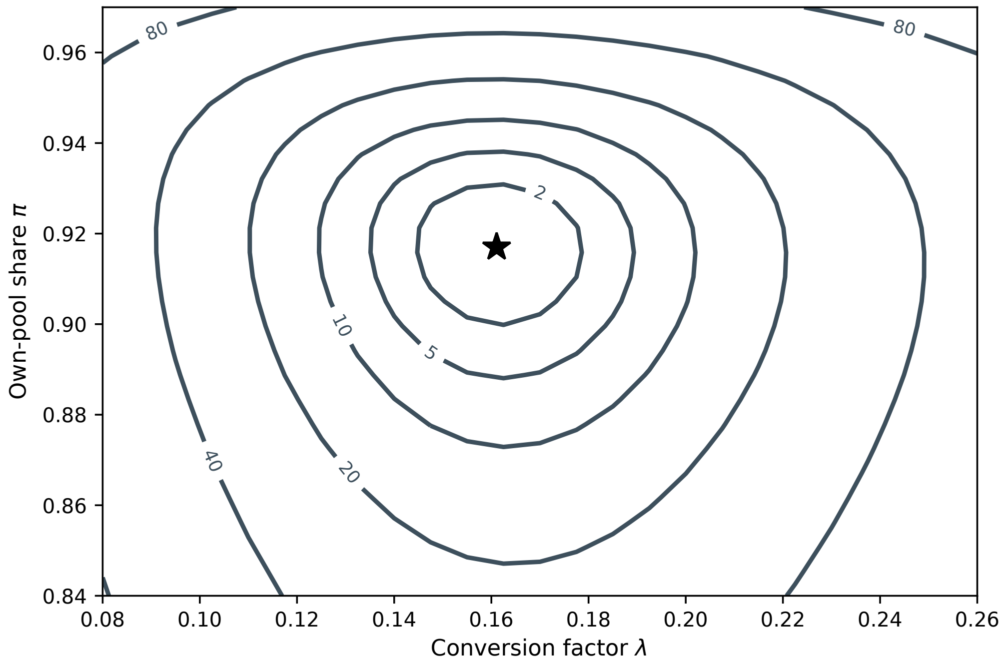

Notes. The star marks the joint estimate. Contours report the increase in Poisson deviance above its value of 40.514, with regional type intercepts recalibrated at each grid point. These contours are not confidence regions; uncertainty is obtained from the joint bootstrap.

E.6 Uncertainty

The bootstrap combines experimental uncertainty with sampling variation in CES, HCMST, ACS, and ANES. Each replication selects one joint posterior draw of the experimental preferences and choice-error scale, resamples all four external sources, reconstructs the regional populations and couple tables, and re-estimates λ, π, and the regional intercepts. The same replication then solves the counterfactual equilibria.

The external data are resampled at the respondent or couple level. Independent mean-one exponential multipliers modify the original survey weights for CES respondents, HCMST respondents, linked ACS couples, and ANES dyads. A respondent's HCMST multiplier is shared across the attraction measure and the couple-table calculations wherever applicable. Both members of an ACS couple receive the same multiplier. CES weights are normalized within survey year, preserving the equal weighting of years used in the point estimate; regional type masses, singlehood rates, and national aggregation weights are then recalculated from the resampled respondents. Each HCMST and ANES resample is shared across all uses of that source, preserving dependence between the population construction and the couple tables used for estimation.

Every replication repeats the partner-age adjustment, party classification, and reconciliation of partnered counts. The feasibility cap described in Appendix E.2 keeps each type's population partnered outside the model nonnegative. The 10 percent reconciliation limit applies to the observed-data construction; bootstrap draws that cross that limit are retained when their populations remain feasible. The cap binds in 1.5 percent of replications.

All 1,000 bootstrap replications completed, with zero failures. Each replication passes checks on population accounting, reproduction of singlehood, and the numerical equilibrium solution. The reported 95 percent intervals are the 2.5th and 97.5th percentiles of the resulting estimates. These intervals describe uncertainty under the maintained model; the specification tests below instead simulate couple tables under a fitted null model.

E.7 The One-Pool Benchmark

Unsegmented meeting provides a benchmark in which couple association has a simpler form. At $\pi = 1/2$, both pools contain half of every type. Adding their equilibria yields the one-pool matching function of Section 2.2. Taking logs gives

$$\log \mu^r_{ij} = a^r_i + b^r_j + \frac{\lambda}{2} S_{ij}, \tag{59}$$

where $a^r_i = (\alpha^r_i + \log \mu^r_{i0})/2$ and $b^r_j = (\gamma^r_j + \log \mu^r_{0j})/2$. Female- and male-type fixed effects absorb these terms in a Poisson regression of regional couple counts. Profiling out those fixed effects gives the same margin-adjusted comparison used by the joint estimator at $\pi = 1/2$.

Cross-product ratios show what identifies λ in this benchmark. For female types i, i' and male types j, j', define

$$\Delta S = S_{ij} + S_{i'j'} - S_{ij'} - S_{i'j}.$$

Then

$$2 \log \left(\frac{\mu^r_{ij} \mu^r_{i'j'}}{\mu^r_{ij'} \mu^r_{i'j}} \right) = \lambda \Delta S. \tag{60}$$

Both the type intercepts and singles terms cancel. This cancellation is specific to one pool. With two differently composed pools, regional counts add two matching functions; participation changes their relative contributions and hence their aggregate association.

Re-estimating the one-pool model gives a deviance of 139.93, compared with 40.514 for the education-based two-pool model. This fitted one-pool benchmark differs from the unsegmented-meeting counterfactual: the former re-estimates λ and the intercepts, whereas the latter changes only π to one half and holds the fitted surplus parameters fixed.

E.8 Fit and Specification Checks

The fit checks compare regional couple tables and examine the restrictions on surplus and pool assignment. Fit statistics use the tables adjusted to HCMST's female- and male-type totals. Counterfactual outcomes instead use the equilibrium couple counts and CES regional population weights. We first describe the two aggregations, then report the fit tests and regional comparisons.

For the fit statistics, national cell shares are

$$q_{ij}^{\text{fit}} = \frac{\sum_r \widetilde{C}_{ij}^r}{\sum_r \sum_{k,l} C_{kl}^r}. \tag{61}$$

This aggregation weights regions by HCMST weighted couple counts. For equilibrium outcomes, let w_r be the CES share of the national population in region r. Since each regional population is normalized to one, national equilibrium couple shares are

$$q_{ij}^{\text{eq}} = \frac{\sum_r w_r \mu_{ij}^r}{\sum_r w_r \sum_{k,l} \mu_{kl}^r}. \tag{62}$$

The cross-party shares are therefore 16.18 percent in the adjusted fit comparison and 15.41 percent in the equilibrium baseline. The two shares use different margins and aggregation weights.

Overall fit test. The fit test uses 1,000 simulated samples. It fixes the experimental surplus, constructed regional populations, source-couple counts, and observed survey and partner-age weights. Within each region, it draws couple types with probabilities proportional to the fitted margin-adjusted cell counts and assigns the observed weights to simulated couples. The simulated margins need not equal the observed margins; they are fitted anew in each replication. Re-estimating λ, π, and the regional intercepts gives a deviance at least as large as the observed 40.514 in 589 samples, yielding the plus-one tail probability $p = (589 + 1)/(1{,}000 + 1) = 0.589$. This fitted-reference check evaluates regional couple-type associations conditional on the population construction and outside-partnership assumptions.

Regional fit. Figure 8 compares the observed and fitted regional tables. The common conversion factor and own-pool share do not reproduce all regional variation in party sorting.

Party-specific pool assignment. Let a_t be the probability that party-by-education type t enters pool 1. The four probabilities are common across sexes and regions, with $1 - a_t$ assigned to pool 2. We estimate them jointly with λ, refitting regional type intercepts and the margin-adjusted couple tables throughout. The education-based model is nested at $a_{D,BA} = a_{R,BA} = \pi$ and $a_{D,NoBA} = a_{R,NoBA} = 1 - \pi$.

Assignment differs mainly by education. The fitted average education gap is 79.8 percentage points, compared with an average party gap of 0.6 point. Table 20 gives the current estimates and counterfactual outcomes. Half and full penalty removal yield similar national effects under the two assignment rules. These are point comparisons; they do not establish that the assignment probabilities are precisely estimated.

Additional meeting pools. The unrestricted three-pool specification adds four assignment parameters. Its deviance improvement over unrestricted two pools is 3.441. In 200 simulated samples

Table 20: Party-specific meeting pools: estimates and counterfactuals

	Education-based	Party-specific
Conversion factor λ	0.161	0.161
Poisson deviance	40.514	39.813
Pool 1: Democratic, BA (%)	91.7	99.5
Pool 1: Democratic, NoBA (%)	8.3	20.3
Pool 1: Republican, BA (%)	91.7	100.0
Pool 1: Republican, NoBA (%)	8.3	19.5
Half removal: singlehood (%)	45.56	45.50
Half removal: local welfare ($1,000)	4.46	4.47
Full removal: singlehood (%)	40.03	39.87
Full removal: local welfare ($1,000)	13.54	13.57

Notes. Both specifications use the current log-income preferences and responsive outside-partnership option, re-estimating conversion factors, assignment, and regional intercepts. Pool labels are aligned so pool 1 is BA-centered. Counterfactuals retain each model's own fitted parameters. Welfare is the local partner-income equivalent at $62,500 per full-population adult. The table reports point estimates without uncertainty intervals.

generated under the fitted two-pool specification, 47 yield at least that improvement after both models are refitted. The plus-one value is $48/201 = 0.239$; simulation uncertainty alone gives a 95 percent interval of [0.178, 0.300] for the tail probability. The conventional chi-squared reference is inappropriate because the third pool is not identified under the null.

With four pools the best attained fit places the conversion factor at zero and makes assignment nearly type-specific. The additional flexibility can therefore attribute almost all sorting to contact. The search does not certify a global optimum. Moreover, a partner-income welfare measure that divides by λ is undefined at zero; no dollar-welfare estimate is reported for that specification. Table 5 compares fit and parameter counts.

Regional conversion factors. The regional exercise holds the own-pool share at its common estimate and re-estimates each region's conversion factor and type intercepts. The factors are 0.096 in the Northeast, 0.181 in the Midwest, 0.157 in the South, and 0.219 in the West. Deviance falls from 40.514 to 29.696. In a separate 1,000-sample fitted-reference test, 22 deviance improvements are at least as large as the observed 10.819, giving $p = 23/1{,}001 = 0.023$. This check rescales one national experimental preference matrix; it does not identify regional preferences.

Known-parameter recovery. A 100-sample simulation generates observations under the fitted model and re-estimates its parameters. Mean estimation errors are 0.0045 for λ and 0.0036 for π, with RMSEs of 0.0136 and 0.0116. For complete penalty removal, the estimated change in full-population singlehood has mean error -0.107 percentage point and RMSE 0.380 point. The local welfare equivalent at $62,500 has mean error $-\$0.130$ thousand and RMSE $0.474 thousand. These simulations assess parameter and counterfactual recovery. They do not measure how often the bootstrap intervals cover the true values.

E.9 External Evidence

The external exercises evaluate implications of the model using data that do not determine the baseline estimates. The North Carolina comparison concerns party sorting, while the ACS comparison concerns the relationship between residential segregation and education sorting.

North Carolina party sorting. The North Carolina voter file provides a held-out statewide comparison of party sorting in a broader population than the model represents. The November 2025 file links opposite-sex adults at the same address with an age gap of at most 15 years, selecting the closest-age pair when several pairs are possible. The comparison includes adults outside ages 25–45 and does not require a shared surname. Party is taken from registration; unaffiliated voters are assigned a party when their primary participation identifies a modal party. We retain pairs in which both people are classified as Democratic or Republican. Address-based links need not identify actual partnerships, and the broader ages and registration-based party measure limit comparability with HCMST. Because the voter file does not record education, the comparison evaluates party sorting alone.

The statewide comparison uses the national log-income preferences, conversion factor, and own-pool share with North Carolina population and singlehood inputs. State type intercepts reproduce the latter; the voter-file couple counts do not enter calibration. The model predicts a cross-party share of 15.39 percent, compared with 17.52 percent in the voter file. Across the four party-pairing cells, the observed and predicted shares have correlation 0.990 and RMSE 2.55 percentage points. Because singlehood is calibrated locally, this comparison evaluates only party sorting among couples.

Residential segregation and education sorting. The ACS shows that education sorting rises with residential separation at both the state and commuting-zone levels. For residential subareas b within area a, let B_{ab} and N_{ab} be the numbers of college and non-college adults aged 25–45, and let B_a and N_a be their area totals. The dissimilarity index is

$$D_a^{\text{res}} = \frac{1}{2} \sum_b \left| \frac{B_{ab}}{B_a} - \frac{N_{ab}}{N_a} \right| .$$

The index is zero when the two groups have the same residential distribution and one when they occupy separate subareas (Duncan and Duncan, 1955). The state analysis measures their distribution across public use microdata areas within each state. The commuting-zone analysis allocates those areas to commuting zones using the Dorn crosswalk.

The education-sorting measure removes the same-education share expected from each place's education composition. We calculate the random-matching benchmark using the local shares of partnered women and men who are college graduates, then subtract it from the observed same-education couple share. The calculation uses co-resident couples in which both partners are aged 25–45 and education is observed for both.

Residential separation is strongly associated with the education-sorting measure under both geographic definitions. Couples-weighted regressions use heteroskedasticity-robust standard errors. Across 51 states, a one-standard-deviation increase in the dissimilarity index is associated with 1.7 percentage points more education sorting ($t = 7.9$). Across 660 commuting zones, the corresponding increase is 3.1 percentage points ($t = 24.5$).

The ACS comparison supports education-segmented contact but does not isolate its causal effect. Residential separation may also covary with preferences or other local characteristics. Public use microdata areas are substantially larger than neighborhoods, so the dissimilarity index measures separation at a coarser geographic level than many couples' meeting environments.

E.10 Counterfactual Identification and Sensitivity

A counterfactual that removes one surplus component requires more structure than a fitted baseline joint-surplus matrix. To see the distinction, choose arbitrary constants c_i and d_j and define

$$S^{p,*}_{ij} = S^p_{ij} + c_i + d_j, \qquad \alpha^{r,*}_i = \alpha^r_i - \lambda c_i, \qquad \gamma^{r,*}_j = \gamma^r_j - \lambda d_j.$$

These changes leave the baseline Φ^r_{ij} unchanged. Removing $S^{p,*}$ while holding the transformed intercepts fixed, however, changes surplus by an additional $-\lambda(c_i + d_j)$ relative to removing S^p. The same baseline matrix can therefore support different depolarization counterfactuals when the party component is left unrestricted.

The experiment and the mapping in Section 5.1 fix this decomposition. They specify the type-specific cross-party terms and set the party component to zero for same-party pairings. A common conversion factor preserves the experiment's relative valuation of party and education. A reduction replaces λS^p with $\rho\lambda S^p$ while retaining the type intercepts, education surplus, type masses, and pool assignments. The historical exercise in Section 8.1 sets $\rho = 0.6$, the half-penalty benchmark sets $\rho = 1/2$, and complete removal sets $\rho = 0$. The baseline normalization and the assumption that the remaining components stay fixed are part of the counterfactual's interpretation.

Figure 19 traces equilibrium outcomes across the full range of common reductions. The response becomes steeper as more of the cross-party penalty is removed, reflecting both logit choice and equilibrium adjustment.

Figure 19: Effects of Reducing the Cross-Party Penalty

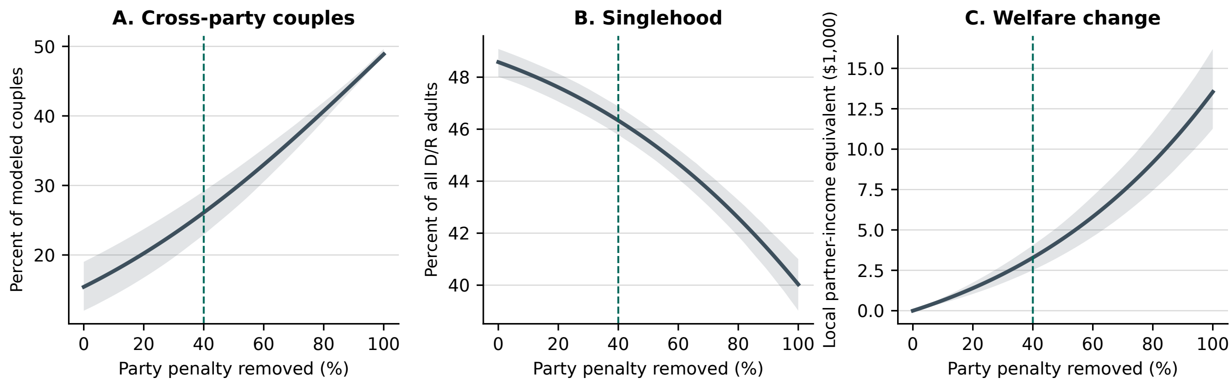

Notes. The horizontal axis is the percentage of the cross-party penalty removed from every type. Curves show point estimates, and bands show pointwise 95 percent intervals from 1,000 joint-bootstrap replications. The dashed line marks the 40 percent aversion reduction. Composition, meeting opportunities, and outside-partnership payoffs remain fixed. Welfare is a local partner-income equivalent at $62,500 per full-population adult.

Table 21 reports estimates for common penalty reductions, residual-label removal, and unsegmented meeting. The 40 percent row changes aversion alone and holds current composition fixed; it is not the combined historical exercise in the main text. The full-removal and meeting rows reproduce the corresponding main-text partnership outcomes.

Table 21: National counterfactual outcomes and uncertainty

	Cross-party couples (%)	Singlehood (%)	Local welfare ($1,000)
40 percent reduction	26.11	46.33	3.28
	[22.93, 29.19]	[45.78, 46.88]	[2.51, 4.05]
Penalty halved	29.43	45.56	4.46
	[26.55, 32.15]	[45.00, 46.10]	[3.46, 5.44]
Entire penalty removed	48.87	40.03	13.54
	[48.52, 49.63]	[39.02, 41.00]	[11.27, 16.20]
Residual label removed	25.72	46.42	3.15
	[21.07, 31.43]	[45.06, 47.28]	[1.91, 5.26]
Unsegmented meeting	15.44	42.47	7.50
	[12.04, 19.09]	[40.85, 43.81]	[5.09, 10.57]

Notes. Point estimates with 95 percent percentile intervals from 1,000 joint-bootstrap replications below. Couple shares are conditional on modeled partnerships; singlehood includes the fixed same-sex-only singles. Welfare is the local partner-income equivalent at $62,500, averaged over the full D/R population and assigning zero change to the fixed groups. Residual-label intervals also draw the within-module label shares and resample beliefs. The pointwise intervals condition on each counterfactual's assumptions.

The mechanism-based counterfactual maps a within-module decomposition onto the structural party penalty. For respondent party p, let L_p denote the residual party-label penalty and I_p the penalty associated with perceived Democratic–Republican differences on the four issue positions. We calculate I_p by multiplying each position's value in the mechanism choices by the respondent's perceived party difference on that position, then summing across issues. The residual-label share is

$$k_p = \frac{L_p}{L_p + I_p}.$$

The estimates are $k_D = 27.1$ percent [19.5, 34.5] and $k_R = 66.8$ percent [38.5, 113.4]. The Republican interval extends above 100 percent because the Republican mechanism sample is small and its issue-position estimates are imprecise; the bootstrap retains rather than truncates those draws.

The residual-label counterfactual removes the estimated label share separately from each partner's side of cross-party surplus. For a cross-party pairing of female type i and male type j, the party component becomes

$$-\big[(1 - k_{p(i)})\psi_i^f + (1 - k_{p(j)})\psi_j^m\big].$$

The share associated with the four issue positions remains fixed, as do the conversion factor and all other model inputs. Transferring the shares assumes that the mechanism module's decomposition applies to the main experiment's penalty. The exercise transfers proportions across experiments; it does not require equality of their dollar-valued penalties, and it does not validate that equality. Uncertainty combines the mechanism posterior and respondent-level belief resampling with all 1,000 structural bootstrap replications. Table 21 reports national intervals. Because k_D and k_R differ, this counterfactual does not lie on the common-reduction curve in Figure 19.

Table 22: Local welfare equivalents by respondent type

	Men	Women
Penalty halved		
Democratic, BA	4.42 [2.61, 6.29]	3.87 [2.64, 5.17]
Democratic, NoBA	2.54 [1.49, 3.72]	3.65 [2.52, 4.85]
Republican, BA	10.86 [7.28, 14.90]	5.75 [2.66, 8.68]
Republican, NoBA	4.89 [3.28, 6.65]	4.01 [2.09, 6.12]
Entire penalty removed		
Democratic, BA	12.71 [7.81, 17.54]	10.37 [7.15, 13.90]
Democratic, NoBA	9.09 [5.82, 12.55]	10.71 [7.32, 14.88]
Republican, BA	31.16 [20.49, 44.06]	17.97 [8.66, 26.96]
Republican, NoBA	14.70 [9.59, 21.06]	14.43 [8.89, 21.05]

Notes. Point estimates with 95 percent joint-bootstrap intervals. Units are thousands of annual partner-income dollars at the $62,500 reference income. Each type average includes its fixed same-sex groups with zero change. BA denotes a bachelor's degree or more. The table reports common half and full penalty removal, not residual-label removal.

Welfare is calculated within each pool before aggregation. Appendix A.5 relates pool-specific welfare to the probability of singlehood. Regional type welfare averages those values using pool-assignment probabilities; it generally differs from the negative log of the type's aggregate singlehood rate. For the fitted income scale $c = 6.5$, a normalized utility change $\Delta\mathcal{W}$ has log-income equivalent $\Delta\mathcal{W}/(c\lambda)$ and local dollar equivalent $62{,}500\,\Delta\mathcal{W}/(c\lambda)$. These transformations preserve the income unit used to estimate preferences.

In composition counterfactuals, outside-partnership payoffs and each type's fixed-group shares remain unchanged, but outside-partnership rates among the equilibrium population respond. National welfare compares population-weighted mean welfare across the two compositions; type-level changes instead compare payoffs using a specified fixed set of regional weights. A national composition effect therefore need not equal an average of within-type changes and is not an individual return to changing party or education.

E.11 Sensitivity to Population and Partner-Party Inputs

Table 23 re-estimates the model under five alternative inputs. The ANES-corrected HCMST case changes the couple table used to estimate sorting while preserving the baseline population construction. The UAS case instead changes the party association used to construct the modeled population, retaining the classified HCMST sorting table. It uses opposite-sex UAS306 couples whose two members are aged 25–45 and report their own party and education (Understanding America Study, 2020). Leaners count as partisans; non-leaning Independents remain a separate category in the population bridge. Household base weights estimate the national association, which is applied to regional CES margins. Differences in party-question wording and selection into a sample in which both partners respond remain limitations.

The remaining cases change the HCMST measure used to set aside same-sex-only singles. One uses gay or lesbian identity rather than reported attraction, one pools the attraction share across sexes, and one restricts the attraction calculation to respondents who report no current partner rather than the broader CES-like single category. In each case, existing same-sex partnerships remain fixed using the same region-by-sex-by-education ACS rates, with BA and non-BA cells kept in their original order.

Table 23: Input sensitivities under complete removal of the cross-party penalty

Input specification	λ	π	Singlehood (%)	Local welfare ($1,000)
Baseline	0.161	0.917	40.03	13.54
ANES-corrected HCMST table	0.133	0.931	40.90	14.57
UAS population construction	0.161	0.915	39.77	13.98
Sexual identity rather than attraction	0.161	0.917	40.09	13.53
Common attraction share across sexes	0.161	0.917	40.03	13.55
Strictly unpartnered attraction sample	0.161	0.917	40.10	13.53

Notes. Point estimates, with parameters re-estimated for each input construction. All rows use log-income preferences and a responsive outside-partnership option. Singlehood is a counterfactual level in the full D/R population. Welfare is the change from each specification's own baseline, in thousands of local partner-income dollars at $62,500. The table does not report uncertainty intervals for these five input comparisons.

These alternatives leave complete-removal singlehood between 39.77 and 40.90 percent, compared with 40.03 percent in the baseline. They do not resolve uncertainty about how outside partners would respond in a fully modeled wider market: the baseline holds their payoff fixed and does not impose a market-clearing condition on their supply.

F Measuring Marriage and Cohabitation

The Cooperative Election Study (CES) does not fully measure the paper's marriage-or-cohabitation outcome: cohabiting respondents need not report a domestic partnership. We use the American Community Survey (ACS) to measure partnership status and the CES to retain population composition and party differences in marriage. This correction changes the masses used to fit the model.

F.1 A common partnership measure

The ACS identifies a cohabiting couple through an unmarried partner of the householder, counting both members. Married people without such a link are married, including those whose spouse lives elsewhere; everyone else is neither married nor cohabiting. An unmarried-partner link takes precedence over reported marital status, so the categories do not overlap. The measure misses cohabiting subfamilies.

The ACS target covers household adults aged 25–45 in 2017–2021 and uses person weights. It classifies 50.81 percent as married, 10.72 percent as cohabiting, and 38.48 percent as neither. The CES classifies only 49.70 percent as married or in a domestic partnership. The CES figure includes every party classification; the ACS does not measure party. The ACS extract does not identify citizenship, and neither sample is restricted to citizens. Treating the CES domestic-partnership category as all cohabitation would therefore substantially overstate singlehood.

F.2 Combining ACS partnership status with CES party

Within region-by-sex-by-education-by-age cells, the adjustment preserves CES party masses and marriage differences while matching ACS status rates. Democrats, Republicans, independents, and people without a usable party classification all enter; we select Democrats and Republicans afterward.

The CES determines party differences in marriage. Let m_{cp} be its married share for party group p in cell c, and let M_c and H_c be the ACS married and cohabiting shares. One shift δ_c, common to parties in the cell, satisfies

$$\widetilde{m}_{cp} = \operatorname{logit}^{-1}(\operatorname{logit}(m_{cp}) + \delta_c), \qquad \frac{\sum_p n_{cp}\widetilde{m}_{cp}}{\sum_p n_{cp}} = M_c, \tag{63}$$

where n_{cp} is the CES population mass. The shift preserves each party group's mass and its marriage log-odds difference from other groups; the ACS supplies no party-specific rate.

Among corrected unmarried people, we assign the ACS cohabiting fraction $H_c/(1 - M_c)$ equally across parties within a cell. The resulting cohabiting and neither shares are

$$\widetilde{h}_{cp} = (1 - \widetilde{m}_{cp})\frac{H_c}{1 - M_c}, \qquad \widetilde{s}_{cp} = (1 - \widetilde{m}_{cp})\left(1 - \frac{H_c}{1 - M_c}\right). \tag{64}$$

The corrected shares match the ACS targets in each calibration band, while party masses remain unchanged. The adjustment preserves CES party differences in marriage. The reported domestic-partnership category does not determine the cohabitation adjustment. Among Democrats and Republicans, 62.70 percent are married or cohabiting.

We combine ages 40–44 and 45 for status calibration because several age-45 CES cells are sparse. Their ACS target uses combined weighted counts; original age masses and partner-age eligibility remain separate. The largest deviation from an original-age ACS target is 3.45 percentage points. Restricting CES to 2017–2021, while holding type masses fixed, yields a corrected Democratic/Republican partnered share of 62.28 percent, versus 62.70 percent with CES through 2024.

F.3 From partnership status to the matching population

Existing same-sex partnerships and singles attracted only to the same sex remain outside the woman–man market. Among HCMST singles, 6.4 percent of men and 3.1 percent of women report attraction only to the same sex. We apply these rates to the corrected neither-married-nor-cohabiting population across region, party, and education because finer HCMST cells are small; people attracted to both sexes remain eligible. ACS partner links identify existing same-sex partnerships separately among married and cohabiting adults. We allow those rates to vary by region, sex, and education but apply them equally across parties within each cell.

The remaining partnerships are divided between the modeled market and an outside option. ACS links give the joint probability that a partner is of the opposite sex and aged 25–45, separately by partnership status and demographics; no second opposite-sex adjustment is applied. For married people whose spouse lives elsewhere, we impute partner age and sex, including the same-sex rate, from otherwise similar linked couples. HCMST and ANES then identify the share of eligible opposite-sex partnerships joining two modeled Democratic/Republican types. Combining these rates assumes partner party is independent of partner sex and age within the demographic cells used for the adjustment. Other opposite-sex partnerships, including those outside the modeled ages or party groups, enter the outside option.

Each region-by-sex-by-party-by-education type has five exhaustive states:

$$N = L^{SS} + L^{OS} + C^{SS} + C + O^{OS}, \qquad N^R = L^{OS} + C + O^{OS}. \tag{65}$$

Here L^{SS} is same-sex-only singles, L^{OS} other singles, C^{SS} existing same-sex partnerships, C modeled opposite-sex partnerships, and O^{OS} other opposite-sex partnerships. Only N^R enters the matching equilibrium. Its outside payoff, calibrated at baseline from $\log(O^{OS}/L^{OS})$, stays fixed, but the outside count can change. The fixed groups are added back to full-population counts and assigned zero modeled welfare change. Modeled partnerships must fit within available opposite-sex partnership mass; this constraint does not bind.

The bootstrap carries uncertainty in these measurements into the model estimates. Each replication draws ACS households jointly for partnership status and partner links, resamples CES, HCMST, and ANES, draws from the experimental posterior, and refits the model. Halving the cross-party penalty reduces full-Democratic/Republican singlehood by 2.734 percentage points. The reduction is 2.746 with CES through 2021, 2.724 when calibrating overall partnership rather than marriage odds, and 2.628 when spouse-absent married people are ineligible. All use the same counterfactual and population denominator.

F.4 The historical comparison

The historical comparison uses marriage plus observed cohabitation, rather than legal marriage alone. In the March Current Population Survey, we count married adults and both members of householder-linked unmarried couples, restrict ages to 25–45, and use person weights. The partnered share falls from 66.75 percent in 1995 to 61.67 percent in 2025, a decline of 5.08 percentage points. This measure does not recover cohabiting couples missed by the relationship code.

The specified aversion counterfactual equals 29.8 percent of the 1995–2025 measured decline. Reducing the estimated cross-party penalty by 40 percent raises partnership by 2.04 percentage points among Democrats and Republicans, or 1.51 points among all adults using their CES population share. The counterfactual holds current meeting opportunities and party composition fixed, and the aversion change is externally motivated. This ratio quantifies a conditional response, leaving other causes of the historical partnership decline outside the exercise. Later direct cohabitation questions cannot reveal how many cohabiting subfamilies the CPS missed in 1995.