Candidate positioning or limited choice? Voter–candidate fit in the 2024 US election

Samir Varma
VS Asset Management, LLC

Accepted Manuscript. This is the author's accepted manuscript (the version revised after peer review and accepted for publication, 22 October 2026), including its online appendix. It is not the publisher's typeset Version of Record.

The Version of Record of this manuscript has been published and is available in *Journal of Elections, Public Opinion and Parties* 6 November 2026 https://www.tandfonline.com/doi/10.1080/17457289.2026.2755744.

Please cite the Version of Record: Varma, Samir. 2026. "Candidate positioning or limited choice? Voter–candidate fit in the 2024 US election." *Journal of Elections, Public Opinion and Parties*. https://doi.org/10.1080/17457289.2026.2755744.

Replication package: https://doi.org/10.5281/zenodo.23358049.

ARTICLE

Candidate Positioning or Limited Choice? Voter–Candidate Fit in the 2024 U.S. Election

Samir Varma

VS Asset Management, LLC

ARTICLE HISTORY
Compiled October 20, 2026

ABSTRACT
When voters say neither major-party candidate represents them, the usual diagnosis is that the candidates stand in the wrong places. There is a second possibility: voters combine their views in so many ways that no pair of candidates could fit them much better. I test both with the 2024 ANES, using seven issues on which respondents placed themselves and both candidates. Setting Harris and Trump aside, I search for the two platforms that would sit closest to the same voters. Harris and Trump, as voters placed them, come close: the best pair found removes only 8.8% of the mismatch voters faced, and moving Trump alone gets most of that. More than nine-tenths of the mismatch remains, and it lies in the number of options. A centrist third candidate, the usual remedy, removes 18.4%. Letting all three options move does better: the best three-option menu removes a further 15.7% of what the best pair leaves, about 23% in all. Its middle option is cross-cutting: relatively conservative on defense and liberal on the environment. Under plurality, electing that middle option would fit voters better than electing either outer one. Better-placed candidates would help little; the larger problem is having only two.

KEYWORDS
voter–candidate congruence; candidate perceptions; candidate positioning; limited party choice; American elections

1. Two Diagnoses of Poor Voter–Candidate Fit

In the 2024 American National Election Studies (ANES), about one respondent in eleven disliked both Harris and Trump.[1] The complaint that neither candidate represents ordinary voters usually arrives with a diagnosis: the candidates stand in the wrong places. On that account, the parties offer platforms farther from their voters than they need to be, and better-placed candidates would close the gap. The remedies follow: different nominees, different primaries, a moderate third candidate.

A second diagnosis is less often stated. If voters' views lined up along a single left–right line, the main question would be where the two candidates stand. American voters' views do not line up so neatly. Consider a voter who favors both higher defense spending and stronger environmental protection: a conventionally conservative platform matches the first view and misses the second, and a conventionally liberal

CONTACT Samir Varma. Email: samir@vsasset.com. ORCID: 0000-0001-6437-5395

[1] Weighted share, using post-election weights, of the 4,686 post-election respondents who rated both candidates on the 0–10 like–dislike scale, counting a rating of 4 or below as dislike: 9.1%. Section 7 uses the same definition.

platform does the reverse. When views combine in many ways across issues, any two platforms leave many voters far from both, wherever the platforms stand. On this account, the constraint is the number of options, and better-placed candidates can do little about it. The remedies differ too: more options, and electoral rules that decide how a richer menu becomes an officeholder.

Both diagnoses predict the same symptom, voters far from both candidates, and ordinary measures of voter–candidate congruence combine them in a single distance. Which one applies matters for anyone who wants to reduce the mismatch. This paper separates them for the 2024 U.S. electorate.

The method is a thought experiment run on survey data. Set Harris and Trump aside and ask: if any two platforms could be designed for these voters, how close could they come? I place voters on seven survey-weighted issue scales, and I place Harris and Trump on the same scales at the average positions where respondents placed them; throughout, "Harris" and "Trump" refer to these average placements. I then search for the pair of platforms, or policy bundles (one position on every issue), that sits closest to those same voters, keeping the voters and the measure of distance fixed. The difference in mismatch between Harris and Trump and the best-found pair measures the room for repositioning, which I call *positioning headroom*. The mismatch that remains under the best-found pair is what two options leave undone in this electorate. I then ask what a third platform would add.

Harris and Trump, as voters saw them, were nearly as well positioned as any pair the search could find. For the 3,487 respondents who placed themselves on all seven issues, replacing Harris and Trump with the best-found pair would remove 8.8% of the mismatch voters faced; more than nine-tenths would remain. Across the weighting, scaling, and issue-omission checks examined here, Harris and Trump produce 7% to 11% more mismatch than the best-found pair. Better-placed nominees, the first remedy, would therefore gain little. The small room for repositioning lies mostly on one side. Trump is more conservative than the best-fitting conservative platform on every issue, most visibly on the environment, abortion, and health insurance, and moving his location alone removes more mismatch than moving Harris's (7.2% against 3.4%). Moving Trump inward improves fit, but even optimally placed platforms must each represent large groups of voters who disagree with one another.

Most of the mismatch therefore lies in the size of the menu. Voters share a broad left–right pattern, but their positions also vary across issues in ways that pattern does not capture. Even the best-found pair removes only about two-fifths of the mismatch left by a single platform at the electorate's average position.

What would a third option add? The usual answer is a centrist. Adding a platform at the electorate's average position on every issue to Harris and Trump comes close to the best single addition found, and it removes about twice as much mismatch as repositioning both candidates. Yet it removes only 18.4% of observed mismatch and leaves four-fifths in place. Letting all three platforms move does better, removing about 23%, and that menu looks different from Harris, Trump, and a moderate. Compared with the best-found pair, the liberal and conservative platforms move outward, and the middle platform pairs a relatively conservative defense position with a relatively liberal environmental position. Its nearest voters are less firmly partisan than those nearest either outer platform. The missing option is moderate on average and cross-cutting issue by issue.

A better menu still has to elect someone. Plurality selects one officeholder, so the improvement in available options does not settle how well the winner fits the electorate. If the middle platform wins, the officeholder's mismatch is within 1% of the best single

platform's; if an outer platform wins, it is 72% or 90% higher. Which platform wins depends on electoral behavior and coordination, which these conditional comparisons leave open.

Research on ideological congruence measures the distance between citizens and the political options they face (Huber and Powell 1994; Golder and Stramski 2010; Ferland 2016). That distance mixes two questions that separate literatures address. Where a given number of options stand is the subject of theories of party competition from Downs (1957) through Adams, Merrill, and Grofman (2005). How many options the system supplies is the subject of work on electoral coordination (Cox 1997) and on its consequences for congruence across electoral systems (Powell 2000). This paper splits the mismatch a congruence measure reports into the part better-placed candidates could remove and the part only more options could. Where cross-national work ranks party systems against one another, here one electorate is measured against the best platforms it could have had, which complements work on perceived underrepresentation (Feitosa, Pilet, and Talukder 2025) and on how citizens infer candidate locations (Balogh et al. 2025; Kurella and Rapp 2024).

The two diagnoses do not share the blame equally. The benchmark describes what freely placed platforms could achieve in this electorate; competing parties need not offer them. It puts a number on the argument of Ahler and Broockman (2018): because citizens combine issue positions in their own ways, no small set of platforms can represent them closely. The cross-cutting middle platform bears on the debate over whether voters between the parties hold genuinely moderate or cross-pressured positions (Fowler et al. 2023), and matches evidence from multiparty systems that voters with cross-cutting positions are the ones the existing supply serves worst (Lefkofridi, Wagner, and Willmann 2014; Hillen and Steiner 2020). When voters say that neither candidate represents them, the evidence here places most of the problem in the number of options: two platforms leave most of the measured mismatch in place, wherever they stand.

Section 2 defines the benchmarks. Section 3 constructs the matched ANES policy space; Sections 4 and 5 compare Harris and Trump with the two- and three-platform benchmarks and examine winner fit; Sections 6 and 7 turn to the cross-year comparison and dissatisfaction; Section 8 gathers the interpretation and limits, and Section 9 concludes.

2. Benchmarking Positioning and Two-Bundle Fit

Let $\theta_i \in \mathbb{R}^p$ denote voter i's position in a common policy space, let s_i denote the voter's survey weight, and let $B = (b_1, \ldots, b_K) \in (\mathbb{R}^p)^K$ be an indexed menu of K policy bundles. A policy bundle is one point in this space: a vector giving one position on every measured issue. Bundle locations may coincide, so a menu with $K+1$ entries includes every menu available with K. I measure nearest-menu mismatch by weighted average squared distance:

$$D(B) = \frac{\sum_i s_i \min_{1 \le j \le K} \|\theta_i - b_j\|^2}{\sum_i s_i}. \tag{1}$$

The best mismatch achievable with K bundles is

$$D_K^* = \min_{B \in (\mathbb{R}^p)^K} D(B). \tag{2}$$

Every comparison below holds the voter sample and distance metric fixed.

The first comparison isolates positioning. Let $\widetilde{D}_2$ be the mismatch generated by average perceived Harris and Trump locations. The *positioning ratio* is

$$\rho_2 = \frac{\widetilde{D}_2}{D_2^*}. \tag{3}$$

A ratio near one leaves little positioning headroom between Harris and Trump, at their average perceived locations, and the best-found pair. The share of observed mismatch removed by replacing Harris and Trump with the best-found pair is $1 - 1/\rho_2$.

Two further comparisons measure fit at different menu sizes. The *retained-mismatch share*, $c_2 = D_2^*/D_1^*$, gives the fraction of best-single-platform mismatch left by the best-found pair. The *third-bundle gain*, $g_3 = 1 - D_3^*/D_2^*$, gives the fraction of best-pair mismatch removed by the best-found three-bundle menu. Together, ρ_2, c_2, and g_3 separate positioning headroom from the fit gained by expanding the menu.

The amount left after two bundles depends on the geometry of voter disagreement. I summarize that geometry with effective dimensionality,

$$d_{\text{eff}} = \frac{(\operatorname{tr}\Sigma)^2}{\operatorname{tr}(\Sigma^2)}, \tag{4}$$

where tr denotes matrix trace and Σ is the survey-weighted covariance matrix of standardized issue positions. In the seven-item application, effective dimensionality ranges from one to seven. It equals one when variation lies on one axis and rises as disagreement spreads across additional directions (Del Giudice 2021; Koedam, Binding, and Steenbergen 2025). Effective dimensionality describes the spread of disagreement; the D_K^* comparisons measure how closely the bundles fit it.

Allowing a third location weakly lowers the global optimum because the two-bundle solution remains available. In the application, the best-found $K = 3$ objective is lower than the best-found $K = 2$ objective. The magnitude of the gain, the location of the third bundle, the movement of the other two bundles, and the voters nearest each bundle are empirical. Computationally, equations (1)–(2) form the weighted K-means problem (MacQueen 1967; Gersho and Gray 1992). The theoretical notation denotes a global minimum. In the empirical results, "best-found" means the lowest numerical objective obtained across the stated independent starts.

3. Data and Measurement

3.1. *The matched ANES policy space*

The analysis uses the August 8, 2025 full release of the 2024 ANES Time Series Study (American National Election Studies 2025). The headline policy space contains the seven pre-election issues for which the survey asks respondents to place themselves, Kamala Harris, and Donald Trump on matched 1–7 scales: spending versus government services, defense spending, government health insurance, abortion, guaranteed jobs

and income, aid to Black Americans, and environmental regulation. All scales are oriented so that higher values represent the more conservative endpoint. The May 19, 2026 re-release changes labels for the political-attention variable V241004 and moves one respondent's post-election codes from one negative missing value to another. The workflow uses V241004's documented numeric ordering and treats all negative codes as missing, leaving analytic inputs and inclusion decisions unchanged (American National Election Studies 2026).

The headline sample contains respondents with valid self-placements on all seven issues and a positive ANES pre-election full-sample weight, yielding $N = 3{,}487$. Hereafter, *electorate* in a headline estimate means this survey-weighted complete self-placement sample. Candidate locations are survey-weighted item-by-item means of the sample's Harris and Trump placements. Weighted valid shares range from 97.5% to 99.5% across the fourteen candidate–issue placements. These means are the Harris and Trump locations used in the headline comparison.

ANES questions about immigration, border policy, and trade use different formats, timing, and political objects from the matched placement scales. The common policy space therefore contains the seven matched issues. The scope subsection in Section 8 returns to the coverage implied by that choice.

3.2. *Survey weights, issue weights, and missingness*

Four weighting operations enter the analysis. First, the headline geometry, candidate means, and benchmarks use the ANES pre-election full-sample survey weight; the attitude models use the post-election full-sample weight. Second, each issue is standardized using its survey-weighted mean and standard deviation. Third, the baseline metric gives each standardized issue equal weight in squared distance. Fourth, a separate swing-weighted targeting exercise in the Online Appendix multiplies survey weights by a model-based pivotality term (Lindbeck and Weibull 1987). The targeting exercise supplies a theoretical comparison alongside the baseline and the issue-weight perturbations.

The ANES importance batteries cover a different set of objects from the seven matched placement items. I therefore report respondent-unweighted estimation, raw scales, seven leave-one-issue-out estimates, and four double-weight perturbations. The perturbations emphasize abortion; spending/services, health insurance, and jobs together; racial equality; or the environment. These specifications show whether the ordering depends on any one issue or metric choice.

Missingness is selective. Across the seven self-placement items, the survey-weighted share coded *haven't thought much* ranges from 6.3% for abortion to 15.7% for the environment. Complete cases account for 63.3% of positive survey weight. They score higher than incomplete cases on political attention (2.63 versus 2.05 on a 0–4 scale) and campaign interest (1.42 versus 1.08 on a 0–2 scale). Pure independents constitute 4.9% of complete cases and 10.1% of incomplete cases. The Online Appendix separates code 99 from refusals and don't-know responses and constructs pairwise-available and partial-case comparisons. Section 4.3 reports how those broader constructions change the results.

3.3. *Perceptions, outcomes, and computation*

Perceived candidate positions are the relevant candidate locations for voter–candidate congruence because voters respond to where they believe the options stand (Balogh et al. 2025; Kurella and Rapp 2024). Partisans may project their preferences onto a favored candidate or place the opposing candidate farther away (Merrill, Grofman, and Adams 2001), and campaign persuasion may move both self-placements and candidate placements. The headline estimate averages perceived locations across the complete self-placement sample. Additional constructions use Democratic identifiers, Republican identifiers, co-partisans, out-partisans, attentive respondents, and respondent-specific complete perceptions.

The post-election models use dual dislike of Harris and Trump and dissatisfaction with the available electoral choice. They include party identification, ideology, age, education, race, political attention, and campaign interest. Pre-existing alienation can affect perceptions, self-placement, and later attitudes, so these models describe associations.

All benchmark solutions are recomputed with weighted K-means. The production run uses 200 independent initializations for each K and retains the lowest numerical objective. Confidence intervals are percentile intervals from 400 Rao–Wu rescaled bootstrap replicates (Rao and Wu 1988) over the ANES full-sample variance primary sampling units (PSUs; V240108c) within variance strata (V240108d). Every replicate re-estimates survey-weighted standardization, candidate means, and the K=2 and K=3 benchmarks using 20 independent starts per K. The attitude-model standard errors use the same ANES strata and PSUs with the post-election weight and condition on the estimated menus. The Online Appendix reports the variable manifest, missingness construction, perception diagnostics, full sensitivity grid, and optimization runs.

4. How Far Are Harris and Trump from the Best-Found Pair?

4.1. *The best-found pair leaves more than half of single-platform mismatch*

Table 1 reports the basic geometry. On the survey-weighted seven-item baseline, effective dimensionality is 2.74 and the first principal component accounts for 57.6% of total standardized variance. The electorate has a dominant first direction alongside consequential variation on others, in line with estimates that American policy preferences require more than one dimension (Treier and Hillygus 2009; Broockman 2016). The best-found two-bundle menu retains 59.0% of the mismatch left by the best single platform, with a 95% design-bootstrap interval from 57.0% to 61.0% (Online Appendix Table OA5). Two bundles remove 41.0% of single-platform mismatch and leave 59.0% in place.

Table 2 holds the menu size fixed at two. Average perceived Harris and Trump locations produce weighted mismatch of 4.531, compared with 4.132 under the best-found pair. The resulting positioning ratio is 1.097, with a 95% design-bootstrap interval reported in the table. Replacing Harris and Trump with the best-found pair would remove 8.8% of observed mismatch. Using average placements answers the question of where one pair of candidates, facing all voters, stands relative to the best pair. Giving each respondent that respondent's own placements of Harris and Trump yields a larger removable share of 17.3% (Section 4.2), but that comparison also incorporates heterogeneity and error in perceived candidate locations and is therefore a different estimand.

Table 1. Preference Geometry in the 2024 ANES Complete-Self-Placement Sample

Specification	N	d_{eff}	First-PC share	D_2^*/D_1^*
2024 weighted 7-item matched baseline	3,487	2.74	0.58	0.590
2024 weighted 6-item sensitivity comparison	3,527	2.60	0.59	0.585
2024 unweighted 7-item matched baseline	3,487	2.63	0.59	0.576
2024 unweighted 6-item sensitivity comparison	3,527	2.50	0.60	0.573

Notes: The headline specification uses complete voter self-placements on seven matched issue items; candidate coordinates use itemwise-available Harris and Trump placements. Weighted rows use the ANES pre-election full-sample weight. The 6-item rows omit abortion and retain the other six matched issues; their sample contains the 3,527 respondents with complete self-placements on those six items, which is why their N exceeds the seven-item baseline. Every standardized issue receives equal weight in the baseline distance metric.

Under classical independent response error, the ratio based on average placements is likely attenuated toward one (Section 8.1), so the 8.8% estimate may be conservative under that error model. Under every construction, repositioning leaves most observed mismatch in place. Candidate locations estimated among highly attentive respondents produce a ratio of 1.129; partisan coalition centroids provide a low-ratio empirical comparison of 1.007.

Table 2. Average Perceived Candidate Locations Versus the Best-Found Two-Bundle Benchmark

Measure	Weighted mismatch	Positioning ratio	Notes
Average perceived candidate locations	4.53	1.10	95% CI for ratio: [1.080, 1.123]
Candidate placements (high-attention respondents)	4.66	1.13	Candidate placements estimated among attentive respondents
Partisan coalition centroids	4.16	1.01	Coalition-centroid comparison

Notes: The positioning ratio is observed nearest-candidate mismatch divided by the best-found two-bundle benchmark in the survey-weighted, standardized, equal-issue seven-item space. The main row uses average perceived major-party candidate locations. The attentive-respondent row re-estimates those locations among respondents who report paying attention to politics always or most of the time.

4.2. *Where Harris and Trump differ from the best-found pair*

Figure 1 shows the comparison on the original oriented scales. The perceived Harris location tracks the lower best bundle fairly closely. Relative to that bundle, Harris is somewhat more conservative on spending and health insurance, more liberal on defense and aid to Black Americans, and closer to the voter median on the environmental scale. The perceived Trump location is more conservative than the upper best bundle on every issue. The largest differences occur on environment (5.92 versus 4.56), abortion (5.57 versus 4.68), and health insurance (5.82 versus 5.05). The upper best bundle is the proximity-minimizing centroid of the voters nearest it. The aggregate positioning headroom is small and concentrated on a few issues. Single-candidate counterfactuals make the asymmetry exact: holding perceived Harris fixed and freeing only the Trump location removes 7.2% of observed mismatch, while the reverse removes 3.4% (Online Appendix Table OA24). Most of the limited positioning headroom lies on the Trump side.

Whose perceptions define the candidates affects the level of mismatch. Each construction of common Harris and Trump locations is evaluated against the same 3,487

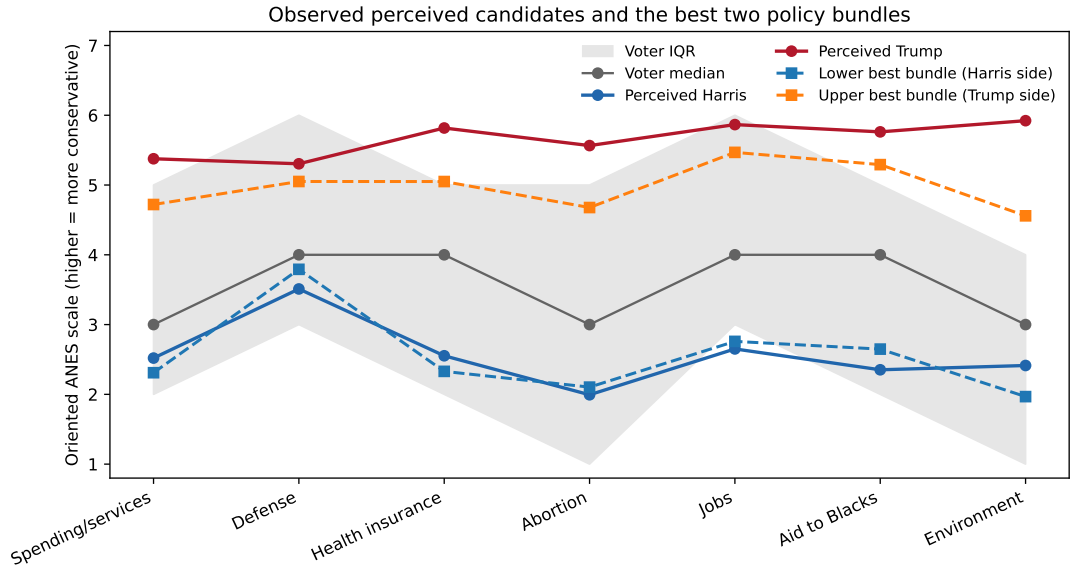


Figure 1. Perceived candidate locations and the best-found two policy bundles in the survey-weighted seven-item space. The grey band shows the weighted interquartile range of voter self-placements and the grey line shows the weighted median. Candidate locations are weighted means of perceptions; best-found bundles are weighted K-means centroids converted back to the oriented 1–7 scales. Higher values indicate the more conservative endpoint. Candidate design-bootstrap intervals appear in the Online Appendix; optimizer stability for the bundle coordinates appears in Table 3.

baseline voters while changing the partisan audience that supplies the candidate means. A menu built from each candidate's co-partisans gives a positioning ratio of 1.070. Menus based on Democratic identifiers' perceptions of both candidates or Republican identifiers' perceptions of both give 1.179 and 1.165. The out-partisan construction gives 1.408. These differences measure the sensitivity of perceived positioning to partisan audience.

The estimate based on average placements also differs from fit measured against each respondent's own placements of the candidates. Among the 3,357 respondents who place themselves and both candidates on all seven issues, re-estimating the geometry gives a mismatch of 4.505 under average Harris and Trump placements and a best common-pair mismatch of 4.109, a ratio of 1.096. Giving each respondent that respondent's own perceived Harris and Trump locations raises mismatch to 4.972 and the same-sample ratio to 1.210. Replacing those respondent-specific menus with the best common pair would remove 17.3% of mismatch. Online Appendix Table OA16 reports both comparisons on the same sample.

4.3. *How stable is the diagnosis?*

Table 3 shows that the K=2 result is stable across local optimization runs. All 200 K=2 runs have a relative objective gap below 10^{-4}, and aligned centroids shift by at most 0.001 standardized units among solutions with relative gaps below 10^{-3}. K=3 has several close local modes: 117 of 200 runs have relative gaps below 10^{-4}, all fall below 10^{-3}, and the largest aligned shift in that set is 0.088. These modes share the same substantive three-bundle structure. K=4 and K=5 vary more across starts, so their aggregate values serve as diminishing-return benchmarks and their coordinates receive no interpretation.

Table 3. Optimization Stability Across 200 Independent Starts

K	Best objective	Median objective	Rel. gap $< 10^{-4}$	Rel. gap $< 10^{-3}$	Max. centroid shift
2	4.131530	4.131530	200/200	200/200	0.001
3	3.482402	3.482429	117/200	200/200	0.088
4	3.195523	3.195621	113/200	157/200	0.128
5	2.985973	2.999958	27/200	91/200	0.267

Notes: Each start uses an independently drawn weighted initialization in the standardized 7-item matched space. The two run-count columns use the relative objective gap from the best-found solution. The final column is the largest aligned Euclidean shift of any bundle among solutions whose relative objective gap is below 10^{-3}. Label switching is removed before the comparison.

The metric family preserves the ordering. Table 4 summarizes the full Online Appendix grid. Across all alternatives, the positioning ratio ranges from 1.073 to 1.107 and the third-bundle gain ranges from 15.0% to 18.6%. Leaving out defense lowers D_2^*/D_1^* most sharply. Dropping abortion yields a positioning ratio of 1.089 and a third-bundle gain of 16.9%; doubling its weight yields 1.097 and 15.4%. Every issue omission and double-weight perturbation leaves retained two-bundle mismatch as the larger margin.

Table 4. Structured Metric Sensitivity

Quantity	Baseline	Minimum	Maximum
D_2^*/D_1^*	0.590	0.541	0.595
Positioning ratio	1.097	1.073	1.107
K=2 to K=3 gain	15.7%	15.0%	18.6%

Notes: The sensitivity family includes respondent-unweighted estimates, raw issue scales in place of standardized scales, seven leave-one-issue-out estimates, and defined double-weight perturbations for abortion, the government-economic-role issue group, racial equality, and the environment. The economic-role group contains spending/services, health insurance, and jobs. The ranges exclude the baseline row; the full specification grid appears in the Online Appendix.

The broader-data checks show the direction of complete-case selection. Effective dimensionality is 2.74 among complete cases, 3.01 in the directly calculated pairwise-available weighted correlation matrix, and 2.98 after weighted mean imputation among respondents observed on at least six issues. The last construction covers 78.2% of positive survey weight (Online Appendix Tables OA14–OA15). Its positioning ratio is 1.098, essentially the complete-case estimate of 1.097. Broader coverage therefore raises measured dimensionality while leaving the positioning diagnosis stable.

5. What Does a Third Policy Bundle Add?

A menu with three entries contains every menu with two, so a third bundle weakly improves nearest-option fit by construction. The questions with empirical content are how large the gain is, where the added bundle stands—in particular, whether a centrist addition suffices—and whether the gain reaches the winner a plurality election selects. This section answers the first two in the fixed geometry; Section 5.3 bounds the third with a winner-mismatch identity and conditional benchmarks.

Table 5 shows that the best-found three-bundle menu reduces best-pair mismatch by 15.7%, with a 95% design-bootstrap interval from 14.5% to 17.2%. The best-found four- and five-bundle menus reduce it by 22.7% and 27.7%. Total gain relative to K=2 rises, while the stepwise reductions in mismatch shrink from 0.65 to 0.29 and 0.21. Even the best-found five-bundle menu leaves 65.9% of observed mismatch in place.

Figure 2 shows how the joint K=3 solution repositions the full menu. The lower

Table 5. Weighted Benchmarks from Adding Policy Bundles

K	Best-found weighted mismatch	Share of D_1^*	Gain relative to $K = 2$
1	7.00	1.00	–
2	4.13	0.59	–
3	3.48	0.50	15.7% [14.5, 17.2]
4	3.20	0.46	22.7%
5	2.99	0.43	27.7%

Notes: Mismatch values are best-found survey-weighted K-means benchmarks in the standardized, equal-issue seven-item ANES space. The last column reports the additional improvement from the best-found two-bundle menu to a richer menu of size K; brackets on the K=3 row give a 95% Rao–Wu design-bootstrap interval.

bundle moves farther toward the liberal endpoints and the upper bundle moves farther toward the conservative endpoints. The new middle bundle is moderate in aggregate and cross-cutting across issues (cf. Fowler et al. 2023): it lies near 3.3–4.1 on most scales, reaches 4.59 on defense, and falls to 2.92 on environment. The joint exercise allows all three bundles to move.

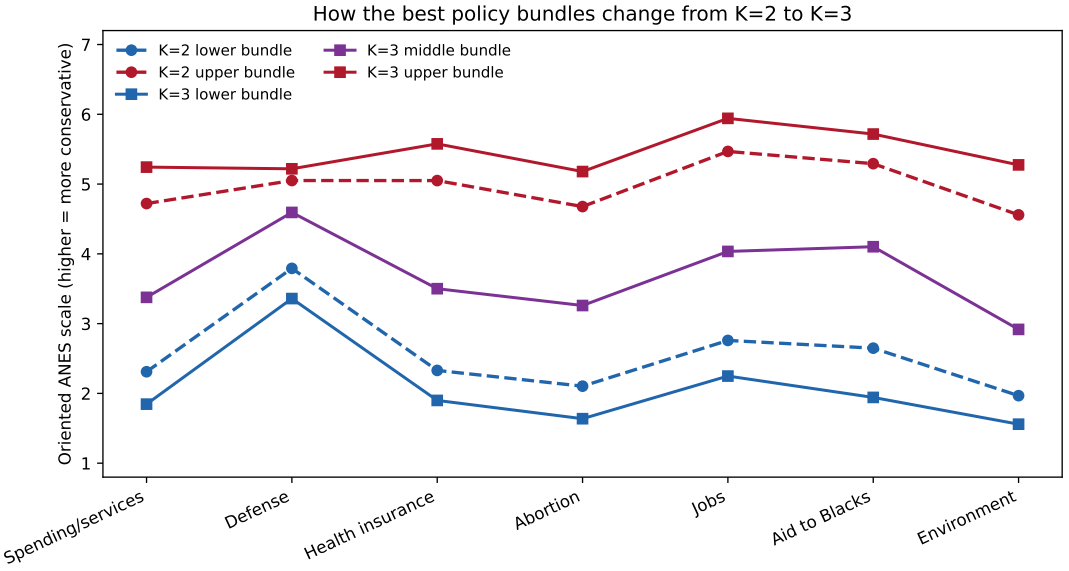


Figure 2. Best-found joint K=2 and K=3 policy bundles converted back to the oriented 1–7 ANES scales. The K=3 solution repositions all bundles. Higher values indicate the more conservative endpoint.

5.1. *Is the missing option a centrist?*

Debates about third-party entry usually name a centrist as the missing option: a divergent pair leaves the center open in the spatial logic of Downs (1957), and third-party support responds to major-party failure (Rosenstone, Behr, and Lazarus 1996). The best single platform gives that intuition an exact form: under squared distance it is the survey-weighted mean of voter positions, a platform at the electorate's center on every issue. Holding the perceived Harris and Trump locations fixed, adding this ideal centrist removes 18.4% of observed mismatch, with a 95% design-bootstrap interval from 16.5% to 20.3%, and becomes the nearest option for 35.8% of weighted voters. The best-found addition to the same fixed pair removes 19.9%, so the centrist cap-

tures 92.5% of the best-found single addition's gain. Among single added platforms, the centrist prescription therefore nearly matches the best-found addition in nearest-option fit, and adding one such platform removes about twice as much mismatch as repositioning both existing candidates.

Even this near-optimal addition leaves 81.6% of observed mismatch in place. That empirical result is consistent with the electorate's multidirectional variation; effective dimensionality summarizes that variation but does not by itself determine the finite-menu gain. A centrist improves the existing pair, but the best three-platform menu also moves the outer platforms; measured against observed mismatch, it removes about 23%, compared with 18.4% for the centrist. The baseline now shifts from Harris and Trump to the best-found pair. Added to that pair, the ideal centrist removes 9.6% of best-pair mismatch and the best-found single addition removes 10.0%, while the jointly repositioned three-bundle menu removes 15.7%. The larger gain requires moving the outer bundles outward as well, and the joint solution's middle bundle sits at the mildly cross-cutting profile of Figure 2, near but not at the electorate mean (Table 7). Plurality adds a further contingency: under the winner identity of Section 5.3, the centrist menu delivers winner-platform mismatch equal to the best single-platform value when the centrist platform wins, since its displacement from the electorate mean is zero by construction, and 1.350 or 2.146 times that value when the perceived Harris or Trump location wins. Online Appendix Table OA24 reports the full menu comparison.

5.2. *Who is nearest the middle bundle?*

Table 6 describes the voters nearest each K=3 bundle. The middle bundle receives 38.0% of weighted assignments, compared with 32.8% and 29.2% for the outer bundles. Leaners constitute 32.3% of the middle group, and pure independents 10.0%; both shares exceed those in either outer group. Most voters nearest the middle bundle are partisans or leaners, but the group has weaker partisan attachment than either outer group (cf. Lefkofridi, Wagner, and Willmann 2014; Hillen and Steiner 2020).

Table 6. Voters Nearest Each Bundle in the Best-Found Three-Bundle Menu

Bundle	*N*	Weighted share	Strong	Weak	Leaner	Independent
Lower	1,084	32.8%	55.5%	17.5%	24.9%	2.1%
Middle	1,298	38.0%	30.6%	27.1%	32.3%	10.0%
Upper	1,105	29.2%	53.4%	16.1%	29.1%	1.2%

Notes: Voters are assigned to their nearest bundle in the best-found joint K=3 menu in the standardized 7-item space. Partisan categories are shares within each bundle's weighted assignment; vote shares require electoral behavior. Party identification is missing for six cases, so the partisan-category percentages can sum below 100.

A separate within-group geometry check points in the same direction. Effective dimensionality is 2.34 among strong partisans, compared with 3.02 among weak partisans and 2.98 among leaners (Online Appendix Table OA18). The pure-independent estimate uses 160 complete cases and has a wide design-bootstrap interval. Voters with weaker major-party attachments show more multidirectional variation within their groups.

5.3. *From a richer menu to one winner*

Nearest-menu mismatch describes the options available to voters. A single-winner election asks how far voters are from the platform of the candidate who wins. If

$w(B) \in \{1, \ldots, K\}$ denotes the winner selected from menu B, winner-platform mismatch is

$$D_{\text{win}}(B) = \frac{\sum_i s_i \|\theta_i - b_{w(B)}\|^2}{\sum_i s_i} = D_1^* + \|b_{w(B)} - \bar{\theta}\|^2, \tag{5}$$

where $\bar{\theta}$ is the survey-weighted electorate mean. The identity separates dispersion around the best single platform from the winning platform's displacement from that mean.

Plurality can improve menu fit and move the winner farther from the electorate mean. Consider a one-dimensional electorate with 45% of voters at -1, 40% at 0.2, and 15% at 1. With candidates at -1 and 0.2, nearest-position voting gives the candidate at 0.2 55% and the win. Adding a candidate at 1 gives the rightmost voters a closer option and splits the previous 55% into 40% and 15%. The candidate at -1 then wins with 45%. The electorate mean is -0.22, so the winner's displacement rises from 0.42 to 0.78 even as nearest-menu fit improves. This is the familiar plurality vote-splitting mechanism (Duverger 1954; Cox 1997).

Table 7 applies equation (5) to the three-bundle solution. If the middle bundle wins, winner-platform mismatch is 1.007 times the best single-platform value. If the lower or upper bundle wins, it is 1.719 or 1.904 times that value. The same menu can therefore deliver close-to-mean or far-from-mean representation by the officeholder, depending on which bundle plurality selects.

Table 7. Winner-Platform Mismatch Conditional on Which K=3 Bundle Wins

Conditional winner	Winner mismatch	Ratio to D_1^*	Squared displacement
Lower K=3 bundle	12.033	1.719	5.033
Middle K=3 bundle	7.051	1.007	0.051
Upper K=3 bundle	13.327	1.904	6.327

Notes: Each row treats one bundle in the best-found joint K=3 menu as the winner and computes average squared distance from every voter to that bundle. The last column is the squared displacement $\|b_{\text{win}} - \bar{\theta}\|^2$ in the identity $D_{\text{win}} = D_1^* + \|b_{\text{win}} - \bar{\theta}\|^2$.

6. How Does Electorate Geometry Compare Across Elections?

The harmonized five-item ANES comparison in Figure 3 supplies descriptive point estimates of electorate geometry and partisan coalition centroids in 2016, 2020, and 2024. The 2020 point estimates show an electorate that is easier to compress into two bundles: $d_{\text{eff}} = 2.26$ and $D_2^*/D_1^* = 0.544$, compared with 2.66 and 0.608 in 2016 and 2.59 and 0.591 in 2024. Relative to 2024, the 2020 pattern has lower best-two mismatch on defense and health insurance, slightly lower mismatch on aid to Black Americans, and slightly higher mismatch on spending and jobs. Partisan-centroid separation is also greater on essentially every common issue. On these point estimates, the common-item evidence describes a more one-dimensional electorate and coalition centroids closer to the best-found two-bundle benchmark in 2020. The separate seven-item 2024 analysis supplies the candidate-specific comparison. The harmonized comparison can decompose the 2020 difference into electorate geometry and coalition separation; it cannot identify whether campaigns, candidate movements, or compositional change produced it.

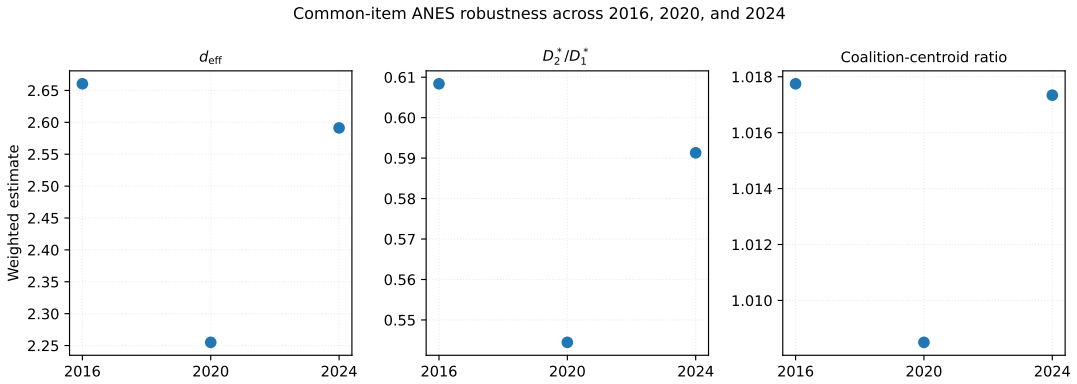


Figure 3. Descriptive point estimates of electorate geometry and partisan-coalition centroids on five ANES cumulative-file items harmonized across 2016, 2020, and 2024. The panels report effective dimensionality, retained two-bundle mismatch, and coalition-centroid positioning. The separate seven-item 2024 analysis supplies the candidate-specific comparison. Issue-level contributions appear in the Online Appendix.

7. Fit Differences and Political Dissatisfaction

The attitude models describe how respondent-level fit differences covary with post-election attitudes. For respondent i, let d_i^{obs} be squared distance to the nearer average perceived candidate, d_i^{2*} distance to the nearer best-found two-bundle point, and d_i^{3*} distance to the nearer best-found three-bundle point. The positioning gap is $p_i = d_i^{\text{obs}} - d_i^{2*}$. The signed distance change is $m_i = d_i^{2*} - d_i^{3*}$. Across the baseline geometry sample under the pre-election weights, the weighted mean of m_i equals the positive absolute reduction $D_2^* - D_3^*$; dividing that mean by D_2^* gives the system-level third-bundle gain. Because the two menus are estimated separately, 33.7% of baseline weighted cases have a negative value of m_i (Online Appendix Table OA19). Both regressors enter the models after standardization under the post-election weight.

Table 8 relates the two fit measures to two post-election outcomes. A one-standard-deviation increase in the signed K=2-to-K=3 distance change is associated with a 2.0-percentage-point increase in the fitted probability of disliking both major-party candidates and a 0.073 increase on the 0–3 dissatisfaction-with-choice scale. Respondents for whom the three-bundle menu offers a larger fit improvement also report greater dissatisfaction. The positioning gap is positively associated with both outcomes as well. K=4 and K=5 signed definitions, a weighted-logit model for dual dislike, and an external-efficacy outcome appear in the Online Appendix.

8. Discussion

Harris and Trump leave little room for repositioning: the gap between them and the best-found pair is small and lies mostly on the Trump side, while most mismatch survives under any pair found. A third option adds fit, most when all three platforms move, and plurality determines whether that fit reaches the officeholder. The cross-year comparison places 2024 between the easier-to-fit 2020 electorate and the harder-to-fit 2016 electorate.

Table 8. Post-Election Attitudes Associated with Signed Distance Changes

Outcome	N	Signed fit change		Positioning gap	
		Coef.	SE	Coef.	SE
Dual dislike of both major-party candidates	2,853	0.020	0.008	0.019	0.009
Dissatisfaction with electoral choice	2,853	0.073	0.026	0.107	0.029

Notes: Entries are weighted least-squares coefficients from models that include the standardized signed K=2-to-K=3 distance change ($d_i^{2*} - d_i^{3*}$) and standardized positioning gap ($d_i^{\text{obs}} - d_i^{2*}$) in the same regression, along with controls for party identification, ideology, age, education, race, political attention, and campaign interest. Because the K=2 and K=3 menus are estimated separately, the first difference can be negative for some respondents. Each row uses baseline geometry cases with positive post-election weight and complete data on that outcome and every control. Each analytic sample represents 56.6% of all positive post-election survey weight. Dual dislike is an indicator for rating both major-party presidential candidates at 4 or below on the 0–10 like-dislike scale. Dissatisfaction with electoral choice is coded from 0 to 3, with higher values indicating greater dissatisfaction. Standard errors use the ANES variance strata and PSUs with the post-election full-sample weight and condition on the estimated menus. The coefficients describe sample associations.

8.1. *Scope and limitations*

The headline benchmarks hold several features fixed. They use a survey-weighted sample of complete self-placements, seven matched issues, an equal-issue standardized metric, and Harris and Trump located at respondents' average placements. A campaign or party can change positions, perceptions, turnout, and voting at the same time. Candidate entry, electoral equilibrium, policy implementation, and welfare therefore require a behavioral and institutional model built for those outcomes. The joint K=3 solution answers the narrower question of how much nearest-option fit three freely positioned bundles can supply.

The analysis treats the 1–7 placements as approximately interval-scaled and assumes compensatory squared distance across issues; the raw-scale, omission, and issue-weight checks in Table 4 show how the conclusions fare under nearby metric choices. The measured policy space leaves immigration, trade, and other unmatched considerations for surveys that place voters and candidates on a broader common battery. Complete cases are more politically engaged and contain fewer pure independents than incomplete cases. Broader observed-data constructions raise effective dimensionality from 2.74 to about 3.0 while leaving the positioning ratio near 1.10 (Online Appendix Tables OA14–OA15); recovering the remaining 21.8% of positive survey weight would require the stronger assumptions of model-based multiple imputation (King et al. 2001). Single-item placements also carry response error (Ansolabehere, Rodden, and Snyder 2008). Random error attenuates cross-item correlations, which raises measured effective dimensionality and the retained-mismatch share, and it enters observed and best-found mismatch alike, which is likely to attenuate the positioning ratio toward one; part of the measured multidimensionality may therefore be noise, and the 8.8% positioning share may be conservative under a classical independent-error model. The cross-year comparison reports point estimates and leaves uncertainty in year differences open. The attitude models add a final boundary. Their estimation sample covers 56.6% of positive post-election weight after outcome and control completeness, so the associations describe that observed subpopulation. Prior alienation can shape perceived fit and later dissatisfaction, so panel or experimental evidence would be needed to establish causal direction. Their standard errors describe uncertainty conditional on the estimated menus.

Within the measured seven-issue space, the diagnosis is one-sided. Complaints about where the candidates stand concern the smaller part of the mismatch, although the

respondent-level positioning gap is largest among pure independents and rises as major-party attachment weakens (Online Appendix Table OA18). The larger part belongs to the menu. Whether a richer menu improves representation by the officeholder depends on the institutions that select the winner.

9. Conclusion

When voters in 2024 say neither candidate represents them, the evidence here points mostly to the number of options; where the candidates stand accounts for little. Harris and Trump, as voters placed them, are close to the best pair that could be found for these voters: replacing them would remove 8.8% of measured mismatch, and more than nine-tenths would remain. Even the best-found pair leaves 59.0% of the mismatch that a single platform at the electorate's center would leave. Respondents who would gain more from a three-platform menu also report greater dissatisfaction with the available choice, conditional on the included controls (Section 7).

Adding a centrist to Harris and Trump nearly achieves the best improvement available from a single addition, removing 18.4% of observed mismatch. Allowing all three platforms to move improves fit further, removing about 23%, with a middle platform that is relatively conservative on defense and relatively liberal on the environment and whose nearest voters have weaker partisan attachment than either outer group. Plurality then selects one officeholder, and which platform wins depends on electoral behavior and coordination that these conditional comparisons leave open. The next question is which institutions turn additional menu fit into representation by the winner.

A second question concerns the electorate. The benchmarks hold the distribution of views fixed, and that distribution can change. The 2016–2024 point estimates of Section 6 show no steady increase in the difficulty of fitting voters with two platforms: the 2020 electorate was easier to fit than the 2016 or 2024 electorates. A longer series, and comparisons across party systems, could establish whether changing patterns of issue disagreement make a fixed number of platforms less representative.

All proofs, simulation evidence, data details, robustness tables, and additional theoretical benchmarks appear in the Online Appendix.

Disclosure Statement

The author reports no potential conflict of interest.

Funding

This research received no specific grant from any funding agency in the public, commercial, or nonprofit sectors.

Ethics and Consent

This study is a secondary analysis of de-identified public-use ANES data. The ANES study documentation governs its consent and data-collection procedures; the analysis involved no author recruitment of or interaction with participants.

Use of Generative AI

During the preparation of this work, the author used AI-assisted tools (including OpenAI and Anthropic systems) for language editing, consistency checks, and adversarial review prompts intended to identify possible errors or overstatements. The author reviewed and revised all AI-assisted output and takes full responsibility for the content of the manuscript, including any remaining errors.

Data Availability Statement

The replication files that support the findings of this study are openly available in Zenodo at `https://doi.org/10.5281/zenodo.23358049`. The archive contains the Python scripts that produce all tables, figures, and numerical results reported here, the generated outputs themselves, and a README describing archive contents and execution order. The underlying survey inputs were derived from the publicly available ANES 2024 Time Series Study and ANES Time Series Cumulative Data File at electionstudies.org.

References

Adams, James F., Samuel Merrill III, and Bernard Grofman. 2005. *A Unified Theory of Party Competition: A Cross-National Analysis Integrating Spatial and Behavioral Factors*. Cambridge, UK: Cambridge University Press.

Ahler, Douglas J., and David E. Broockman. 2018. "The Delegate Paradox: Why Polarized Politicians Can Represent Citizens Best." *The Journal of Politics* 80 (4): 1117–1133. https://doi.org/10.1086/698755.

American National Election Studies. 2025. "ANES 2024 Time Series Study Full Release" (dataset). Version of August 8, 2025. American National Election Studies. `https://electionstudies.org/data-center/2024-time-series-study/`.

American National Election Studies. 2026. "2024 Time Series: Updates & Errata." American National Election Studies. `https://electionstudies.org/data-center/2024-time-series-study/2024-time-series-updates-errata/`.

Ansolabehere, Stephen, Jonathan Rodden, and James M. Snyder. 2008. "The Strength of Issues: Using Multiple Measures to Gauge Preference Stability, Ideological Constraint, and Issue Voting." *American Political Science Review* 102 (2): 215–232. https://doi.org/10.1017/S0003055408080210.

Balogh, Timea, Will Horne, James Adams, Simon Weschle, and Christopher Wlezien. 2025. "Does Political Sophistication Moderate How Citizens Use Information to Infer Left-Right Distances between Parties?" *Journal of Elections, Public Opinion and Parties* 35 (2): 308–325. https://doi.org/10.1080/17457289.2024.2429533.

Broockman, David E. 2016. "Approaches to Studying Policy Representation." *Legislative Studies Quarterly* 41 (1): 181–215. https://doi.org/10.1111/lsq.12110.

Cox, Gary W. 1997. *Making Votes Count: Strategic Coordination in the World's Electoral Systems*. Cambridge, UK: Cambridge University Press.

Del Giudice, Marco. 2021. "Effective Dimensionality: A Tutorial." *Multivariate Behavioral Research* 56 (3): 527–542. https://doi.org/10.1080/00273171.2020.1743631.

Downs, Anthony. 1957. *An Economic Theory of Democracy*. New York: Harper and Brothers.

Duverger, Maurice. 1954. *Political Parties: Their Organization and Activity in the Modern State*. London: Methuen.

Feitosa, Fernando, Jean-Benoît Pilet, and David Talukder. 2025. "Perceived Underrepresenta-

tion and Populist Voting." *Journal of Elections, Public Opinion and Parties* 36 (2): 430–450. https://doi.org/10.1080/17457289.2025.2483672.

Ferland, Benjamin. 2016. "Revisiting the Ideological Congruence Controversy." *European Journal of Political Research* 55 (2): 358–373. https://doi.org/10.1111/1475-6765.12127.

Fowler, Anthony, Seth J. Hill, Jeffrey B. Lewis, Chris Tausanovitch, Lynn Vavreck, and Christopher Warshaw. 2023. "Moderates." *American Political Science Review* 117 (2): 643–660. https://doi.org/10.1017/S0003055422000818.

Gersho, Allen, and Robert M. Gray. 1992. *Vector Quantization and Signal Compression*. Boston: Kluwer Academic Publishers.

Golder, Matt, and Jacek Stramski. 2010. "Ideological Congruence and Electoral Institutions." *American Journal of Political Science* 54 (1): 90–106. https://doi.org/10.1111/j.1540-5907.2009.00420.x.

Hillen, Sven, and Nils D. Steiner. 2020. "The Consequences of Supply Gaps in Two-Dimensional Policy Spaces for Voter Turnout and Political Support: The Case of Economically Left-Wing and Culturally Right-Wing Citizens in Western Europe." *European Journal of Political Research* 59 (2): 331–353. https://doi.org/10.1111/1475-6765.12348.

Huber, John D., and G. Bingham Powell. 1994. "Congruence between Citizens and Policymakers in Two Visions of Liberal Democracy." *World Politics* 46 (3): 291–326. https://doi.org/10.2307/2950684.

King, Gary, James Honaker, Anne Joseph, and Kenneth Scheve. 2001. "Analyzing Incomplete Political Science Data: An Alternative Algorithm for Multiple Imputation." *American Political Science Review* 95 (1): 49–69. https://doi.org/10.1017/S0003055401000235.

Koedam, Jelle, Garret Binding, and Marco R. Steenbergen. 2025. "Multidimensional Party Polarization in Europe: Cross-Cutting Divides and Effective Dimensionality." *British Journal of Political Science* 55: e24. https://doi.org/10.1017/S0007123424000474.

Kurella, Anna-Sophie, and Milena Rapp. 2024. "Combining Voter Preferences with Party Position Estimates from Different Sources for Studying Voting Behavior and Representation." *Electoral Studies* 87: 102734. https://doi.org/10.1016/j.electstud.2023.102734.

Lefkofridi, Zoe, Markus Wagner, and Johanna E. Willmann. 2014. "Left-Authoritarians and Policy Representation in Western Europe: Electoral Choice across Ideological Dimensions." *West European Politics* 37 (1): 65–90. https://doi.org/10.1080/01402382.2013.818354.

Lindbeck, Assar, and Jörgen W. Weibull. 1987. "Balanced-Budget Redistribution as the Outcome of Political Competition." *Public Choice* 52 (3): 273–297. https://doi.org/10.1007/BF00116710.

MacQueen, J. B. 1967. "Some Methods for Classification and Analysis of Multivariate Observations." In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1*, edited by L. M. Le Cam and J. Neyman, 281–297. Berkeley, CA: University of California Press.

Merrill, Samuel, Bernard Grofman, and James Adams. 2001. "Assimilation and Contrast Effects in Voter Projections of Party Locations: Evidence from Norway, France, and the USA." *European Journal of Political Research* 40 (2): 199–221. https://doi.org/10.1111/1475-6765.00594.

Powell, G. Bingham. 2000. *Elections as Instruments of Democracy: Majoritarian and Proportional Visions*. New Haven, CT: Yale University Press.

Rao, J. N. K., and C. F. J. Wu. 1988. "Resampling Inference with Complex Survey Data." *Journal of the American Statistical Association* 83 (401): 231–241. https://doi.org/10.1080/01621459.1988.10478591.

Rosenstone, Steven J., Roy L. Behr, and Edward H. Lazarus. 1996. *Third Parties in America: Citizen Response to Major Party Failure*. 2nd ed. Princeton, NJ: Princeton University Press.

Treier, Shawn, and D. Sunshine Hillygus. 2009. "The Nature of Political Ideology in the Contemporary Electorate." *Public Opinion Quarterly* 73 (4): 679–703. https://doi.org/10.1093/poq/nfp067.

Online Appendix

for Candidate Positioning or Limited Choice? Voter–Candidate Fit in the 2024 U.S. Election

Samir Varma
VS Asset Management, LLC

This Online Appendix reports supplementary proofs, robustness checks, measurement diagnostics, and replication details referenced in the main paper. The theoretical benchmarks in Sections 1 and 7 provide interpretive context; the empirical estimates come from the ANES workflow documented below.

1. Proofs

Proposition 1 (Properties of effective dimensionality)**.** *Let $\lambda_1, \ldots, \lambda_n$ be the eigenvalues of Σ. The effective dimensionality satisfies $1 \leq d_{\text{eff}} \leq n$, equals n when all eigenvalues are equal, and approaches 1 as a single eigenvalue dominates.*

Proof. **Part (i):** By Cauchy-Schwarz, $(\sum_i \lambda_i)^2 \leq n \sum_i \lambda_i^2$, so $d_{\text{eff}} \leq n$. Since all $\lambda_i \geq 0$ and at least one is positive, $d_{\text{eff}} \geq 1$.

Part (ii): If $\lambda_i = \bar{\lambda}$ for all i, then $\operatorname{tr}(\Sigma) = n\bar{\lambda}$ and $\operatorname{tr}(\Sigma^2) = n\bar{\lambda}^2$, so $d_{\text{eff}} = n$.

Part (iii): If $\lambda_1 / \operatorname{tr}(\Sigma) \to 1$, then $\lambda_2, \ldots, \lambda_n \to 0$ relative to λ_1. Thus $\operatorname{tr}(\Sigma^2) \to \lambda_1^2$, and $d_{\text{eff}} \to (\operatorname{tr}\Sigma)^2/\lambda_1^2 \to 1$. □

Proposition 2 (Gaussian $K = 2$ halfspace benchmark)**.** *If issue positions are Gaussian, then the best two-halfspace benchmark through the mean is*

$$D_2^{\text{HB}} = \operatorname{tr}(\Sigma) - \frac{2}{\pi}\lambda_1. \tag{1}$$

Proof. Translate so $\mu = 0$. For any two-cell partition P with best cell centroids, mismatch satisfies the variance decomposition

$$D(P) = \mathbb{E}\|\theta\|^2 - \mathbb{E}\big\|\mathbb{E}[\theta \mid P(\theta)]\big\|^2.$$

Restrict now to halfspace partitions $\{v \cdot \theta \geq 0\}$ through the mean; symmetry implies equal cell mass and opposite centroids. For such a partition one can compute

$$m(v) \equiv \mathbb{E}[\theta \mid v \cdot \theta \geq 0] = \sqrt{\frac{2}{\pi}}\frac{\Sigma v}{\sqrt{v^\top \Sigma v}}, \qquad \|m(v)\|^2 = \frac{2}{\pi}\frac{v^\top \Sigma^2 v}{v^\top \Sigma v}.$$

See Tallis (1965) for plane-truncation moments of the multivariate normal. Maximizing the Rayleigh quotient $(v^\top \Sigma^2 v)/(v^\top \Sigma v)$ yields λ_1, attained at $v = v_1$. Since $\mathbb{E}\|\theta\|^2 = \operatorname{tr}(\Sigma)$, equation (1) follows for the best partition in this restricted mean-through-halfspace class. The empirical D_2^* minimizes mismatch over the unrestricted partition class. □

Proposition 3 (Two-party minimum differentiation under logit)**.** *Under central symmetry with $K = 2$, suppose each party maximizes expected vote share $S_k(\mathbf{B}) =$*

$\mathbb{E}[p_k(\theta; \mathbf{B})]$. *The centrist minimum-differentiation profile is a Nash equilibrium of the logit competition game where voters choose among parties with probabilities*

$$p_k(\theta; \mathbf{B}) = \frac{\exp(-\|\theta - b_k\|^2/\tau)}{\sum_{j=1}^{K} \exp(-\|\theta - b_j\|^2/\tau)}. \tag{2}$$

Proof. Translate coordinates so that $\mu = 0$. Fix party 2 at $b_2 = 0$ and let party 1 choose $b \equiv b_1$. From (2), party 1's choice probability can be written as

$$p_1(\theta; b, 0) = \frac{1}{1 + \exp((\|\theta - b\|^2 - \|\theta\|^2)/\tau)} = \sigma\left(\frac{2b \cdot \theta - \|b\|^2}{\tau}\right)$$

where $\sigma(x) \equiv 1/(1 + e^{-x})$ is the logistic CDF. Let $Z \equiv b \cdot \theta$. Central symmetry implies $Z \overset{d}{=} -Z$ for any fixed b. Define $a \equiv \|b\|^2/\tau \geq 0$ and $t \equiv 2Z/\tau$ so that $p_1 = \sigma(t - a)$. Then

$$S_1(b, 0) = \mathbb{E}[\sigma(t - a)] = \frac{1}{2}\mathbb{E}[\sigma(t - a) + \sigma(-t - a)].$$

Because $\sigma(\cdot)$ is increasing and $a \geq 0$,

$$\sigma(t - a) + \sigma(-t - a) \leq \sigma(t) + \sigma(-t) = 1$$

with strict inequality whenever $a > 0$ and $\mathbb{P}(t \neq 0) > 0$. Hence $S_1(b, 0) \leq 1/2$ with equality if and only if $b = 0$ under nondegeneracy. Therefore $b_1 = \mu$ is the unique best response to $b_2 = \mu$, and by symmetry (μ, μ) is a Nash equilibrium. □

This equilibrium is a separate theoretical benchmark; convergence toward the mean under probabilistic spatial voting is a known result (Hinich 1977; Adams, Merrill, and Grofman 2005). The stated logit game produces co-location at the electorate mean. The empirical analysis instead uses proximity-minimizing bundles, average perceived candidate locations, and plurality winner selection.

2. Simulation Evidence for Effective Dimensionality

Figure OA1 reports synthetic Gaussian simulations with fixed $\mathrm{tr}(\Sigma) = 1$ and three different eigenvalue spectra. For each spectrum, D_K^* is approximated via K-means on a large simulated sample and compared to the $-2/n$ and $-2/d_{\text{eff}}$ benchmarks on a log–log scale (Gersho and Gray 1992). The effective-dimensionality slope tracks the small-K regime when the spectrum is concentrated.

Table OA1 reports the approximation error for the same simulation designs.

3. Data and Measurement Details

3.1. *Matched ANES 2024 variables*

Table OA2 records the voter and candidate variables that define the common seven-item space. The spending/services triplet is reverse-coded so higher values have the

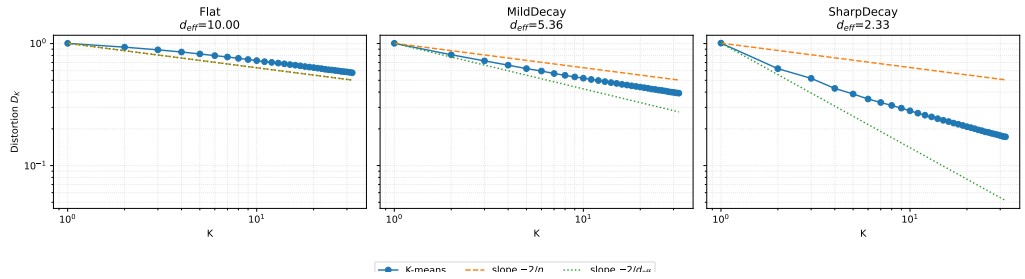


Figure OA1. Simulated Gaussian mismatch for three eigenvalue spectra (flat, mildly decaying, sharply decaying) with $n = 10$ dimensions and $N = 6{,}000$ draws. Log–log scale; lines show reference slopes $-2/n$ and $-2/d_{\text{eff}}$ anchored at $K = 1$. Source: Simulations.

Table OA1. Effective-Dimensionality Prediction Error in Gaussian Simulations

Spectrum	K	D_K (K-means)	$\widehat{D}_K$	% error
Flat	2	0.935	0.874	6.5
Flat	3	0.888	0.806	9.2
Flat	5	0.822	0.728	11.5
MildDecay	2	0.810	0.776	4.2
MildDecay	3	0.722	0.667	7.6
MildDecay	5	0.623	0.551	11.6
SharpDecay	2	0.623	0.557	10.6
SharpDecay	3	0.519	0.393	24.2
SharpDecay	5	0.386	0.254	34.3

Notes: $\widehat{D}_K = D_1 K^{-2/d_{\text{eff}}}$, with d_{eff} computed from the spectrum and D_1 estimated empirically from the simulated sample. Percentage error is absolute.

conservative orientation used by the other six scales. Every main-text distance uses survey-weighted standardization followed by equal issue weights.

The analysis file is the August 8, 2025 ANES full release named in the replication manifest. ANES subsequently issued a May 19, 2026 re-release. Its published correction notice identifies variables V241106x, V241352, V241353x, V241445x, V241620, V242051, V242213, V242233, V242234x, and V242243–V242248x; none enters the voter coordinates, candidate coordinates, weights, controls, or outcomes used here. A separate label-only change affects the political-attention control V241004 while preserving its numeric ordering.[1]

Table OA2. Matched ANES 2024 Policy Placement Variables

Issue	Self	Harris	Trump
Spending/services	V241239	V241240	V241241
Defense spending	V241242	V241243	V241244
Government health insurance	V241245	V241246	V241247
Abortion	V241248	V241250	V241251
Guaranteed jobs and income	V241252	V241253	V241254
Aid to Black Americans	V241255	V241256	V241257
Environmental regulation	V241258	V241259	V241260

Notes: Each row uses matched pre-election 1–7 self, Democratic presidential candidate, and Republican presidential candidate placements. On abortion, 1 means always permitted without restrictions and 7 means never permitted. Values −9, −8, and −1 denote refusal, don't know, and inapplicability; self-placement code 99 denotes haven't thought much.

[1] `https://electionstudies.org/anes_timeseries_2024/`

Table OA3 records the raw endpoints and the single orientation reversal used in the analysis.

Table OA3. Orientation of the Seven Matched Issues

Issue	Raw endpoints (1, 7)	Recoding
Spending/services	Fewer services | more services	$x \mapsto 8 - x$
Defense	Decrease | increase	None
Health insurance	Government plan | private plan	None
Abortion	Always permitted | never permitted	None
Jobs	Government guarantee | individual responsibility	None
Aid to Black Americans	Government help | self-help	None
Environment	Tougher regulation | regulation too burdensome	None

Table OA4 records the survey-design variables, outcomes, controls, and coding conventions.

Table OA4. Weights, Variance Design, Outcomes, and Controls

Role	ANES variable(s)	Coding and use
Pre-election weight	V240108a	Full-sample weight for policy geometry and candidate perceptions.
Post-election weight	V240108b	Full-sample weight for attitude models.
Variance design	V240108c, V240108d	Full-sample variance PSU and variance stratum, respectively.
Party identification	V241227x	Seven categories; categorical control with pure independent (4) as reference.
Ideology	V241177	Seven-point scale, survey-weight standardized; enters linearly and quadratically.
Age	V241458x	Survey-weight standardized continuous control.
Education	V241465x	Five categories; categorical control with category 1 as reference.
Race/ethnicity	V241501x	Six categories; categorical control with category 1 as reference.
Political attention	V241004	Reversed so higher means more attention, then survey-weight standardized.
Campaign interest	V241005	Reversed so higher means more interest, then survey-weight standardized.
Dual dislike	V242434, V242435	Indicator equal to one when both candidate ratings are 4 or below on the 0–10 scales.
Choice dissatisfaction	V242429	The four codebook values 1, 2, 4, and 5 map to 0, 1, 2, and 3, respectively.
External efficacy	V242200	Five categories reversed to 0–4 so higher values mean stronger agreement that officials do not care.
Time comparison	VCF0004, VCF0009z, VCF0301, VCF0839, VCF0843, VCF0806, VCF0809, VCF0830	Election year, cumulative-file weight, party identification, and five harmonized policy items.

Notes: Attitude models require a complete seven-item self-placement record, a positive post-election weight, valid outcome and controls, and the stated reference categories. The weighted least-squares and logit specifications include the signed fit change, positioning gap, all listed controls, and an intercept. The replication archive records every recode and sample restriction.

3.2. *Design-bootstrap intervals and candidate means*

The analysis uses 400 Rao–Wu rescaled-bootstrap replicates (Rao and Wu 1988). Each replicate resamples ANES variance PSUs within variance strata, rescales the selected PSUs, and re-estimates survey-weighted item standardization, average candidate perceptions, and the weighted K=2 and K=3 benchmarks. Tables OA5 and OA6 report percentile intervals from the resulting replicate distributions.

Table OA5. Design-Bootstrap Summary for the 2024 Headline Baseline

Statistic	Estimate	95% CI
d_{eff}	2.74	[2.60, 2.89]
D_2^*/D_1^*	0.59	[0.57, 0.61]
Positioning ratio, candidate placements	1.10	[1.08, 1.12]
K=2 to K=3 gain	15.7%	[14.5%, 17.2%]

Notes: Intervals are percentile intervals from 400 Rao–Wu rescaled-bootstrap replicates over the ANES full-sample variance PSUs within variance strata. Each replicate re-estimates standardization, covariance, the K=2 and K=3 benchmarks, and candidate-placement means. Singleton strata are retained as certainty units.

Table OA6. Weighted Candidate Placement Means on the Matched 7-Item Space

Actor	Issue	Mean [95% CI]
Harris	Spending/services	2.52 [2.45, 2.59]
Harris	Defense	3.51 [3.43, 3.60]
Harris	Health insurance	2.55 [2.47, 2.64]
Harris	Abortion	1.99 [1.94, 2.05]
Harris	Jobs	2.65 [2.57, 2.74]
Harris	Aid to Blacks	2.35 [2.27, 2.43]
Harris	Environment	2.41 [2.35, 2.49]
Trump	Spending/services	5.38 [5.29, 5.46]
Trump	Defense	5.30 [5.21, 5.38]
Trump	Health insurance	5.82 [5.74, 5.88]
Trump	Abortion	5.57 [5.49, 5.65]
Trump	Jobs	5.87 [5.80, 5.93]
Trump	Aid to Blacks	5.76 [5.69, 5.83]
Trump	Environment	5.92 [5.84, 5.99]

Notes: Means are on the oriented 1–7 issue scales used in the analysis, where higher values are more conservative. Confidence intervals use the Rao–Wu ANES design bootstrap described in Table OA5.

3.3. *Best-found policy-bundle profiles*

Table OA7 converts the best-found K=2, K=3, and K=4 bundles back to the oriented 1–7 scales. The K=3 middle bundle is moderate in aggregate and differs from the voter median issue by issue.

Table OA7. Best-Found Weighted Policy-Bundle Profiles in the Matched Seven-Item Space

K	Bundle	Spending	Defense	Health	Abortion	Jobs	Aid to Blacks	Environment
2	Lower	2.31	3.79	2.33	2.10	2.76	2.65	1.97
2	Upper	4.72	5.05	5.05	4.68	5.47	5.29	4.56
3	Upper	5.24	5.22	5.58	5.18	5.94	5.72	5.27
3	Middle	3.38	4.59	3.50	3.26	4.03	4.10	2.92
3	Lower	1.85	3.36	1.90	1.64	2.25	1.94	1.56
4	ID 1	2.28	5.23	2.56	2.43	2.67	2.90	2.14
4	ID 2	1.94	2.42	1.75	1.60	2.34	1.92	1.52
4	ID 3	3.86	4.44	4.07	3.58	4.72	4.53	3.32
4	ID 4	5.44	5.39	5.78	5.44	6.03	5.88	5.62

Notes: Entries are weighted K-means centroids converted back to the oriented 1–7 issue scales. Higher values are more conservative. The coordinate-level discussion covers K=2 and K=3, whose rows use lower, middle, and upper labels. Higher-K rows retain arbitrary numerical IDs.

4. Optimization and Metric Robustness

4.1. *Independent starts*

The main text reports the 200-start summary for K=2 through K=5. The underlying CSV records every objective and the aligned maximum bundle shift. These diagnostics establish numerical stability around the reported best-found solutions. Global optimality remains a theoretical property of D_K^*.

Table OA8. Optimization Stability Across 200 Independent Starts

K	Best objective	Median objective	Rel. gap $< 10^{-4}$	Rel. gap $< 10^{-3}$	Max. centroid shift
2	4.131530	4.131530	200/200	200/200	0.001
3	3.482402	3.482429	117/200	200/200	0.088
4	3.195523	3.195621	113/200	157/200	0.128
5	2.985973	2.999958	27/200	91/200	0.267

Notes: Each start uses an independently drawn weighted initialization in the standardized 7-item matched space. The two run-count columns use the relative objective gap from the best-found solution. The final column is the largest aligned Euclidean shift of any bundle among solutions whose relative objective gap is below 10^{-3}. Label switching is removed before the comparison.

4.2. *Metric and issue-weight sensitivity*

Table OA9 reports the structured metric perturbation family. Raw mismatch levels move with the scale, while every specification preserves the paper's ordering. The government-economic-role perturbation doubles spending/services, health insurance, and jobs; the other named perturbations double a single issue. A matched importance battery would be needed to estimate voter-specific salience.

Table OA9. Full Metric and Issue-Weight Sensitivity Grid

Specification	Dimensions	D_1^*	D_2^*	D_2^*/D_1^*	Positioning ratio	K=3 gain
Baseline: survey-weighted, standardized, equal issue weights	7	7.000	4.132	0.590	1.097	15.7%
Respondent-unweighted, standardized, equal issue weights	7	7.000	4.029	0.576	1.107	16.2%
Survey-weighted raw 1-7 scales, equal issue weights	7	25.247	14.691	0.582	1.100	16.0%
Leave one issue out: Spending/services	6	6.000	3.539	0.590	1.103	15.9%
Leave one issue out: Defense	6	6.000	3.247	0.541	1.107	18.6%
Leave one issue out: Health insurance	6	6.000	3.528	0.588	1.094	16.7%
Leave one issue out: Abortion	6	6.000	3.493	0.582	1.089	16.9%
Leave one issue out: Jobs	6	6.000	3.571	0.595	1.099	16.0%
Leave one issue out: Aid to Blacks	6	6.000	3.544	0.591	1.082	15.0%
Leave one issue out: Environment	6	6.000	3.551	0.592	1.073	15.0%
Double weight: abortion	7	8.000	4.646	0.581	1.097	15.4%
Double weight: government economic role bundle	7	10.000	5.671	0.567	1.090	16.1%
Double weight: racial equality	7	8.000	4.626	0.578	1.102	17.2%
Double weight: environment	7	8.000	4.620	0.577	1.103	17.1%

Notes: Survey-weighted specifications use the ANES pre-election full-sample weight. Standardized specifications use survey-weighted item means and standard deviations; the respondent-unweighted specification re-standardizes with equal respondent weights, so each standardized issue again has unit variance and D_1^* equals the number of issues. Equal issue weights mean that every included standardized issue contributes symmetrically to squared distance. Double-weight specifications multiply the stated issue contribution to squared distance by two. The government-economic-role bundle contains spending/services, health insurance, and jobs. These perturbations test metric dependence; respondent-specific salience would require a matched importance battery.

4.3. *Alternative position benchmarks*

Table OA10 compares average perceived Harris and Trump locations with partisan coalition centroids and the pre-existing swing-weighted targeting benchmark. For the last row, a survey-weighted logit predicts two-party vote or intention from all seven standardized policy coordinates. If p_i is the fitted probability of a Republican choice, the pivotality term is $q_i = p_i(1 - p_i)$, in the swing-voter tradition of Lindbeck and Weibull (1987). Democratic and Republican targeting locations are the policy means of strong and weak identifiers (party-identification categories 1–2 and 6–7), weighted by $s_i q_i$ within coalition. The resulting locations are evaluated using the baseline survey weights. The row supplies a theoretical targeting comparison; average perceived Harris and Trump locations remain the main empirical measure.

Table OA10. Alternative Perceived-Position Benchmarks on the Weighted Seven-Item Baseline

Benchmark	Weighted mismatch	Positioning ratio	Role
Candidate placements	4.53	1.10	Empirical baseline
Partisan coalition centroids	4.16	1.01	Coalition-centroid comparison
Swing-weighted targeting benchmark	5.14	1.24	Theoretical targeting benchmark

Notes: All rows use the same survey-weighted seven-item voter geometry and the same best-found two-bundle benchmark. The candidate-placement row is the empirical baseline used in the main text; the coalition-centroid row is a low-mismatch empirical comparison. The swing-weighted row is a theoretical targeting benchmark constructed from a weighted policy-only logit fit on 3,065 two-party vote or intent observations (McFadden pseudo-$R^2 = 0.534$).

5. Missingness and Perception Diagnostics

5.1. *Item missingness and engagement*

Table OA11 separates code 99 from other self-placement nonresponses. Table OA12 reports valid Harris and Trump placements both among all positive-weight cases and in the exact seven-item analytic baseline. Table OA13 then shows the selection into

complete cases, the 78.2% of positive survey weight covered by the at-least-6/7 partial-case sensitivity, and the still-uncovered group with at most five items.

Table OA11. Item-Level Missingness in the Matched Policy Battery

Issue	Valid	Haven't thought	Don't know	Refused	Inapplicable
Spending/services	85.5%	14.2%	0.1%	0.2%	0.0%
Defense	84.4%	15.2%	0.1%	0.3%	0.0%
Health insurance	89.8%	9.9%	0.1%	0.2%	0.0%
Abortion	93.2%	6.3%	0.0%	0.5%	0.0%
Jobs	89.1%	10.6%	0.2%	0.2%	0.0%
Aid to Blacks	86.4%	12.4%	0.3%	0.9%	0.0%
Environment	83.9%	15.7%	0.2%	0.2%	0.0%

Notes: Shares are survey-weighted among respondents with positive pre-election full-sample weights. Code 99, refusal, don't know, and inapplicability appear as separate categories.

Table OA12. Availability of Candidate Placements by Issue

Issue	All positive-weight cases		7/7 analytic baseline	
	Harris	Trump	Harris	Trump
Spending/services	96.9%	97.8%	99.2%	99.5%
Defense	95.2%	96.4%	97.5%	98.2%
Health insurance	96.3%	96.6%	98.6%	98.2%
Abortion	97.5%	97.5%	99.1%	99.3%
Jobs	96.2%	96.8%	99.1%	99.4%
Aid to Blacks	96.6%	96.7%	99.1%	99.3%
Environment	95.5%	95.8%	98.9%	99.1%

Notes: Shares are survey-weighted. The last two columns use the exact seven-self-placement complete-case population in which average candidate locations are estimated; within it, all valid candidate placements enter issue by issue.

Table OA13. Political Engagement Across Missingness Groups

Group	N	Weighted share	Attention	Interest	Independent
Complete 7/7	3,487	63.3%	2.63	1.42	4.9%
Incomplete matched battery (<7/7)	1,789	36.7%	2.05	1.08	10.1%
At least 6/7 (partial-case sensitivity)	4,266	78.2%	2.57	1.38	5.1%
At most 5/7 (outside partial-case sensitivity)	1,010	21.8%	1.88	0.99	12.9%

Notes: Higher attention and campaign-interest values indicate greater engagement. The pure-independent share uses the ANES seven-category party-identification measure. The incomplete row is the complement of the complete-case baseline among respondents with positive survey weight. The at-least-6/7 row overlaps the complete row and identifies the population covered by the partial-case sensitivity; the at-most-5/7 row identifies the cases outside its coverage.

5.2. *Pairwise and partial-case geometry*

The pairwise-available weighted correlation matrix is positive semidefinite, with minimum eigenvalue 0.377. Table OA14 compares its effective dimensionality with the complete-case and partial-case constructions. Table OA15 adds the positioning ratio and K=2-to-K=4 gain for the partial-case sensitivity.

5.3. *Partisan perception constructions*

Table OA16 first reconstructs the common Harris and Trump locations from different partisan audiences. It then compares average and respondent-specific placements

Table OA14. Geometry Under Alternative Missing-Data Constructions

Construction	N or pair-N range	d_{eff}	Min. eigenvalue	PSD
Complete-case weighted correlation	3,487	2.739	–	Yes
Pairwise-available weighted correlation	4,185–4,606	3.009	0.377	Yes
Weighted mean imputation among respondents with >=6/7 items	4,266	2.980	–	Yes

Notes: The pairwise-available matrix uses all positive-weight respondents observed on each item pair and is checked for positive semidefiniteness before its eigenvalue summary is interpreted. Weighted mean imputation is limited to respondents observed on at least six of seven issues.

Table OA15. Partial-Case Robustness for the Weighted Seven-Item Baseline

Specification	N	d_{eff}	D_2^*/D_1^*	Positioning ratio	K=2 to K=4
Complete cases (7/7 matched items)	3,487	2.74	0.59	1.10	22.7%
Partial cases (>=6/7 matched items; weighted mean imputation)	4,266	2.98	0.62	1.10	21.5%

Notes: The headline row is the complete-case weighted seven-item baseline used in the main text. The partial-case row retains respondents with at least six of seven self-placement items and fills the remaining item with weighted item means before re-estimating the geometry, best-found menus, and candidate-placement benchmark.

on the exact complete-perception sample. The pooled ratio remains 1.096, while the respondent-specific ratio is 1.210; the latter combines perception heterogeneity with common-menu positioning. The large out-partisan ratio likewise shows why projection and partisan contrast have no simple monotone implication for the aggregate positioning gap.

5.4. *Actual-winner perceived-location benchmark*

Main-text equation (5) can be evaluated for Trump as the actual 2024 winner. Table OA17 reports this supplementary benchmark. The ratio to the best single platform ranges from 1.728 under Republican perceptions of Trump to 2.763 under Democratic perceptions. The variation shows how the measurement changes with the audience supplying Trump's perceived location.

6. Heterogeneity, Attitudes, and Time

6.1. *Partisan-strength heterogeneity*

Table OA18 reports ANES design-bootstrap intervals for the partisan-strength estimates. The pure-independent estimate is based on 160 complete cases and should be treated as descriptive.

6.2. *Post-election attitude robustness*

Table OA19 first reports how often the respondent-level distance difference is negative, zero, or positive when the K=2 and larger menus are optimized separately. Table OA20 reports the K=3, K=4, and K=5 signed distance changes for dual dislike, dissatisfaction with choice, and external efficacy. Table OA21 uses weighted logit for dual dislike.

Table OA16. Sensitivity to Whose Candidate Perceptions Define the Menu

Panel A: Alternative sources for common Harris and Trump locations

Perception construction	N	Observed mismatch	Positioning ratio
All respondents	3,487	4.531	1.097
Democratic identifiers perceive both	3,487	4.871	1.179
Republican identifiers perceive both	3,487	4.815	1.165
Co-partisan candidate perceptions	3,487	4.420	1.070
Out-partisan candidate perceptions	3,487	5.816	1.408

Panel B: Common and respondent-specific menus on the same complete-perception sample

Perception construction	N	Best-pair D	Observed D	Ratio	Removable
Average placements (all respondents)	3,357	4.109	4.505	1.096	8.8%
Respondent-specific perceived locations	3,357	4.109	4.972	1.210	17.3%

Notes: Panel A evaluates every common menu against the same 3,487 baseline voters and the same best-pair denominator. Democratic and Republican identifiers are categories 1–2 and 6–7 of the ANES party-identification scale. The co-partisan construction uses Democratic perceptions of Harris and Republican perceptions of Trump; the out-partisan construction reverses those sources. Panel B restricts the data to the 3,357 respondents with complete Harris and Trump placements on every issue, then re-estimates survey-weighted standardization, the best common pair, and the pooled candidate means on that exact sample. The respondent-specific row lets each voter face that voter's own perceived Harris and Trump locations. Its ratio combines within-person perception heterogeneity with the distance of average Harris and Trump locations from the best common pair and therefore supplies a broader perception-sensitive comparison.

Table OA17. Actual-Winner Perceived-Location Benchmark

Trump perception source	Winner mismatch	Ratio to D_1^*	Excess from electorate mean
All respondents	15.024	2.146	8.024
Democratic identifiers perceive both	19.342	2.763	12.342
Republican identifiers perceive both	12.097	1.728	5.097
Co-partisan candidate perceptions	12.097	1.728	5.097
Out-partisan candidate perceptions	19.342	2.763	12.342

Notes: Trump is the actual 2024 winner. Mismatch is average squared distance to the stated average perceived Trump location. Its decomposition is $D_{\text{win}} = D_1^* + \|b_{\text{win}} - \bar{\theta}\|^2$. Perception sensitivity keeps this descriptive benchmark outside the headline comparisons.

Table OA18. Signed Fit Differences by Partisan Strength in the Matched 7-Item Space

Group	N	d_{eff} [95% CI]	Mean positioning gap, common menus [95% CI]	Signed change, common menus (2 to 4) [95% CI]
Strong partisans	1,670	2.34 [2.20, 2.46]	0.13 [0.04, 0.25]	0.96 [0.87, 1.06]
Weak partisans	709	3.02 [2.75, 3.33]	0.56 [0.42, 0.74]	0.93 [0.82, 1.07]
Leaners	942	2.98 [2.72, 3.25]	0.59 [0.43, 0.75]	0.86 [0.76, 0.98]
Independents	160	5.49 [4.56, 5.71]	1.06 [0.76, 1.46]	1.23 [0.96, 1.51]

Notes: Effective dimensionality is re-estimated within each partisan-strength subgroup. The other two columns are subgroup weighted means evaluated against the common full-sample menus: the average perceived candidate pair and the best two- and four-bundle menus. In each Rao–Wu replicate, survey-weighted standardization and those common full-sample menus are re-estimated before the subgroup means are taken. These columns describe where subgroups fare under the same system-level menu; subgroup-specific optimization would define a different estimand. Party-identification codes are missing for six of the 3,487 complete cases, so subgroup counts sum to 3,481.

Table OA19. Respondent-Level Signs of Separately Optimized Menu Comparisons

K comparison	Negative share	Zero share	Positive share	Weighted mean
K=2 to K=3	33.7%	0.0%	66.3%	0.649
K=2 to K=4	28.1%	0.0%	71.9%	0.936
K=2 to K=5	26.8%	0.0%	73.2%	1.146

Notes: The respondent-level quantity is $d_i^{2*} - d_i^{K*}$. The system-level weighted mean is positive because the larger menu lowers aggregate optimized mismatch. Independently optimized menus can differ respondent by respondent, so a negative value means that respondent i is farther from the reported K-bundle solution than from the reported two-bundle solution.

Table OA20. Post-Election Attitude Models: K-Sensitivity and External-Efficacy Robustness

Outcome	K benchmark	N	Distance-change coef.	SE	Positioning coef.	SE	R^2
Dual dislike of both major-party candidates	3	2,853	0.020	0.008	0.019	0.009	0.080
Dual dislike of both major-party candidates	4	2,853	0.028	0.009	0.019	0.009	0.084
Dual dislike of both major-party candidates	5	2,853	0.029	0.009	0.017	0.009	0.085
Dissatisfaction with electoral choice	3	2,853	0.073	0.026	0.107	0.029	0.179
Dissatisfaction with electoral choice	4	2,853	0.063	0.027	0.098	0.029	0.179
Dissatisfaction with electoral choice	5	2,853	0.043	0.028	0.088	0.028	0.177
Officials do not care what people like me think	3	2,872	0.067	0.034	0.045	0.035	0.055
Officials do not care what people like me think	4	2,872	0.069	0.033	0.040	0.034	0.056
Officials do not care what people like me think	5	2,872	0.062	0.031	0.032	0.034	0.055

Notes: Each row is a weighted least-squares model with the standardized signed distance change $d_i^{2*} - d_i^{K*}$, standardized positioning gap, and the same control set used in the main text: party identification, ideology, age, education, race, political attention, and campaign interest. Separately optimized menus can differ respondent by respondent. The third outcome recodes the ANES external-efficacy item so higher values indicate stronger agreement that public officials do not care what people like me think. Standard errors use ANES variance strata and PSUs and condition on the estimated menus.

Table OA21. Weighted Logit Robustness for Dual Dislike

K benchmark	N	Distance-change coef.	SE	Positioning coef.	SE
3	2,853	0.242	0.102	0.228	0.112
4	2,853	0.356	0.113	0.224	0.106
5	2,853	0.349	0.101	0.179	0.104

Notes: Entries are weighted logit coefficients for the dual-dislike outcome using the same regressors and controls as the main-text models. Positive coefficients indicate a higher probability of disliking both major-party candidates. Standard errors use ANES variance strata and PSUs and condition on the estimated menus.

6.3. *Harmonized time comparison*

Table OA22 reports the five-item geometry and partisan-centroid comparison. Table OA23 and Figure OA2 show the issue contributions. At the descriptive point-estimate level, 2020 has substantially lower best-found K=2 mismatch than 2024 on defense and health insurance, slightly lower mismatch on aid to Black Americans, and slightly higher mismatch on spending and jobs. This common-item comparison measures electorate geometry and coalition alignment; the separate seven-item 2024 analysis measures candidate-specific fit.

Table OA22. Common-Item ANES Robustness Across Election Years

Year	N	d_{eff}	D_2^*/D_1^*	Coalition-centroid ratio
2016	2,886	2.66	0.61	1.02
2020	5,718	2.26	0.54	1.01
2024	3,655	2.59	0.59	1.02

Notes: This table uses the harmonized 5-item ANES cumulative-file battery available in 2016, 2020, and 2024, weighted by the cumulative-file full-sample weight. The final column reports the partisan-coalition-centroid comparison; the 2024 headline analysis uses candidate placements.

Table OA23. Issue-Level Decomposition of the 2016–2024 Common-Item Comparison

Year	Issue	|Partisan separation|	Partisan-menu contribution	Best-found K=2 contribution
2016	Spending/services	1.243	0.601	0.628
2016	Defense	0.816	0.794	0.770
2016	Health insurance	1.154	0.544	0.530
2016	Jobs	1.053	0.585	0.556
2016	Aid to Blacks	1.141	0.573	0.556
2020	Spending/services	1.295	0.555	0.580
2020	Defense	0.988	0.737	0.706
2020	Health insurance	1.348	0.451	0.450
2020	Jobs	1.300	0.492	0.487
2020	Aid to Blacks	1.305	0.511	0.499
2024	Spending/services	1.213	0.558	0.556
2024	Defense	0.779	0.866	0.864
2024	Health insurance	1.165	0.554	0.540
2024	Jobs	1.123	0.506	0.468
2024	Aid to Blacks	1.300	0.524	0.529

Notes: All quantities use within-year survey-weighted standardization on the five harmonized cumulative-file items. The comparison uses partisan coalition centroids because the common cross-wave battery contains voter placements and party identification. It describes electorate geometry and coalition alignment. The separate seven-item 2024 analysis supplies the Harris–Trump candidate comparison.

6.4. *Centrist-third-candidate counterfactual*

Table OA24 reports the menu comparison behind the centrist analysis in the main text. The ideal centrist is the survey-weighted mean of voter positions, which under squared distance is the best single platform. Each fixed-pair addition minimizes nearest-option mismatch with the stated pair held fixed, using multi-start alternating optimization; the joint K=3 row repositions all three bundles. The design bootstrap re-estimates survey-weighted standardization, average perceived candidate locations, the centrist location, and the fixed-pair addition in every replicate.

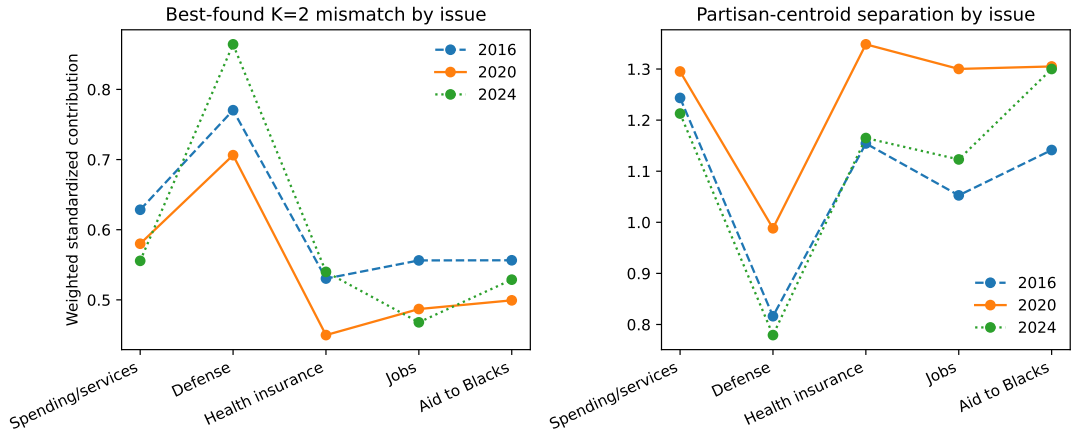


Figure OA2. Issue contributions to the harmonized time comparison. The left panel decomposes best-found K=2 mismatch; the right panel reports absolute separation between Democratic and Republican coalition centroids. All quantities use survey-weighted within-year standardization.

Table OA24. Centrist-Third-Candidate Counterfactual in the Matched 7-Item Space

Menu	Mismatch	Gain	Share nearest added option
Harris and Trump	4.531	—	—
Harris and Trump + ideal centrist	3.698	18.4% [16.5, 20.3]	35.8% [32.9, 38.6]
Harris and Trump + best-found addition	3.630	19.9% [18.2, 21.8]	34.9% [31.9, 37.4]
Best-found pair	4.132	—	—
Best-found pair + ideal centrist	3.734	9.6%	—
Best-found pair + best-found addition	3.719	10.0%	—
Best-found joint K=3 menu	3.482	15.7% [14.5, 17.2]	—
Perceived Harris fixed, Trump location freed	4.204	7.2%	—
Perceived Trump fixed, Harris location freed	4.378	3.4%	—

Notes: The ideal centrist is the survey-weighted mean of voter positions on every issue, which is the best single platform under squared distance. Gains in the upper panel are relative to mismatch under Harris and Trump; gains in the lower panel are relative to best-found-pair mismatch. Fixed-pair additions are optimized by multi-start alternating assignment and weighted-mean updates with the pair held fixed, using 56 starts for the point estimates and 13 per bootstrap replicate; the joint K=3 row repositions all three bundles and repeats the main text's Table 5 gain. Bracketed intervals are 95% percentile intervals from 400 Rao–Wu design-bootstrap replicates that re-estimate standardization, candidate means, the centrist, and the fixed-pair addition; best-found-pair rows condition on the point-estimate best-found pair and are therefore reported without intervals. The centrist captures 92.5% [88.5, 95.0] of the best-found addition's gain against Harris and Trump. Both additions gain more against the wider Harris–Trump pair than against the best-found pair, which leaves less interior room for a third option. The best-found addition to Harris and Trump lies at spending/services 3.69, defense 4.79, health insurance 3.94, abortion 3.64, jobs 4.53, aid to Black Americans 4.46, and environment 3.14 on the oriented scales, 0.56 standardized units from the electorate mean. The single-candidate rows hold one perceived location fixed and optimize the other alone; their gains are relative to mismatch under Harris and Trump and are point estimates. Applying the winner identity of the main text to the Harris–Trump-plus-centrist menu gives winner-platform mismatch of 1.000, 1.350, and 2.146 times the best single-platform value when the centrist, Harris, or Trump platform wins.

7. Supplementary Benchmark Detail

The following supplementary Gaussian bound provides interpretive context connecting effective dimensionality to menu mismatch.

Proposition 4 (Three-bundle upper bound for a multivariate Gaussian)**.** *For $\theta \sim \mathcal{N}(\mu, \Sigma)$, placing three bundles optimally along the first principal component yields*

$$D_3^* \leq \operatorname{tr}(\Sigma) - 0.8098\,\lambda_1, \tag{3}$$

where $0.8098 \approx 1 - 0.1902$ *and 0.1902 is the one-dimensional K=3 Gaussian distortion coefficient (Gersho and Gray 1992).*

Proof. For a standard normal scalar Z, symmetry gives three optimal centers $(-a, 0, a)$ with thresholds at $\pm a/2$. The centroid condition is $a = \mathbb{E}[Z \mid Z > a/2] = \phi(a/2)/[1 - \Phi(a/2)]$, whose positive solution is $a = 1.224006$; direct integration over the three cells gives mismatch 0.190174. Construct three bundles that vary only along the first eigenvector, using this quantizer scaled by $\sqrt{\lambda_1}$. The resulting mismatch is $\sum_{j \geq 2} \lambda_j + 0.190174\lambda_1 = \operatorname{tr}(\Sigma) - 0.809826\lambda_1$. The displayed coefficient is rounded conservatively. The unrestricted optimum can only be lower. □

8. Replication Manifest for Analysis Outputs

Table OA25 maps each output family to its production script and principal generated files.

Table OA25. Replication Manifest for Analysis Outputs

Output family	Production script	Principal artifact stems
Seven-item baseline, candidate positions, K=1–5 menus, bootstrap, subgroups, signed distance changes, attitudes, and time summary	code/jepop_revision_analysis_PP22.py	jepop_2024_baseline_PP22, jepop_positioning_PP22, jepop_kmenus_PP22, jepop_signed_distance_changes_PP22, jepop_consequences_PP22
Optimizer runs, metric sensitivities, issue profiles, K=3 composition and conditional-winner benchmarks, perception diagnostics, missingness, and time decomposition	code/jepop_pp22_diagnostics.py	jepop_optimizer_*_PP22, jepop_metric_sensitivity_PP22, jepop_issue_profiles_PP22, jepop_k3_conditional_winner_PP22, jepop_perception_*_PP22
Main and supplementary empirical tables	code/make_jepop_tables_PP22.py	tab_jepop_*_PP22, oa_jepop_*_PP22
Supplementary diagnostic tables	code/make_jepop_diagnostic_tables_PP22.py	tab_jepop_optimizer_stability_PP22, tab_jepop_k3_conditional_winner_PP22, oa_jepop_metric_sensitivity_PP22, oa_jepop_missingness_*_PP22
Centrist-third-candidate counterfactual	code/jepop_centrist_counterfactual_PP31.py	jepop_centrist_counterfactual_PP31, oa_jepop_centrist_counterfactual_PP31
Numerical verification trace	code/verify_pp22_numbers.py, code/verify_pp31_centrist.py	verification/NUMBER_VERIFICATION_PP22.md, verification/NUMBER_VERIFICATION_PP22.tsv, verification/NUMBER_VERIFICATION_PP31_centrist.tsv

Notes: PP22 is the frozen production-pipeline identifier retained by the current submitted archive. run_all_PP22_JEPOP.sh regenerates the complete output tree, and verify_pp22_numbers.py is the archive's executable numerical verifier. The core script implements the ANES variance-design bootstrap and stratum-centered PSU sandwich calculations. CSV files retain full precision; TeX tables apply display rounding.

References

Adams, James F., Samuel Merrill III, and Bernard Grofman. 2005. *A Unified Theory of Party Competition: A Cross-National Analysis Integrating Spatial and Behavioral Factors*. Cambridge, UK: Cambridge University Press.

Gersho, Allen, and Robert M. Gray. 1992. *Vector Quantization and Signal Compression*. Boston: Kluwer Academic Publishers.

Hinich, Melvin J. 1977. "Equilibrium in Spatial Voting: The Median Voter Result Is an Artifact." *Journal of Economic Theory* 16 (2): 208–219. https://doi.org/10.1016/0022-0531(77)90005-9.

Lindbeck, Assar, and Jörgen W. Weibull. 1987. "Balanced-Budget Redistribution as the Outcome of Political Competition." *Public Choice* 52 (3): 273–297. https://doi.org/10.1007/BF00116710.

Rao, J. N. K., and C. F. J. Wu. 1988. "Resampling Inference with Complex Survey Data." *Journal of the American Statistical Association* 83 (401): 231–241. https://doi.org/10.1080/01621459.1988.10478591.

Tallis, G. M. 1965. "Plane Truncation in Normal Populations." *Journal of the Royal Statistical Society: Series B (Methodological)* 27 (2): 301–307. https://doi.org/10.1111/j.2517-6161.1965.tb01497.x.