What Do Populist Voters Want?
A Nonparametric Revealed Preference Approach

Nina Nikiforova*

November 9, 2026

Click here for the latest version

Abstract

This paper proposes a nonparametric revealed-preference approach for identification and estimation of voter preferences from aggregate election returns in a multidimensional spatial voting model. I apply it to study the rise of populism in Germany, using the municipal-level results for four federal (Bundestag) elections between 2013 and 2025 and modeling political competition along the Left–Right and Anti-Elite dimensions. The results reveal a transformation in the basis of support for the radical-right populist party Alternative for Germany (AfD): what began primarily as anti-elite voting became increasingly aligned with the party's right-wing positions. By 2021, the left–right dimension accounts for as much as 70% of AfD support. Counterfactuals in 2025 suggest that a modest rightward shift by the center-right Christian Democratic Union (CDU) attracts about one-fifth of AfD voters if the AfD's position is fixed. When allowed to respond strategically, however, the AfD moves further right and reverses these gains; in equilibrium, the CDU does not move right. These results may explain why establishment parties struggle to reclaim populist voters despite apparent ideological openings.

*Bocconi University. E-mail: nina.nikiforova@unibocconi.it. I am grateful to Guido Tabellini for his constant support and guidance; and to Jérôme Adda, Cailin Slattery, Pietro Tebaldi and Francesco Trebbi for their generous advice and feedback. I thank Matilde Bombardini, Christian Cox, Catherine De Vries, Francesco Decarolis, Sarah Eichmeyer, Matthew Gentzkow, Sergei Izmalkov, Rafael Jiménez-Durán, Thomas Le Barbanchon, Nicolas Longuet-Marx, Massimo Morelli, Eduard Osipov, Carlo Schwarz, Piero Stanig, Edoardo Teso and Christopher Turansick for their helpful comments and suggestions. Part of this paper was written while I was a visiting scholar at the Haas School of Business, UC Berkeley, which I thank for its hospitality.

1 Introduction

Western democracies are experiencing profound political transformations, with the rise of populism being a prominent example. The populist vote share in Western democracies has roughly tripled since the 1990s and now exceeds a quarter of the electorate (Guriev and Papaioannou, 2022). This trend shows no signs of reversing: Trump returned to the White House in 2024, Giorgia Meloni has governed Italy since 2022, and the Alternative for Germany doubled its vote share between the 2021 and 2025 federal elections. Existing work links this phenomenon to issues such as trade (Autor et al., 2020; Colantone and Stanig, 2018), immigration (Dustmann, Vasiljeva and Piil Damm, 2019), and identity (Norris and Inglehart, 2019; Bonomi, Gennaioli and Tabellini, 2021). Yet evidence on individual drivers does not by itself provide a unified empirical account of how voters' preferences have evolved (Margalit, 2019).

In this paper, I ask and empirically investigate two questions: first, *what do populist voters want?* and second, building on this, *how should mainstream parties respond to the rise of populism?* To answer these questions, I develop an empirical model of multidimensional spatial voting (Downs, 1957) that can be identified and estimated using aggregate vote shares. The model accommodates instrumental variables, minimizes reliance on functional-form and distributional assumptions, and enables counterfactual analysis of spatial competition. Recent empirical work in political economy develops econometric models to estimate voters' preferences (Berry, Cox and Haile, 2025; Danieli et al., 2026; Longuet-Marx, 2026). I connect this empirical agenda to spatial voting theory.[1]

I apply this framework to study the rise of populism in Germany, where the radical-right populist party Alternative for Germany (AfD) leads the national polls and has won several state elections. I begin by recovering the preferences of AfD voters over 2013–2025 along the Left–Right and Anti-Elite dimensions. I find that the party's earlier support was grounded in anti-elite sentiment, whereas its more recent support is rooted in its radical-right stance. Counterfactuals for 2025 suggest that, holding the AfD's position fixed, a modest rightward shift (for instance, on immigration policy) by the center-right Christian Democratic Union (CDU) could attract about a fifth of AfD voters. However, in a sequential repositioning game based on the estimated voters' preferences, the AfD responds to a rightward CDU shift by moving further right, largely reversing the CDU's gains. Anticipating this response, the CDU stays put in most equilibria. To my knowledge, this paper provides the first revealed-preference evidence on what populist voters want, and the first empirical equilibrium analysis of strategic repositioning between an establishment and a populist party.

The first part of the paper develops an approach to identifying and estimating voters' preferences using variation in aggregate election returns, the socio-demographic composition of the electorate, and the positions of the candidates in a multidimensional policy space. Each voter is represented by an ideal point in the policy space and chooses the candidate closest to their ideal point in weighted Euclidean distance. Voters'

[1] Merlo and De Paula (2017) pioneered the estimation of multidimensional spatial voting models, but under specific functional-form and distributional assumptions and without addressing the endogeneity arising from the correlation between voters' preferences and candidates' positions.

preferences are therefore described by two objects: the distribution of ideal points and the relative weights of the dimensions. Both may vary across voter types, defined by observable socio-demographic characteristics. The approach applies to any setting in which voters of the same type face different sets of candidates in the same election, such as simultaneous elections for several offices (the United States, Italy), ranked-choice voting (Australia, Ireland), or district-based representation, as in my application, where voters of the same type face different candidates across electoral districts within a state.

Two observations are key. First, the weights determine the regions of the policy space in which voters prefer one candidate over another. Because voters of the same type face different candidates in different districts, their choices across districts reveal which weights are consistent with the data. This variation set-identifies the weights. Second, there are generally infinitely many ideal-point distributions consistent with the data, as the data pin down only the mass of voters on each joint voting pattern, not the exact distribution that generated it. Point identification of a single distribution would therefore require additional structure, either on the choice model or on unobserved heterogeneity. Instead, I bound the target parameters, building on the method of Tebaldi, Torgovitsky and Yang (2023). Target parameters can describe both voters' preferences and counterfactual behavior, such as the share of voters with ideal points in a given area of the policy space, or the share who would abandon a party if it shifted its position by a given amount.

In the second part of the paper, I apply this approach to study the rise of the Alternative for Germany (AfD). Following the conceptualization of populism in political science (Mudde, 2004), I define the policy space by two dimensions: Left–Right and Anti-Elite. In a companion paper, Nikiforova (2026), I show that these are the two central dimensions of vote choice in Germany in recent years. I estimate the preferences of AfD voters separately for each federal election over 2013–2025 to trace their evolution. I then use the estimated preferences to evaluate counterfactual party repositioning.

Germany's mixed electoral system combines a plausibly sincere measure of party choice with geographic variation in political supply. Voters simultaneously cast one vote for a district candidate, nominated by a party and elected by plurality, and another for a party list, which determines the proportional allocation of seats. In the district vote, supporters of a candidate with no chance of winning may switch strategically to a more viable one; the list vote is free of this incentive, as long as the party is expected to clear the 5% threshold. I therefore use party-list vote shares and treat choices among the six major parties as sincere. District candidates run under party labels and carry their party's campaign in their districts, so I proxy locally perceived party positions with their positions, allowing the same party to be perceived differently across districts. I support this proxy with survey evidence and a permutation test.

For estimation, I use municipality-level socio-demographic data and election returns, together with rich within-state variation in the positions of local candidates. On the voter side, I model political preferences as i.i.d. within a socio-demographic group. These groups are defined as the interaction of age, gender, education, state, and year.[2] I restrict the sample to municipalities with fewer than 10,000 inhabitants, whose

[2]Since the model abstracts from turnout decisions, I use administrative turnout rates by state, year, and age–gender group

compositions vary the most and therefore carry the identifying variation for group-specific preferences. On the candidate side, I measure the positions of about 4,000 district candidates on two dimensions: Left–Right, which combines economic and cultural positions, and Anti-Elite, which captures distrust of political elites, of institutions, and of representative democracy. I combine two *candidate* surveys: the GLES Candidate Study and the nonpartisan *Whom to Vote For?* survey. Because the same survey items recur across election years, the resulting positions are comparable over time. The model relates the vote shares observed in each municipality to the latent ideological preferences of its voters, by socio-demographic group, and to the positions of the candidates running in its district.

To achieve identification, I assume that preferences are common across electoral districts within each socio-demographic group, state, and year. This implies that within-state variation in candidate positions is unrelated to unobserved voter heterogeneity across districts, conditional on the observed socio-demographic composition of municipalities. This is consistent with institutional details. Because the proportional allocation of Bundestag seats is determined by party-list votes, district candidacies are embedded in a broader state-level process of selecting candidates for parliament, in which state party leaderships act as gatekeepers. Candidates are therefore not solely tailored to the preferences of specific local electorates. It is also consistent with the data: within the sample of municipalities below 10,000 inhabitants and conditional on state and party, candidates' positions on both dimensions are not correlated with observed socio-economic characteristics of municipalities, such as unemployment, the migration rate, and the share of residents receiving basic income support. I test the robustness of the results to this assumption in four ways: first, by focusing on very small municipalities, which are least likely to be targeted by candidates; second, by restricting the analysis to districts with non-contender AfD candidates; third, by allowing for preference heterogeneity between electoral districts in which AfD candidates are more or less conservative; and fourth, by focusing on municipalities near district borders within states. The results remain qualitatively similar across all specifications.

I start by estimating the preferences of AfD voters. I first estimate the weight they attach to the Anti-Elite dimension relative to the Left–Right dimension. In Western Germany, the Left–Right dimension carries roughly three to five times as much weight as the Anti-Elite dimension throughout the period, whereas in Eastern Germany the Anti-Elite dimension carries substantial weight in every year and is estimated to outweigh the Left–Right dimension in 2013 and 2025. I then turn to the distribution of AfD voters' ideal points. The results reveal a recomposition of the party's electorate: in 2013, up to 45% of AfD voters held strongly anti-elite views and at least a quarter held only moderate center-right positions, whereas by 2025 the strongly anti-elite share had fallen and the radical-right share had risen sharply, while the party's own position moved only modestly. Combining the estimated weights and ideal points with party positions, I then ask what draws voters to the AfD: its radical-right position or its anti-elite stance. Over time, AfD support has become increasingly grounded in the former, with up to 70% of AfD voters in 2021–2025 supporting the party primarily for its radical-right ideology.

to reconstruct the socio-demographic composition of actual voters.

I next run two repositioning exercises to understand the intensity of competition between the establishment parties and the AfD. First, I explore the sensitivity of the AfD electorate to shifts in the AfD's own position towards more pro-elite or more moderate Left–Right stances (own-position *elasticity*). The results suggest that until 2025 voters were relatively insensitive to small shifts in party positions. This indicates that competition between the AfD and other parties used to be relatively weak: few of its voters lay near the boundaries with other parties' attraction regions. The picture changes in 2025, when even a small moderation in the AfD's Left–Right position would lead to substantial electoral losses.

Second, I ask how establishment parties could compete with populist challengers. I focus on the largest center-right establishment party, the CDU, and, holding the AfD's position fixed, estimate how much of the AfD electorate it could capture by shifting towards more conservative and more anti-elite positions (cross-position *elasticity*). The results suggest that a small repositioning on the Left–Right dimension would capture no more than 10% of the AfD electorate before 2025. In 2025 the same shift would allow the CDU to capture about 20% of the AfD electorate. Crucially, hardening the party's anti-elite positions is not nearly as effective in any year.

I then use the recovered preferences as primitives to empirically estimate the equilibrium of a game with strategic interactions for 2025: both the CDU and the AfD can reposition, while the positions of the remaining parties are held fixed. I consider a Stackelberg game in which the CDU leads and the AfD follows. If the AfD stays put, it is optimal for the CDU to move rightward, because the gains on the right flank (from the AfD) exceed the losses on the left. The AfD's optimal response to a rightward move by the CDU, however, is to move further right itself, recapturing the most right-wing (and relatively pro-elite) voters. In the Stackelberg equilibrium, therefore, the CDU stays put. This exercise offers one answer to the question of why mainstream parties do not adjust their platforms as much as a static calculation would suggest: the anticipation of the populists' response.

Finally, I discuss the assumptions underlying the framework and what they imply for the results. Among other things, I show that (i) the model can be read as a probabilistic voting model in which the error distribution is induced by the distribution of ideal-point shocks rather than assumed directly; (ii) measurement error in candidate positions shifts the bounds towards the uninformative composition; (iii) the consequences of omitting a dimension depend on its relevance and on its correlation with the included ones; and (iv) estimates obtained with the approach of Merlo and De Paula (2017) lead to partly different conclusions from those of my approach. I also re-estimate the model with Economics and Culture as the two dimensions, and find a rise in the role of cultural concerns relative to economic ones.

Related Literature. First, this paper contributes to a growing body of research that estimates structural models of voting in elections from observed election returns (Kawai and Watanabe, 2013; Ujhelyi, Chatterjee and Szabó, 2021; Iaryczower, Montero and Kim, 2024; Kawai and Sunada, 2025; Berry, Cox and Haile, 2025; Longuet-Marx, 2026). This paper contributes by developing a novel nonparametric approach to identifying

and estimating a multidimensional spatial voting model from aggregate election results, in which individuals vote for the candidate closest to them in the weighted policy space. Of these papers, only Iaryczower, Montero and Kim (2024) estimates a spatial voting model, and only in one dimension. As noted in Merlo and De Paula (2017), estimation of a single-dimensional model is considerably easier, since the problem of separately identifying *weights* from ideal points does not arise.

This paper is closely related to Degan and Merlo (2009), who first formalize the partitioning of the policy space under the multidimensional spatial voting model using Voronoi cells; and to Merlo and De Paula (2017), who exploit this partitioning to prove point identification of the ideal-point distribution and the weights under a large-support assumption on candidates' positions, and jointly estimate the best-fitting preference distribution and weights by minimum distance with a sieve density. The large-support assumption delivers point identification, but is untenable in electoral settings.

This paper, in turn, relaxes the large-support assumption, at the cost of set rather than point identification, and asks what the observed candidate configurations alone identify and how to estimate the resulting objects without imposing a functional form on the distribution of ideal points. Building on Tebaldi, Torgovitsky and Yang (2023), I establish set identification of the dimension weights and ideal-point distributions and adapt a linear-programming procedure for estimating sharp bounds on parameters describing both voters' preferences and counterfactual voting outcomes.[3] The approach exploits a similarity between the geometry of choice induced by spatial voting in my setting and that induced by quasilinear demand models in their applications: both partition the space of latent preferences into regions that imply the same choices across choice environments. The weights and ideal-point distributions are informed by distinct features of this geometry.

Second, this paper relates to a literature that employs a range of methodologies to recover voters' preferences, including the use of survey responses on political attitudes and self-placement, which are then mapped into ideological positions (Fiorina and Abrams, 2008; Algan et al., 2017; Hangartner et al., 2019; Gethin, Martínez-Toledano and Piketty, 2022; Desmet, Ortuño-Ortín and Wacziarg, 2024; Draca and Schwarz, 2024; Danieli et al., 2026). Such methods face two common shortcomings: voters' views do not necessarily translate into voting choices, and some views may be stigmatized or socially desirable, and thus not reported truthfully. Survey-based stated preferences can therefore differ substantially from the revealed preferences inferred from electoral outcomes, and the extent of this divergence can only be assessed by estimating preferences from actual election data.

Third, this paper contributes to the large and growing literature on the causes of the rise and success of populist parties and leaders in the Western world. An extensive body of work examines the various drivers of populist voting, including economic factors such as trade and the 2008–2009 Global Financial Crisis (Algan et al., 2017; Colantone and Stanig, 2018; Autor et al., 2020; Guiso et al., 2021; Dippel et al., 2022), immigration (Dustmann, Vasiljeva and Piil Damm, 2019; Alesina and Tabellini, 2024), and cultural factors (Algan et al., 2018; Margalit, 2019; Gennaioli and Tabellini, 2025), among others. While these drivers can be mapped

[3]See also Gu, Russell and Stringham (2026) for a more general treatment of this class of models.

onto distinct ideological dimensions (see, e.g., Rodrik, 2021, for a clear formalization of this logic and a meta-analysis of mechanisms related to globalization shocks), they do not reveal the relative importance of each, or their interplay. This paper addresses this gap by estimating voters' ideal-point distributions and the relative weights of policy dimensions in a multidimensional spatial model, and tracing how preferences evolve over time. In this respect, the paper is closely related to Danieli et al. (2026), who recover the preferences of voters of radical-right populist parties in Europe. Unlike the present paper, however, that study relies on survey-based self-reported vote choices and issue preferences.

I also contribute to the literature on how established parties respond to new entrants and to niche or populist challengers (Palfrey, 1984; Meguid, 2005; Abou-Chadi and Krause, 2020; Spoon and Klüver, 2020; Krause, Cohen and Abou-Chadi, 2023), by empirically estimating best responses and the Stackelberg equilibrium of a platform-choice game between an establishment party and a populist challenger.

Lastly, this paper contributes to the debate on the relative importance of ideological and Anti-Elite dimensions in populist voting. A large part of the literature emphasizes the role of right (or left) positions in shaping populist support (Acemoglu, Egorov and Sonin, 2013; Mudde, 2013; Norris and Inglehart, 2019; Guenther and Nunnari, 2025), whereas another strand highlights the role of populist attitudes distinct from Left–Right ideology (Mudde, 2004; Akkerman, Mudde and Zaslove, 2014; Van Kessel, 2015). Several studies have sought to compare the importance of these two factors directly using survey and experimental data (Silva, Fuks and Tamaki, 2022; Neuner and Wratil, 2022; Neuner et al., 2025). This paper shows that populist voting can take different forms, with different factors prevailing across regions and over time.

The rest of the paper is structured as follows. Section 2 lays out an approach to recover voters' preferences, and Section 3 discusses the scope of the approach. Section 4 introduces the German institutional context and the data. Section 5 presents the estimation procedure and the identification strategy. Section 6 reports the results on the estimated preferences; Section 7 runs spatial competition counterfactuals; Section 8 studies the repositioning game between the CDU and the AfD. Section 9 discusses possible extensions and adaptations of the model. Section 10 concludes.

2 Econometric Model

2.1 Voters, candidates and elections

There is a set of elections $\mathcal{E}$, each indexed by e. In election e, each candidate j from the set $\mathcal{J}_e$ has positions Y_{je}. This is an element of the policy space $\mathcal{P}$, a closed and bounded subset of $\mathbb{R}^K$, known to the econometrician. Each dimension $k = 1, \ldots, K$ represents positions on a set of related issues. For example, on a Left–Right dimension, $Y^1_{je} > Y^1_{\ell e}$ could indicate that candidate j is further to the right than candidate ℓ. While my analysis can be extended to higher dimensions, I focus on $K = 2$ throughout, a choice largely dictated by computational tractability.

Voters, indexed by i, have spatial preferences over $\mathcal{P}$. In particular, each i has an ideal point $X_i \in \mathcal{P}$,

and votes sincerely for the candidate whose positions are the nearest under the weighted Euclidean distance (Downs, 1957; Hinich and Munger, 1996):

$$dist(X_i, Y_{je}) = \sqrt{(X_i - Y_{je})' W (X_i - Y_{je})}\,, \tag{1}$$

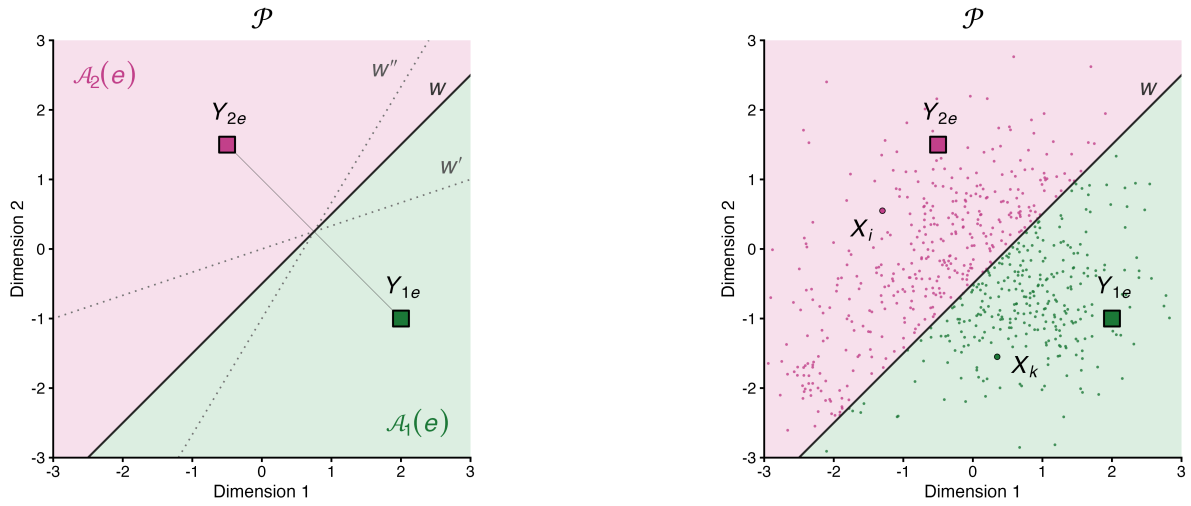
where $W = \text{diag}(1, w)$ for a positive weight $w \in \mathbb{R}_+$.[4] The choice of voter i is then

$$i \text{ votes for } j^* \text{ in election } e \quad \Longleftrightarrow \quad j^* = c(X_i, e) \equiv \underset{j \in \mathcal{J}_e}{\arg\min}\ dist(X_i, Y_{je})\,. \tag{2}$$

In practice, given the weight w, the collection of positions of all candidates in election e, $\mathbf{Y}_e = \{Y_{je},\ j \in \mathcal{J}_e\}$ partitions $\mathcal{P}$ into regions of voters who choose different candidates. The *attraction region* of candidate j in election e is the set of ideal points X_i that vote for j:

$$\mathcal{A}_j(e) = \{X_i \in \mathcal{P} : c(X_i, e) = j\}. \tag{3}$$

Figure 1 illustrates the setup. Figure 1a shows the attraction regions of a two-candidate election for a fixed weight w, and how alternative weights w', w'' would change their geometry. Figure 1b plots hypothetical ideal points of voters in the policy space $\mathcal{P}$, each colored by the closest candidate.



(a) Two-candidate attraction regions (b) Voters' ideal points

Figure 1. Illustration of a two-candidate election setup

Notes: Election e with candidates $j = 1, 2$ in a two-dimensional policy space $\mathcal{P}$ is illustrated. Squares mark the candidates' positions Y_{1e} and Y_{2e}. The background colors show the attraction regions $\mathcal{A}_1(e)$ (green) and $\mathcal{A}_2(e)$ (pink) under weight w, separated by the black solid line. The gray dotted lines show the boundary between the two attraction regions under different weights w' and w''. Changing the weight rotates the boundary around the midpoint of the thin line connecting the two candidates. In panel (b), each dot is a voter's ideal point X_i, colored by the candidate the voter is closest to under w; X_i and X_k are two example voters.

The primitives of interest are the weight w and the distribution of voters' ideal points X_i. These

[4] Only the relative weights matter. Hence, normalizing one of the weights to 1 is without loss of generality.

primitives determine election outcomes for any configuration of candidate positions $\mathbf{Y}_e$.

2.2 Election returns and identifying assumption

The setup so far concerns individual voters, but the econometrician typically observes election returns: vote shares aggregated over administrative units. Anticipating the application, I refer to states, municipalities, and socio-demographic groups, though other settings may use different units and groups (see Section 3).

Formally, each voter i lives in municipality M_i, located in state S_i, and belongs to the observed socio-demographic type D_i. Let (s, m, d) denote a realization of (S_i, M_i, D_i). Each state is divided into electoral districts, each with its own election. Each municipality m belongs to a single district, and is associated with one election, $e(m)$. The data consist of the share of voters choosing candidate j in municipality m, $\mathcal{s}_j(m)$, and the share of eligible voters in municipality m and socio-demographic group d, $q_{m,d}$, $\sum_d q_{m,d} = 1$.[5]

Let $f(x \mid m, d)$ denote the density of ideal points of type-d voters in municipality m.[6] Let $\mathcal{F}$ denote the set of all such conditional density functions, which I leave otherwise unrestricted. Under the choice rule (2), the observed vote share of candidate j in municipality m is $\mathcal{s}_j(m)$:

$$\mathcal{s}_j(m) = \sum_d q_{m,d} \int_{\mathcal{A}_j(e(m))} f(x \mid m, d)\, dx. \tag{4}$$

My key identifying assumption links voters across municipalities and elections within a state.

Assumption 1a. For every state s, type d, and municipalities m, m' in state s, $f(x \mid m, d) = f(x \mid m', d)$ for all x. I write $f(x \mid s, d)$ for this common density.

Intuitively, the assumption says that, conditional on types, preferences have the same distribution across municipalities within a state. This provides identifying variation since, within a state, different elections display different candidates' positions. Therefore, one can learn about preferences within a state by comparing electoral outcomes across the state's districts, conditioning on differences in socio-demographic composition. This is a selection-on-observables assumption: candidates' positions may differ across elections for reasons related to the type composition of the electorate, but not to voters' ideal points beyond their type. It rules out that voters sort into elections according to the candidates on offer, and that parties tailor candidates to local preferences that the types do not capture.

Assumption 1a can be weakened to an instrumental-variable condition. Suppose there is an observed variable Z_i, and let $f(x \mid z, s, d)$ denote the density of ideal points of type-d voters in state s with $Z_i = z$.

[5] The model does not accommodate turnout decisions. In this section I assume that voters who turn out have the same type composition as the eligible population. In the application, I relax this by reconstructing the type composition of actual voters from administrative data on turnout rates by socio-demographic group and state.

[6] I assume that, conditional on $(M_i, D_i) = (m, d)$, the distribution of X_i is absolutely continuous with respect to Lebesgue measure on $\mathbb{R}^K$. Since $S_i = s(M_i)$, this also ensures the existence of a state-level conditional density $f(x \mid s, d)$. Both densities are defined for conditioning values in their respective supports.

Assumption 1b. For every state s, type d, and z, z' such that $f(\cdot \mid z, s, d)$ and $f(\cdot \mid z', s, d)$ are defined, $f(x \mid z, s, d) = f(x \mid z', s, d)$ for all x.

Assumption 1b has the two parts of a standard instrument. *Exogeneity* is the stated condition: the distribution of ideal points does not vary with Z_i, so the instrument is unrelated to the unobserved voter heterogeneity that could be correlated with the platforms. *Relevance* is required for the condition to have identifying content: the platforms must vary with z. Assumption 1a is the special case $Z_i = M_i$, in which the municipality is the instrument: it determines the candidates faced and, within a state and type, is unrelated to ideal points. Section 5.2 motivates Assumption 1a in the application and discusses robustness tests, some of which make use of Assumption 1b.

Match the data. Under Assumption 1a, a density $f \in \mathcal{F}$ predicts the vote shares in municipality m in state s:

$$\sigma_j(m; f) = \sum_d q_{m,d} \int_{\mathcal{A}_j(e(m))} f(x \mid s, d)\, dx,$$

and is consistent with the data if

$$\underbrace{s_j(m)}_{\text{observed}} = \underbrace{\sigma_j(m; f)}_{\text{predicted}} \qquad \text{for all } j \in \mathcal{J}_{e(m)} \text{ and all } m. \tag{5}$$

Target Parameter. To learn about voters' preferences and conduct counterfactual analysis, I consider quantities of interest that are functionals of preferences, that is, of f, w or both. For instance, the share of type-d voters in state s whose ideal points are left-leaning on the first dimension, $T = \{x \in \mathcal{P} : x_1 \leq \underline{x}_1\}$ for a known threshold $\underline{x}_1$, is

$$\theta(s, d; f) = \int_T f(x \mid s, d)\, dx.$$

Other examples are the share of a specific candidate's voters with ideal points in a given region, or vote shares under counterfactual platforms. I refer to such quantities as *target parameters*, and write $\theta(f)$ when the arguments (s, d) are integrated over. The goal is to characterize the values of $\theta(f)$ consistent with the data and the assumptions.

The sharp identified set of ideal-point densities is $\mathcal{F}^* = \{f \in \mathcal{F} : f \text{ satisfies Assumption 1a and (5)}\}$, and the sharp identified set of the target parameter is $\Theta^* = \{\theta(f) : f \in \mathcal{F}^*\}$, the set of all values of the target parameter consistent with the data and Assumption 1a, and no others.

2.3 Identification of the ideal-point distributions and counterfactual quantities

How can one characterize Θ^*? Since $\mathcal{F}$ is infinite-dimensional, a direct search over unrestricted densities is computationally infeasible. In the context of quasilinear discrete choice models, Tebaldi, Torgovitsky and Yang (2023) show that a similar infinite-dimensional problem admits a finite-dimensional formulation. Checking that a value belongs to Θ^* is equivalent to checking feasibility of a finite-dimensional system of

linear equations. I adapt their argument to the multidimensional spatial voting model (2). I begin by treating w as known and relax this assumption in Section 2.4.

Minimal relevant partition. Under Assumption 1a, voters of type d in state s share the distribution $f(x \mid s, d)$. The type can therefore be treated *as if* it were a single population voting in every election of the state. Denote by $\mathcal{E}(s)$ the set of elections of state s. For a voter with a fixed ideal point, a voting profile $\mathbf{j} = (j_e), e \in \mathcal{E}(s)$ records the candidate $j_e \in \mathcal{J}_e$ that the voter would choose if faced with each election's candidates. The set of ideal points producing profile $\mathbf{j}$ (a cell) is

$$\mathcal{X}_{\mathbf{j}} = \bigcap_{e \in \mathcal{E}(s)} \mathcal{A}_{j_e}(e).$$

The collection of available cells, $\mathbb{X} = \{\mathcal{X}_{\mathbf{j}} : \mathcal{X}_{\mathbf{j}} \neq \varnothing\}$, forms a partition of $\mathcal{P}$, with one cell for each feasible voting profile. Following Tebaldi, Torgovitsky and Yang (2023), I call this collection the *minimal relevant partition* (MRP). Each cell corresponds to a distinct voting profile: voters whose ideal points lie in the same cell would vote identically in every election, while voters in different cells would choose different candidates in at least one election.[7]

Figures 2a and 2b show the attraction regions in two hypothetical elections, e_1 and e_2, with candidate positions (Y_{1e_1}, Y_{2e_1}) and (Y_{1e_2}, Y_{2e_2}), respectively. Figure 2c shows the resulting MRP, whose four cells correspond to four distinct voting profiles. For example, a voter whose ideal point lies in the pink cell would choose candidate 2 in either election, whereas a voter whose ideal point lies in the purple cell would choose candidate 1 in election e_2. The MRP leaves the precise location of ideal points within each cell unrestricted. Figure 2d refines the partition by the boundary of a region T defining a target parameter, which becomes the union of areas $\mathcal{X}_1, \mathcal{X}_2, \mathcal{X}_3$.

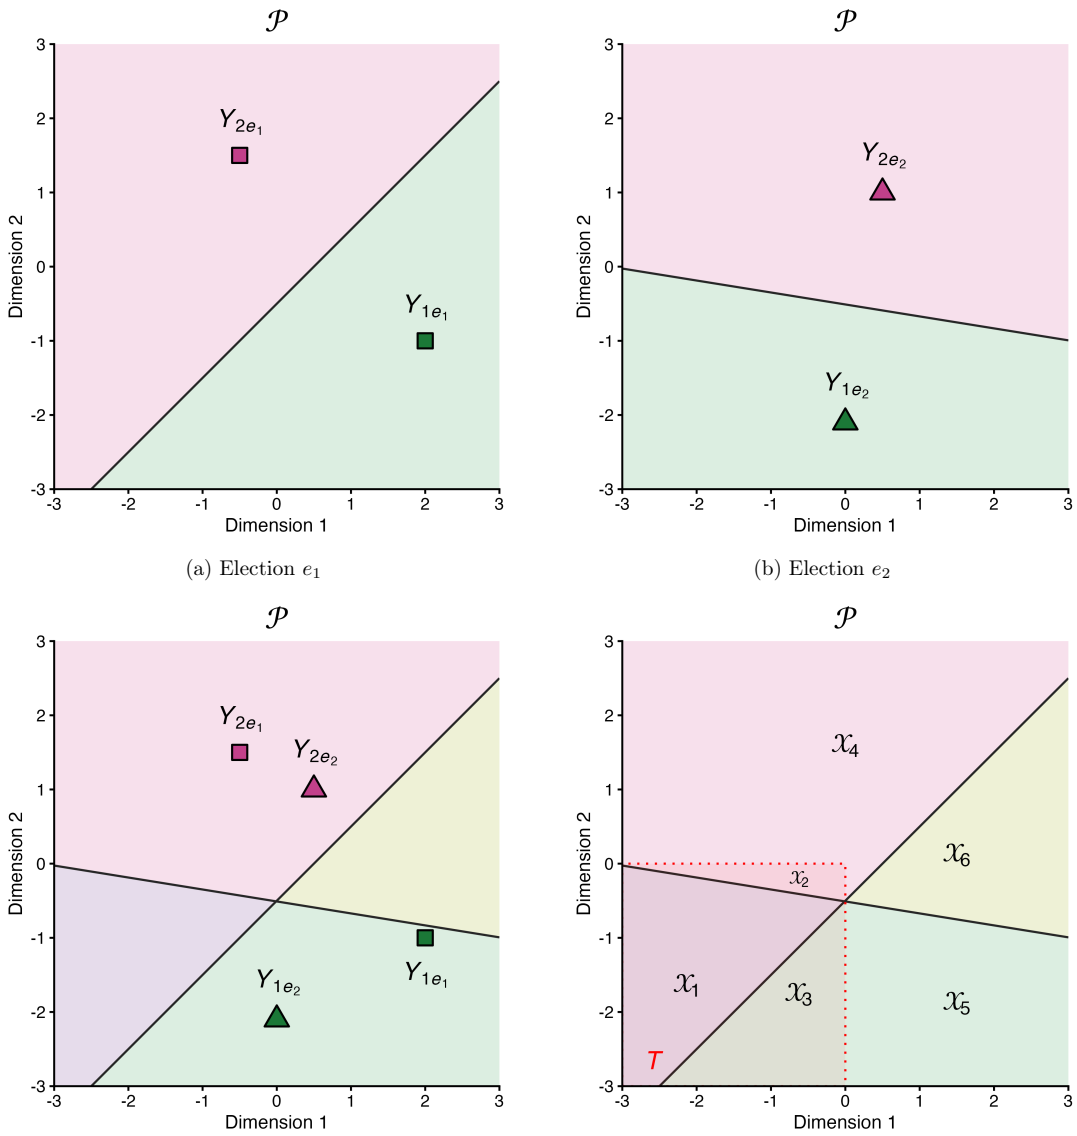
From densities to cell masses. The MRP matters because the returns depend on f only through the mass it places on each cell. For election e, let $\mathbb{X}_j(e) = \{\mathcal{X} \in \mathbb{X} : \mathcal{X} \subseteq \mathcal{A}_j(e)\}$ denote the cells contained in j's attraction region. By (5), the share of type-d voters choosing j in election e is the mass of $f(\cdot \mid s, d)$ on the cells in $\mathbb{X}_j(e)$. The same holds for the target parameter once the partition is refined by the regions that define it. In Figure 2d, T is the union of $\mathcal{X}_1, \mathcal{X}_2, \mathcal{X}_3$, so

$$\theta(s, d; f) = \int_{\mathcal{X}_1} f(x \mid s, d)\, dx + \int_{\mathcal{X}_2} f(x \mid s, d)\, dx + \int_{\mathcal{X}_3} f(x \mid s, d)\, dx,$$

which depends not on f per se, but only on the mass that f places on the three cells.

Attention can therefore shift from densities to *cell masses*, $\phi(\mathcal{X} \mid s, d) = \int_{\mathcal{X}} f(x \mid s, d)\, dx$. A value t belongs to Θ^* if and only if some vector of cell masses reproduces the observed returns and yields t. For

[7] For target parameters involving counterfactual platforms, $\mathcal{E}(s)$ is augmented by the corresponding hypothetical elections. For targets involving regions of the policy space, such as T above, $\mathbb{X}$ is further refined along their boundaries.



(a) Election e_1

(b) Election e_2

(c) Minimal Relevant Partition (MRP)

(d) MRP with Target Parameter

Figure 2. Example of partitions of the policy space in two elections with two candidates each, shown separately in panels (a) and (b), combined into the minimal relevant partition in panel (c), and refined by a target parameter in panel (d).

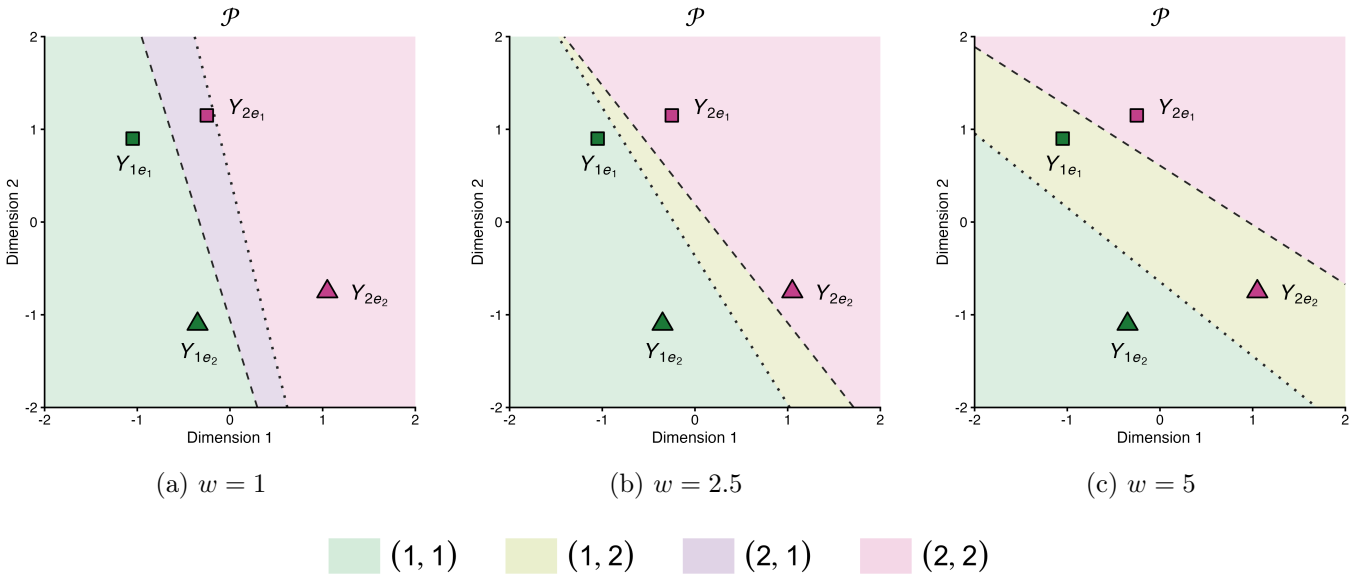
Notes: The figure illustrates a hypothetical setting with two elections, e_1 and e_2, each between two candidates, in a two-dimensional policy space $\mathcal{P}$, for a fixed weight w. Squares mark the positions Y_{1e_1} and Y_{2e_1} of the candidates in election e_1, and triangles the positions Y_{1e_2} and Y_{2e_2} of the candidates in election e_2. In panels (a) and (b), the green region is the attraction region of candidate 1 and the pink region that of candidate 2, separated by the solid line. Panel (c) overlays the partitions of panels (a) and (b), yielding the Minimal Relevant Partition (MRP): green (pink) cells attract voters to candidate 1 (2) in both elections, and the yellow-green (purple) cell attracts voters to candidate 1 (2) in e_1 and to candidate 2 (1) in e_2. Panel (d) presents a finer MRP that also includes the region T corresponding to a target parameter. $\mathcal{X}_n$ denotes the unique cells that characterize observed behavior in elections e_1 and e_2 or the target parameter.

a scalar target parameter, Θ^* is then an interval whose endpoints are the minimum and maximum of the target parameter over the cell masses consistent with the election returns; when the target parameter is linear (or linear-fractional) in the masses, as in the example and in all target parameters of the application, both endpoints can be computed by linear programming. Appendix A demonstrates these results formally in a spatial voting context by adapting an argument from Tebaldi, Torgovitsky and Yang (2023).

The approach therefore answers the questions of interest without restricting the ideal-point distribution to any functional form, and the bounds it delivers are sharp. The width of the bounds reflects what the election returns leave unidentified, and the conclusions about voters' preferences and counterfactuals rest on the observed returns rather than on an imposed functional form.

2.4 Identification of the weight

The previous section treated w as known. When w is unknown, the same cell-mass representation applies, but each value of w yields a different Minimal Relevant Partition, $\mathbb{X}_w$, and hence a different set of cells that can carry mass. A weight is consistent with the data if some cell masses on $\mathbb{X}_w$ reproduce the observed returns. The identifying information thus comes from *which* cells, or voting profiles, a weight permits. This section provides the intuition; Appendix B formalizes the argument, and Appendix B.3 derives sharp bounds on target parameters that allow for all weights consistent with the data.



(a) $w = 1$ (b) $w = 2.5$ (c) $w = 5$

Figure 3. Minimal relevant partitions induced by the same candidate positions under different values of w.

Notes: The figure illustrates the MRP in a hypothetical setting with two two-candidate elections, e_1 and e_2, under different values of the weight w. Squares mark the positions Y_{1e_1} and Y_{2e_1} of the candidates in e_1, and triangles the positions Y_{1e_2} and Y_{2e_2} of the candidates in e_2. Each panel overlays the two elections in the two-dimensional policy space $\mathcal{P}$. Each shaded cell corresponds to a voting profile (j, k): voters whose ideal points lie in it choose candidate j in e_1 and candidate k in e_2.

The key is to observe that marginal vote shares (observed) impose conditions on joint vote shares across elections, and, hence, on the cells that the minimal relevant partition must contain. For illustration, consider

a state with a single voter type and two elections between two candidates: e_1 between Y_{1e_1} and Y_{2e_1}, and e_2 between Y_{1e_2} and Y_{2e_2}. If candidate 1 receives 70% of the vote in e_1 and candidate 2 receives 90% in e_2, at least 60% of the common ideal-point distribution must generate the corresponding voting profile $(1, 2)$, choosing candidate 1 in e_1 and candidate 2 in e_2, and at least 20% the profile $(2, 2)$ (Appendix Table I1 provides these Fréchet bounds); hence, a weight whose MRP lacks either cell cannot reproduce the returns. By contrast, the profiles $(1, 1)$ and $(2, 1)$ are bounded below by 0 and need not belong to the MRP.

Figure 3 shows the MRP for the hypothetical elections e_1 and e_2 under different weights w. The cell $(2, 2)$ is available under every weight, so its lower bound has no identifying content here. The informative cell is $(1, 2)$: the partition under $w = 1$ does not contain it, so that weight cannot reproduce the observed returns and is ruled out. The partitions under $w = 2.5$ and $w = 5$ both contain it, so both can reproduce the observed returns and the two weights are not distinguishable.

The example shows that the weight is identified through the cells available in the MRP, not through their size in the policy space, and that it is set- rather than point-identified.[8] How informative the identified interval is remains an empirical question.

3 Scope of the Approach

The approach allows one to recover voters' preferences in a spatial voting framework from aggregate voting data. Identification rests on comparing voters who share the same preferences across different choice environments. This requires either a plausible selection-on-observables argument or an instrument that shifts the candidates faced without shifting preferences. A single choice environment is a special case in which this condition holds vacuously: one election already partitions the policy space and bounds the ideal-point distribution for a given weight, but with no second environment to compare, the weight itself is not identified and must be fixed a priori. Settings with one or several choice environments are pervasive, and closely related identifying assumptions are already in use in political economy. I organize the discussion around three such settings.

First, consider *simultaneous voting*, where the same voters face different choice environments because several elections take place at once. The main advantage of this setting is that selection on observables holds by construction: the same voters face every ballot, so the preference distribution is identical across the elections being compared, up to differences in turnout across races. Its main limitation is that the number of simultaneous elections is typically small, which yields a coarse partition of the policy space and is potentially less informative. This limitation does not apply where voters face many ballot propositions on the same day, as in California: Snyder (1996) and Gerber and Lewis (2004) exploit such returns with models in the spirit of NOMINATE (Poole and Rosenthal, 1985), estimating proposition cutpoints jointly with voter preferences.

[8] The identification argument does not require the weight to be common across types within a state, and goes through when w varies by type. In practice, however, type-specific weights are far more demanding on the data: informative type-level intervals would require either type-level choice data or much richer variation in candidate positions than is available in the application.

More generally, simultaneous voting arises in the United States, where multiple offices are often filled on the same day; in Australia, where both chambers of parliament are elected together; in Italy, where general elections fill the Chamber of Deputies and the Senate at once; and in Switzerland, where voters face several electoral and referendum ballots on the same day.

Second, *ranked-choice voting* (RCV) provides another setting in which the same voters face varying choice environments. A ranking can be read as a sequence of choices from shrinking choice sets: the ballot contains the same candidates throughout, but a candidate who has already been ranked is no longer available at subsequent stages. Lower-ranked choices can therefore be especially informative about voters' positions in the policy space. As with simultaneous elections, the identifying condition holds by construction. RCV is used in national elections in Australia, Ireland, Malta, and Sri Lanka, among others; in Alaska for state and federal elections; in Maine for primaries and federal general elections; and for local races in several US cities.

Third, consider *geographic representation*, in which different voters face different choice environments across localities. Such variation arises by design in many electoral systems: first-past-the-post, as in U.S. congressional elections and House of Commons elections in Canada and the United Kingdom; mixed systems combining first-past-the-post and proportional representation, as in Germany, Hungary, Japan and Mexico; and proportional representation with multiple districts, as in elections to Brazil's Chamber of Deputies.

In this setting, however, the identifying condition no longer holds by construction. Political economy models tie politicians' positions to voters' preferences, whether through opportunistic incentives, intrinsic congruence, or voters' responsiveness to politicians' cues (Osborne and Slivinski, 1996; Persson and Tabellini, 2002; Callander and Carbajal, 2022). A researcher must therefore either justify the comparison across choice environments or restrict attention to a single one. I now discuss approaches in recent empirical political economy, focusing on work in which candidates' characteristics include an ideological component.

Longuet-Marx (2026) exploits congressional district borders within states, comparing small groups of contiguous precincts that lie in different electoral districts but are otherwise very similar. The model in that paper differs from mine, but the same source of variation could be used within the framework developed here. In the terminology of this paper, the border design corresponds to Assumption 1b with the border group as the instrument: conditional on the group, observed demographics and time-invariant precinct effects, adjacent precincts share the same preference distribution while facing different candidate sets. The overlay of choice situations must then be done only within groups of precincts, which limits the available identifying variation. By pooling many such groups, or by imposing additional restrictions on how preferences may differ across groups, an econometrician may still recover more about the primitives of interest than from single elections alone.

Merlo and De Paula (2017) estimate voters' ideal-point distribution from aggregate returns across member states in the 1999 European Parliament elections, assuming party positions are independent of ideal points conditional on observed precinct characteristics. Danieli et al. (2026) combine individual survey data

with manifesto-based party positions across European countries and over time, and estimate the weights voters place on these positions in a probabilistic voting model with party fixed effects. Both rest on the assumption that, conditional on observables, residual preference heterogeneity is unrelated to party positions.

Iaryczower, Montero and Kim (2024) estimate voters' preferences from aggregate vote shares in elections to Brazil's Chamber of Deputies. With over one hundred candidates per state-year on average and a single ideological dimension, one choice environment already partitions the policy space finely; variation across environments is used mainly to recover how preferences shift with state demographics. Their instrument addresses a different concern, correlation between positions and unobserved candidate attributes such as charisma: each federal candidate's position is instrumented with the positions of same-party mayoral and state-level candidates in the most recent local cycle, which co-move with federal positions through the party but are plausibly unrelated to these attributes.

The model can also be extended to other features of voting behavior. Under strategic voting, some individuals abandon their most preferred candidate if a vote for her is perceived as *wasted*, which generates different effective choice sets across voters. Those who vote for underdog candidates can be treated as sincere and as facing the full choice set; for those who support viable contenders, the choice set should be restricted to viable candidates. This coarsens the partition and may widen the bounds, but makes them robust to strategic voting. The model can also accommodate valence as an additive, candidate-specific term shifting the boundaries of the partition.

I now turn to German federal elections, held under a mixed electoral system. Identifying variation comes from geographic representation: candidate positions vary across electoral districts within a federal state, while the preference distribution is modeled at the state level. Within a state, voters of a given socio-demographic group are assumed to draw their ideal points from the same distribution in every district. I first argue that this condition is plausible in the German setting, and then gauge the robustness of the results.

4 Application: The Rise of Populism in Germany

One of the main puzzles in the political economy literature is the rise of populism, and in particular radical-right populism, in Western democracies. The empirical literature has largely focused on identifying the causal mechanisms underlying the rise of populist politics, examining their effects on voters' preferences, and analyzing the contextual and individual-level factors that shape these preference responses (see Guriev and Papaioannou (2022) for an overview).

While understanding specific mechanisms is essential, it is equally important to develop a broader understanding of the preferences of populist voters, so as to draw policy implications and inform the competition strategies of mainstream parties. One recent attempt to recover populist voters' preferences is Danieli et al. (2026), who study the rise of populist radical-right parties in Europe using survey evidence over time. To my knowledge, my paper is the first to provide revealed-preference evidence on the ideological preferences of

populist voters at large.[9]

4.1 Context

Over recent decades, the German electoral landscape has undergone important changes, most notably the rise of the Alternative for Germany (AfD), a populist radical-right party. In 2017, for the first time since the 1950s, a radical-right party cleared the 5% threshold and entered the Bundestag. In the most recent federal election, in 2025, the AfD finished second, ahead of several established parties, including the Social Democratic Party (SPD), one of Germany's two major parties. Figure 4 shows the rise of the AfD over 2012–2026, plotting responses to the polling question *"Which party would you vote for if the Bundestag election were held next Sunday?"* alongside the results of the four Bundestag elections used in the paper (2013, 2017, 2021, and 2025). As of October 2026, the AfD leads national voting-intention polls with around 30%, about 10 percentage points ahead of the second-placed Christian Democrats (CDU/CSU, hereafter CDU).[10]

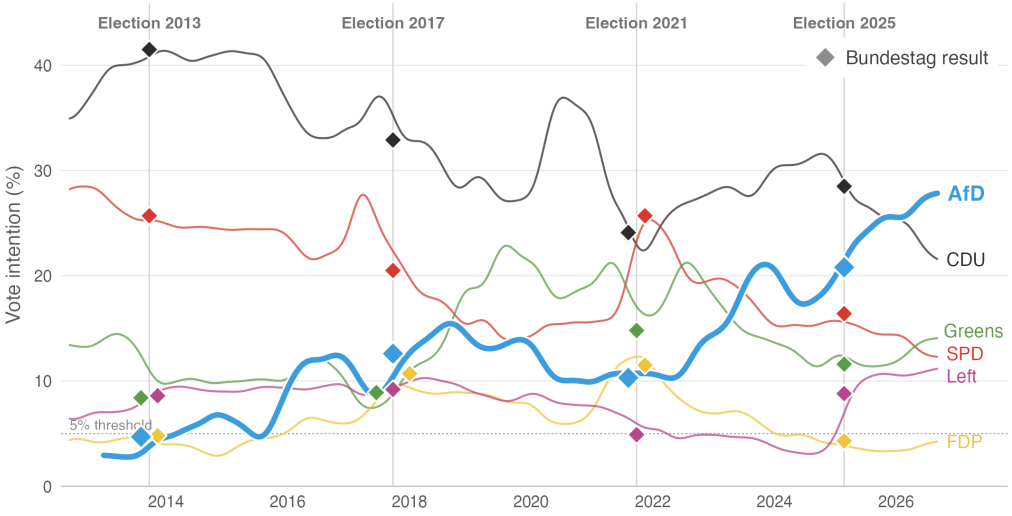


Figure 4. Federal voting intentions, June 2012 – September 2026.

Notes. Lines show the share of respondents intending to vote for each party in the question "Which party would you vote for if the Bundestag election were held next Sunday?", smoothed across 3,435 national polls by eight polling institutes (Allensbach, Forsa, Forschungsgruppe Wahlen, GMS, Infratest dimap, INSA, Kantar/Emnid, YouGov), compiled from ***wahlrecht.de****. The smoothed value on a given day is a Gaussian-weighted mean of all polls with a standard deviation of 60 days. Diamonds mark the official second-vote (party list) shares in the Bundestag elections of 2013, 2017, 2021 and 2025. The dotted line marks the 5% threshold for parliamentary representation. CDU denotes CDU/CSU; Greens and Left denote Bündnis 90/Die Grünen and Die Linke.*

The Alternative for Germany was founded in 2013 in opposition to the German government's response to the euro crisis. Ahead of the 2013 Bundestag election, its main concerns were the Greek bailout, adherence to the rule of law, and stronger participatory democracy (AfD, 2013). In 2015, the party split: many founding members left, and the conservative, anti-immigrant wing took over the leadership, which it has held since.

[9]Revealed-preference data sidestep concerns about survey misreporting, in particular the possible underreporting of populist vote choice, and make it possible to recover the relative importance of different dimensions of political competition, which is difficult to identify from survey data alone. Voters tend to adjust their reported positions towards those of their party, particularly on issues that matter little to them; when relative weights are estimated from survey data, such proximity may then be read as evidence that the dimension carries more weight than it in fact does.

[10]Politico, Poll of Polls, https://www.politico.eu/europe-poll-of-polls/germany/ (accessed October 2026).

Throughout the study period (2013–2025), the AfD was consistently ranked as the most right-wing and anti-elite of Germany's main parties (Jolly et al., 2022), and the only such party with a realistic chance of winning Bundestag seats.

Support for the AfD varies widely across regions and population groups. In 2025 it was the strongest party in all five eastern states, with 32–39% of the party-list vote, and in none of the western ones (Appendix Figure I10); within regions, it is more popular among middle-aged than among younger and older voters, and among men than among women (Figure 5). The estimation below accounts for these differences and highlights the East–West contrast, which may also inform the study of populism across countries with different institutional environments, such as eastern and western Europe.

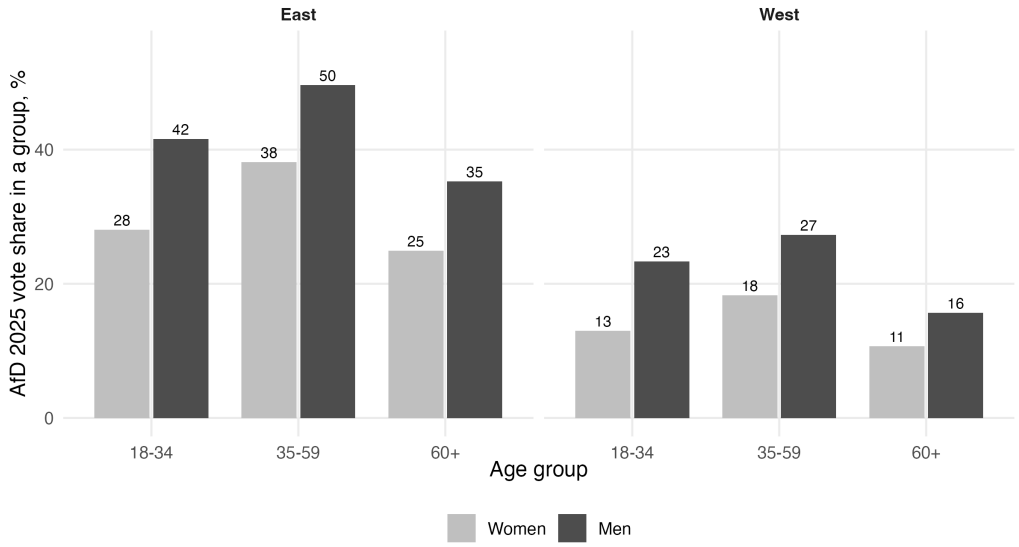


Figure 5. AfD party vote share by region, age group and gender, 2025.

Notes. Bars show the AfD share of valid second votes cast within each group. Source: Repräsentative Wahlstatistik of the Federal Returning Officer (Die Bundeswahlleiterin, 2025), an official statistic in which ballots in a random sample of polling districts are marked by the voter's gender, birth-year group and state. The six published age groups (18–24, 25–34, 35–44, 45–59, 60–69, 70+) are pooled into three by summing votes for visualization. East comprises Brandenburg, Mecklenburg-Vorpommern, Saxony, Saxony-Anhalt and Thuringia; West comprises the remaining states, including Berlin. Appendix Figure I2 shows the same patterns for 2013, 2017, and 2021.

My analysis draws on data from German federal parliamentary (Bundestag) elections. Under Germany's mixed electoral system, each voter casts two votes: a first vote for a direct candidate in a single-member district, and a second vote for a state-level party list. First votes are likely to be heavily shaped by the strategic incentives of first-past-the-post competition, whereas second votes among viable parties can be plausibly treated as sincere (more on this below). Although the two votes are distinct, the candidates on district ballots and on state lists largely overlap: around two-thirds of candidates on state lists also run in single-member districts, and district candidates are on average ranked higher on the lists (Appendix K.1). In what follows, I rely exclusively on second-vote party shares.

I assume that *candidates' positions in single-member districts shape perceived party positions*: I use

estimated candidate positions as proxies for parties' positions in each district. This allows me to exploit cross-district variation in candidate positions while preserving the assumption of sincere voting for the party list. To support this assumption, I use survey data from the 2021 Bundestag election, the richest available, which record each respondent's electoral district. Within federal states and within parties, the estimated Left–Right position of a party's district candidate predicts respondents' Left–Right placement of that party (Appendix Table I2). The relationship holds for the party a respondent votes for with the second vote, and is strongest among AfD voters. It is absent for parties other than the chosen one, and the candidate's position has no bearing on the respondent's own Left–Right self-placement, which speaks against local sorting of voters as an explanation. The results are robust to controlling for gender, age, and education. The same survey also shows that voters interact with and are familiar with their local candidates (Appendix Figure I4).[11]

Finally, I conduct a permutation test, and show that the observed assignment of candidates to districts fits the data better than random reassignments of the same candidates within each state. Specifically, I re-estimate the model with candidate positions randomly reassigned across districts within each state, so that every municipality keeps its own vote shares and demographics but faces another district's candidates. Across 100 such permutations for each region and election year, the observed assignment has the second-lowest fit error of the 101 values, averaged over all region-years ($p = 0.02$). Appendix Figure I1 plots the distribution.[12] Noise in the estimated candidate positions would, if anything, bias the results towards the uninformative composition, in the spirit of attenuation bias in linear regression (see Section 9).

Two additional institutional features are important for the analysis. First, state-level party elites play a substantial role in candidate selection. Although German election law formally assigns nominations in single-member districts to local party chapters, state party organizations act as de facto gatekeepers: they can block district candidates they consider unsuitable, and they determine the state party lists. The latter matters because candidates may enter the Bundestag through either district races or state lists, and most list candidates also run in districts. Candidate selection is, moreover, often uncontested, as elites tend to preselect candidates in advance, particularly when incumbents are involved (Detterbeck, 2016; Schüttemeyer et al., 2024). These considerations motivate treating states as separate units for preference estimation and inform the identification argument. Appendix K.1 provides further evidence on nomination practices, documenting persistence in nominations and the interplay between single-member districts and state lists.

Second, electoral districts are drawn primarily on the basis of population size. Unlike in the United States, redistricting in Germany is nonpartisan and district boundaries are largely stable. A permanent Constituency Commission reviews population changes in each electoral cycle and recommends adjustments

[11] In the political science literature, the influence of single-member district candidates on parties' list-tier outcomes in mixed electoral systems is known as the *contamination effect* (Ferrara, Herron and Nishikawa, 2005). It is documented across mixed systems (Herron and Nishikawa, 2001; Cox and Schoppa, 2002) and, for Germany, by Klingemann and Wessels (2001), Hainmueller and Kern (2008), and Gschwend and Zittel (2015); this literature mostly examines the effects of incumbency, candidate evaluations, and campaign effort. My assumption differs in that I focus on policy positions, assuming that candidates shape voters' perception of the party's position.

[12] Testing each of the eight region-years separately and combining the cell-level p-values with Fisher's method gives $p = 0.04$. The higher value reflects that the observed positions are not the best fit in 2013 and in Eastern Germany in 2017 and 2025.

to the Bundestag. Districts should remain within 15% of the average population, must be redrawn if the deviation exceeds 25%, and must lie entirely within a single state. In practice, only a small share of districts changes in any given electoral period.[13]

4.2 Data

Germany's 16 states are divided into 299 electoral districts, each containing several municipalities (about 11,000 in total). Voting and socio-demographic data are observed at the municipality level, and each electoral district has its own set of candidates. Appendix Figure I3 illustrates these levels.

Municipal-level data. I collect second-vote (party-list) shares at the municipality level for the six main parties: *AfD, CDU/CSU, FDP, Greens, Left,* and *SPD*. I abstract from coalition considerations and treat these parties as viable, so that voters have no strategic reason to vote for a party other than their preferred one. This allows me to assume sincere voting and to estimate voters' ideological preferences from their vote choices.[14,15,16]

I also collect socio-demographic data at the municipality level for each election year. For each municipality, I construct the composition of eligible voters across 24 age–gender–education groups: six age groups (18–24, 25–34, 35–44, 45–59, 60–69, and 70+), two genders, and two education levels (with and without higher education).[17] Each of these groups within a state is a unit of analysis, for example *women aged 18–24 with higher education in Saxony*, and I estimate the distribution of preferences separately for each group, state, and year. These data are measured precisely for 2021, using the 2022 Census. For 2013, 2017, and 2025, I take the 2021 municipal composition as a baseline and adjust it by the change in the size of each age–gender group in the corresponding county (*Kreis*; there are about 400 counties in Germany).

I restrict the benchmark sample to municipalities with fewer than 10,000 inhabitants.[18] Smaller units retain more distinctive socio-demographic compositions, which increases the variation available for estimating group-specific preferences. Appendix Table I4 reports how voting patterns differ across municipalities

[13] See, for example, Die Bundeswahlleiterin (n.d.) for the changes made before the 2025 Bundestag election.

[14] These parties jointly receive more than 90% of second votes in all years. I renormalize vote shares over them, so the analysis excludes the small group of voters who supported non-viable parties; in 2025, this group includes voters of the anti-elite BSW.

[15] Viability also requires that voters expected these parties to clear the 5% threshold. The exceptions are the FDP and the AfD in 2013 and the Left in 2021, at 4.7–4.9%, and the FDP in 2025, at 4.3%.

[16] I use municipalities rather than polling stations for two reasons. First, municipalities are geocoded, whereas polling-station locations must be recovered manually from addresses scattered across state and city portals. This is laborious for 2025 and effectively impossible for earlier elections. Second, polling-station returns cover only in-person votes, while a large share of Germans vote by mail (about a quarter in 2013 and 2017, nearly half in 2021 during the COVID-19 pandemic, and about a third in 2025), and postal votes are in most cases reported only at the municipality level. I am currently constructing polling-station-level results for 2025 by geolocating polling stations manually and imputing postal votes to them based on their observed in-person vote shares. These data will be especially useful for the border robustness check (see below).

[17] Age–gender composition is available for all municipalities, whereas age–gender–education composition is available only for municipalities with more than 10,000 inhabitants, municipal associations (*Gemeindeverbände*), and counties (*Kreise*). I combine the two sources: within each age–gender cell, I split a municipality's population into individuals with and without higher education using the education share of the finest unit for which it is published, that is, the municipality itself, otherwise its municipal association, and otherwise its county.

[18] This restriction effectively captures almost exclusively rural municipalities, making a separate urban–rural distinction unnecessary in this context. It also naturally excludes the city-states (Berlin, Bremen, and Hamburg) from the analysis.

below and above this threshold. Small municipalities lean more strongly towards the AfD (+2.1pp) and the CDU/CSU (+4.4pp) at the expense of the other parties, but this partly reflects state and district composition effects: once those are accounted for, the gaps fall to +0.7 and +3.6pp.

Candidate positions. Voters face distinct party-affiliated candidates across single-member districts. To measure candidates' ideological positions, I combine two candidate surveys:

1) the GLES Candidate Studies for the Bundestag elections 2013–2025, administered by the German Society for Electoral Research and the Leibniz Institute for the Social Sciences (GLES, 2014, 2018, 2023, 2025);

2) *Wen Wählen?* (*Whom to Vote For?*), a nonpartisan platform that fields a standardized questionnaire to election candidates.

Both surveys elicit candidates' own positions on a battery of policy issues and target candidates from parties likely to clear the 5% threshold, so candidates from all parties in my analysis were contacted in each election year.

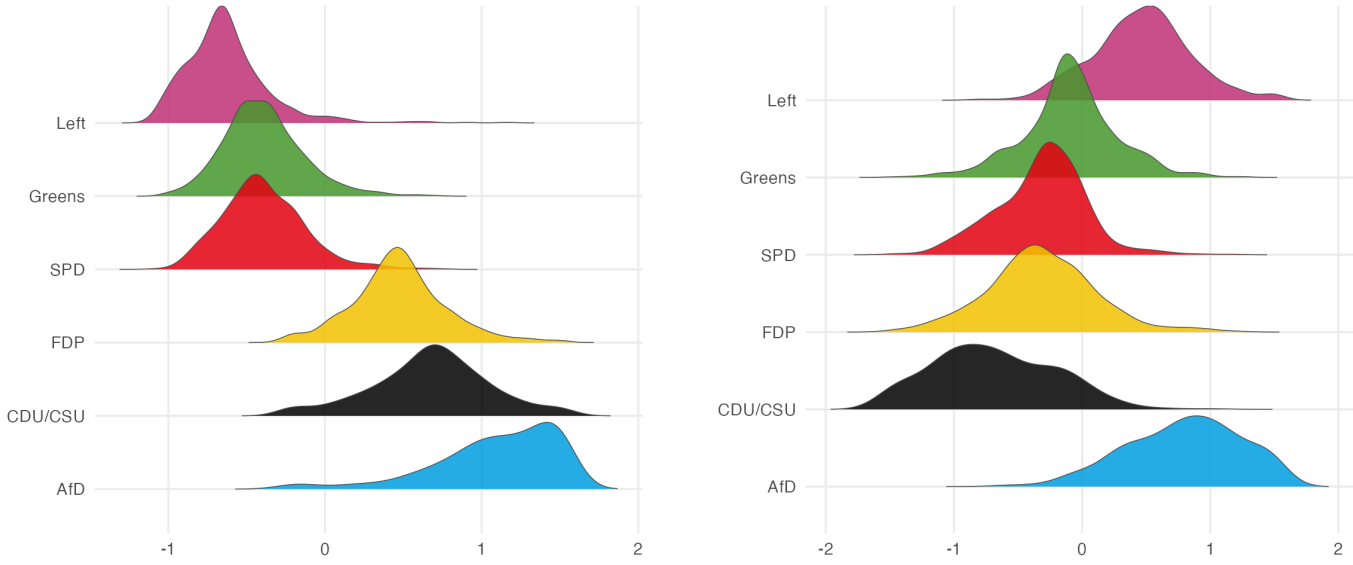
First, using the GLES Candidate Study, I construct a Left–Right score and an Anti-Elite score for each candidate from multiple survey items per dimension (Appendix Table K2 lists the core items and their exact wording). I treat the Left–Right dimension as a bundle of economic, cultural, and immigration policies, built from items on economic issues (e.g., whether the government should reduce income differences or stay out of the economy), cultural issues (e.g., whether same-sex marriage should be banned or whether women should be free to decide about abortion), and immigration (e.g., whether immigrants should adapt to German culture or whether immigration benefits the German economy). In line with the literature, I treat the Anti-Elite dimension as capturing antagonism toward elites, the claim that the people rather than professional politicians or technocrats should hold authority, and general distrust of democratic institutions (Mudde, 2004; Guriev and Papaioannou, 2022). It is measured with items such as whether democracy offers citizens enough opportunity to participate, whether there should be nationwide referendums, and whether citizens have lost trust in politicians. The core items recur across election years, and each wave also includes several year-specific items. I standardize each item using its pooled mean and standard deviation across the years in which it appears, and average each candidate's available standardized items within each dimension to obtain a single index.[19] This yields measures that are comparable across time and space.

Participation in the candidate survey is voluntary, and the response rate among single-member district candidates remains moderate (about 35% across the years). To improve data coverage, I supplement the GLES Candidate Study with a second source, *Wen Wählen?* (Whom to Vote For?). Because the two surveys use distinct question sets, I bridge them using an elastic net regression (Appendix K.2 provides details on the procedure). The 10-fold cross-validated correlation between the predicted and the GLES-based positions

[19] I prefer this method for transparency. The correlation with indices based on Item Response Theory exceeds 0.85 across all years and dimensions.

is about 0.87 for the Left–Right dimension and 0.73 for the Anti-Elite dimension.

Figure 6 plots the distribution of candidate positions by party on the Left–Right and Anti-Elite dimensions, pooled across the four elections (2013–2025). Within each party, positions span a broad range on both dimensions, but the ordering of parties is preserved: on the Left–Right dimension, the Left is furthest left and the AfD furthest right, with the Greens, SPD, FDP, and CDU/CSU in between; on the Anti-Elite dimension, the Left and the AfD are clearly more anti-elite, while the remaining parties sit on the pro-elite side. Appendix Figures K1 and K2 show these distributions and the party medians by election year.[20] These candidate positions are the reference points in the policy space relative to which I locate voters' ideal points in the analysis that follows.



(a) Left–Right (b) Anti-Elite

Figure 6. Distribution of candidate positions by party, pooled across elections

Notes. Kernel density estimates of the positions of single-member district candidates, by party, pooling the 2013, 2017, 2021 and 2025 elections and all federal states. Positions are estimated from the GLES Candidate Survey and the nonpartisan platform *Wen Wählen?*. Each density uses a Gaussian kernel with the bandwidth selected by Scott's rule, computed separately for each party and dimension.

Combining the two surveys yields comparable positions for about 60% of candidates across years. Coverage differs by party: it is lowest for the CDU (37%) and highest for the Greens (77%), with about 50% of AfD candidates covered. Importantly, within each party, respondents and non-respondents obtain comparable first-vote shares and have similar probabilities of winning. I impute missing positions with the median position of the same party's candidates in other districts of the same state, and use the national median only when no state-level information is available. For the estimation, I keep only districts with at most three imputed candidates. This leaves about two thirds of districts across years, in which roughly 30%

[20] External validation is difficult. One benchmark is the within-party evolution of positions in the Chapel Hill Expert Survey (CHES) (Bakker et al., 2015), with which my estimates are broadly consistent. The CDU's rightward and anti-elite shift between 2021 and 2025, and the moderation of the AfD's anti-elite positioning over the same period, are mirrored between the 2019 and 2024 CHES waves; the hardening of AfD positions between 2017 and 2021 is mirrored between the 2014 and 2019 waves. Since CHES waves do not coincide with federal elections, the correspondence is approximate.

of candidate positions are imputed (fewer than 2% from the national median).

An important concern is that two dimensions may be too few. While the framework in principle allows for more dimensions, the variation in candidate positions available here is too limited to yield informative bounds beyond two. Moreover, many other dimensions, such as nationalism, globalism, and immigration, are highly correlated with the two considered, which makes them hard to separate in preference estimation. Appendix Figure K3 documents the correlations between expert-estimated party positions across relevant policy dimensions in Europe; by contrast, the Left–Right and Anti-Elite dimensions are only weakly correlated, both in my data and across European parties.[21] Section 9 discusses these issues and how they could affect my results and their interpretation.

Other data sources. I also use administrative turnout rates by age, gender, state and election year from the Representative Electoral Statistics (Die Bundeswahlleiterin, 2025). Unlike survey measures, these rates are not self-reported: in a representative sample of polling districts, ballot papers are marked with the voter's age group and gender. Since the statistics do not record education, I apply the same turnout rate to both education groups within each age–gender group. This lets me construct, for each municipality, the socio-demographic composition of actual voters rather than of the eligible population. As a result, I can abstract from the turnout decision, and the estimated preferences are those of the population that voted.

5 Estimation & Identification

5.1 Estimation and model fit

The goal of estimation is to recover voters' preferences, namely the weight and the ideal-point distributions, and to use them to bound target parameters θ. For example, one such parameter could be

$$\theta_{Right} = \Pr\left[i \text{ has far-right views } | i \text{ votes for AfD in year } t\right].$$

To do so, I translate the empirical restrictions implied by the spatial voting model into a finite-dimensional optimization problem as discussed in Section 2. Specifically, I match the vote shares predicted by the model to the observed ones, measuring fit by the mean absolute error per municipality. I do this in two steps.[22]

First, I recover the weight of the Anti-Elite dimension relative to the Left–Right dimension by a one-dimensional grid search: for each value w_{AE} on the grid $\mathcal{W}$, I build a partition of ideal points with fixed size, where I remove sets of negligible measure. I then solve a linear program that best fits the election returns,

[21] I rely on CHES for these party positions because the candidate survey I use lacks items on some important dimensions: for example, it contains no questions on nationalism and few on the environment.

[22] In both steps, when I focus on the AfD only, the error is computed on AfD vote shares alone.

and select the weight $\hat{w}_{AE}$ that minimizes the resulting fit error:

$$\hat{w}_{AE} = \arg\min_{w_{AE} \in \mathcal{W}} \left[\min_{\phi \in \Phi_{w_{AE}}} \underbrace{\frac{1}{|\mathcal{M}|} \sum_{m} \sum_{j} |\text{ऽ}_j(m) - \sigma_j(m; \boldsymbol{\phi}, w_{AE})|}_{\text{mean absolute error } Q(w_{AE})} \right],$$

where $\Phi_{w_{AE}}$ collects the probability mass functions over the cells of the partition $\mathbb{X}_{w_{AE}}$, $\sigma_j(m; \boldsymbol{\phi}, w_{AE})$ is the predicted election return of j in municipality m when voters have weight w_{AE} and ideal points distributed according to $\boldsymbol{\phi}$, and $\text{ऽ}_j(m)$ is its observed counterpart. I estimate the weight separately for each region (East or West) and year, so that it is common to all voters within a region and year, whereas the distributions of ideal points remain group-specific.

Second, I fix the weight at $\hat{w}_{AE}$ and compute the best attainable fit $Q(\hat{w}_{AE})$, the smallest mean absolute error per municipality. The bounds on a target parameter θ are then the minimum and maximum of θ over all ideal-point masses whose mean absolute error exceeds $Q(\hat{w}_{AE})$ by no more than ϵ_{abs}.[23] Appendix C provides further details on the estimation, and Appendix D demonstrates the weight-estimation procedure with simulations.

Appendix C.4 discusses the fit of the model. Across municipalities, the correlation between predicted and observed AfD vote shares exceeds 0.9 for 2017–2025 and is 0.59 in 2013 (Appendix Figures C1 and C2). The model also reproduces how AfD support varies with local demographics: the correlations of predicted and observed municipal AfD shares with, for instance, the share of young or university-educated residents are close (Appendix Figure C3). This suggests that the model fits the data well and can be used for the preference estimation that follows.

5.2 Identification

I exploit variation in candidate positions across electoral districts within a state and year to estimate the preferences of AfD voters. Recall that x denotes a voter's ideal point, m the municipality (which determines the voter's district and hence the candidates on the ballot), s the state, and d the socio-demographic cell. The identifying assumption (Assumption 1a) is

$$f(x \mid m, d) = f(x \mid m', d) \quad \text{for municipalities } m, m' \text{ of state } s,$$

that is, within the same state and socio-demographic cell, the distribution of ideal points does not depend on which candidates the municipality's district fields. I first discuss why this assumption is plausible, and then how I assess the robustness of the results to it.

[23] Following Tebaldi, Torgovitsky and Yang (2023), I keep ϵ_{abs} small (10^{-5}) and verify that the bounds change little when it is raised by an order of magnitude.

The assumption rules out two possibilities: first, that voters sort into districts according to the candidates on the ballot; second, that parties choose candidates in response to local preferences that differ across districts within a state beyond what socio-demographic characteristics capture.

First, politically motivated relocation across districts within states is implausible in this setting. Under mixed-member proportional representation there is little scope for district-specific representation; the policy-relevant subnational tier is the state, not the district. Within-state mobility is also low and rarely political. In 2019, roughly 2% of the population moved between counties (Kreise) within the same state, mostly between cities and their surrounding areas, with comparable magnitudes throughout 1991–2024 (Deutscher Bundestag, Wissenschaftliche Dienste, 2022; Statistisches Bundesamt, 2025). Reported motives are family- and dwelling-related in over 70% of moves and job- or education-related in fewer than one in five; the SOEP records no political motive (Hofmann and Rüger, 2024).

Second, district candidacies are embedded in a state-level process of selecting candidates for the Bundestag, which limits local discretion over who runs. State party executives must co-sign district nominations and may object to a local nomination, forcing a repeat vote (§§ 20(2), 21(4) Bundeswahlgesetz), and nominations are shaped by state-list considerations, since seats are allocated through party-list votes. Candidate choice is further constrained by the relatively small pool of potential candidates (see, e.g., Höhne, 2023). Appendix K.1 documents nomination practices in more detail.

As anticipated, I also restrict the sample to municipalities with fewer than 10,000 inhabitants. Because the marginal electoral return to targeting a locality is low when its electorate is small, candidates have less incentive to adapt their platforms to these voters, so observed vote shares in small municipalities are less likely to reflect strategic positioning than those in larger, more contested urban centers.

I present two pieces of observational evidence that strengthen the assumption of little candidate response to local preferences. First, across the municipalities used in the estimation, the candidate positions of a given party are largely uncorrelated with local socio-economic characteristics within a state. The same is not true at the district level: positions correlate with a district's socio-economic characteristics within party and state, with the expected signs. For example, candidates are more left-leaning in districts with higher unemployment or larger shares of benefit recipients. At both levels, candidates' positions correlate with some socio-demographic characteristics, such as age shares and the female share, which is not a threat to identification since the model conditions on them. Appendix Figures I5 and I6 report the balance of socio-economic and socio-demographic characteristics at the municipal and district levels. Conditional on the socio-demographic composition that the model accounts for (age, gender, and education), however, the remaining socio-economic characteristics show no systematic association with candidates' positions on either dimension at the municipal level, and most associations are muted at the district level (Appendix Figures I7 and I8). Second, Appendix Table I3 shows that candidates' positions are uncorrelated with individuals' self-reported views in surveys.

Robustness. I assess the robustness of the identification assumption in four ways, each targeting a different reason why candidate positions might be correlated with local preferences.

First, I restrict the analysis to municipalities with fewer than 5,000 inhabitants. Parties are likely to target more populous, and hence more electorally relevant, constituencies; small municipalities therefore face candidates selected with larger constituencies in mind, so their candidates' positions are less likely to reflect local preferences.

Second, I exclude districts in which the AfD candidate was competitive. Candidates' incentives to cater to local voters increase with their chances of winning, so positions in non-competitive districts are less likely to be tailored to the constituency. Specifically, I drop every district in which the AfD candidate trailed the winner by no more than 5 (10) percentage points in the East (West). The thresholds differ because the party is far more popular in the East. In the East, the exercise covers only 2017 and 2021: in 2013 no AfD candidate was competitive, and in 2025 almost all were, as the AfD won nearly every district. In the West, it applies only to 2025, since in earlier years no AfD candidate was competitive and the restricted sample coincides with the benchmark.

Third, I relax the assumption of preference similarity across districts within a state, following the logic of imperfect instruments in Tebaldi, Torgovitsky and Yang (2023). I classify a district as *conservative* if its AfD candidate lies above the within-state median AfD position on either the Left–Right or the Anti-Elite dimension, and as *moderate* otherwise. I then allow the mass of voters of a given socio-demographic group in each partition cell to differ between conservative and moderate districts of the same state by at most a factor of $1+\kappa$, with $\kappa = 2$, so that masses may differ by up to a factor of three in either direction. If candidates were placed in response to local preferences, the conservative and moderate districts would differ systematically, and this exercise accommodates such differences.

Finally, I restrict the analysis to municipalities on the boundaries between electoral districts within a state, and replace the assumption that preferences are identical across the whole state with a local one. I group municipalities by border segment, defined as the boundary between two adjacent districts, and require the preference distribution of a given socio-demographic group to be identical only among municipalities on either side of the same segment (the median number of municipalities per segment is seven). Across segments within a state, preference distributions may deviate from the state mean by up to 20% in either direction.[24] Neighboring municipalities on either side of a border are similar but face different candidates, so district assignment acts as an instrument: it shifts the candidates a voter faces (*relevance*) but, within a segment, is assumed independent of preferences (*exogeneity*). The identifying variation therefore comes from comparisons in which preferences are most plausibly similar.

The analysis that follows leverages the variation in estimated preferences not only within states, but

[24] Ideally, the bounds would be profiled over this tolerance. I have not yet done so, as this is the most computationally intensive exercise in the paper. At a tolerance of 100%, the results are uninformative.

also across states. This is possible because candidate positions are constructed from survey items harmonized and standardized across elections, so the estimated ideal-point distributions lie in a common policy space and can be compared and aggregated across states, and compared across elections. Target parameters can therefore be defined relative to national benchmarks, such as the national median AfD position, and traced over 2013–2025 (Section 6 onwards). None of these comparisons requires candidate positions to be exogenous across states: ideal-point distributions remain state-specific.

6 Results: Voters' Preferences

This section uses revealed-preference evidence to shed new light on the rise of populism. I first estimate voters' preferences: the weight on the Anti-Elite dimension relative to the Left–Right dimension, w_{AE}, and the distribution of ideal points of each socio-demographic group d in each state s, $f(x \mid s, d)$. I then use these estimates for counterfactual exercises.

6.1 Relative importance of dimensions

The weights determine the shape of the partition of the policy space used to evaluate the target parameters below. Only relative weights are identified, so I normalize $w_{LR} = 1$ and search for the best-fitting w_{AE}. Since precise estimation of the weights requires rich data (Appendix D), I impose a common w_{AE} on all groups within a given year and region (East or West), while allowing each socio-demographic group d in each state s to have its own ideal-point distribution. I search over a grid of relative weights $w_{AE} \in \{0.2, 0.3, \ldots, 0.9, 1\}$ and their inverses. The best-fitting weight is the one that yields the smallest model-fit error.[25]

Figure 7 presents the results for the weights. Diamonds denote the point estimates corresponding to the best-fitting weight. The results suggest that the Anti-Elite dimension carried more weight in eastern than in western Germany throughout the period: in the West, the Left–Right dimension was consistently the more important of the two, whereas in the East, the Anti-Elite dimension carried substantial weight in all years.

These results are consistent with the different roles the two dimensions play in each part of the country. A broad literature documents Eastern Germans' widespread perception of being neglected by federal elites, a perception grounded in measurable underrepresentation: as of 2024, Eastern Germans held only 12.1% of positional elite roles across twelve societal sectors, well below their population share of roughly 20% (Bundesministerium für Wirtschaft und Energie, 2025). In the West, where no comparable region-wide sense of underrepresentation exists, the Anti-Elite dimension is much more muted.

[25] I do not extend the grid further because a weight is identified only insofar as it changes the partition of the policy space, which for each pair of candidates depends on the ratio of their position differences, $|y_j^{LR} - y_k^{LR}|/|y_j^{AE} - y_k^{AE}|$. More than 80% of these ratios for pairs involving AfD candidates lie within $[0.2, 5]$, so more extreme weights would be identified from few pairs and sensitive to noise. As a check, I extend the grid by one step whenever the best-fitting weight lies on its boundary. In the West, adding 0.1 confirms 0.2 in 2021 and 2025, while in 2017 the fit improves slightly at 0.1 (Appendix Figure J3a). This does not materially affect the bounds on the target parameters as I show below, I use 0.2 in all three years.

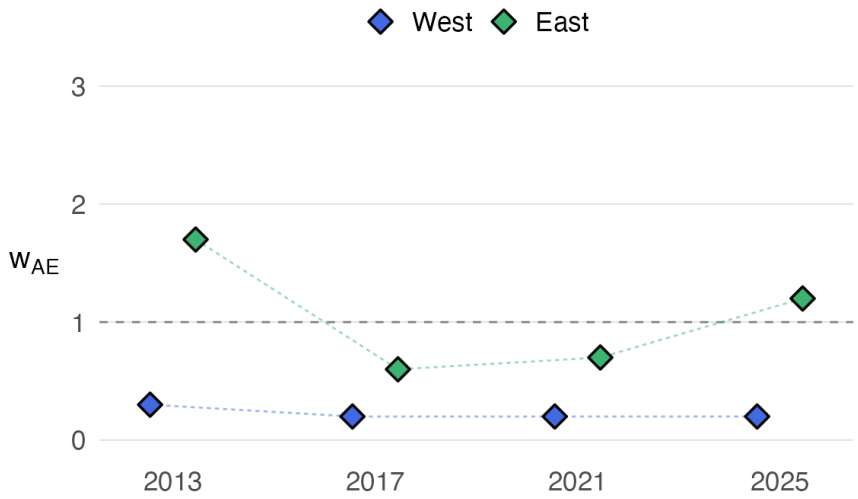


Figure 7. Weights for the Anti-Elite dimension w_{AE} relative to Left–Right $w_{LR} = 1$ across years and regions (East/West)

Notes. $w_{LR} = 1$ is fixed throughout; w_{AE} should be interpreted as the relative weight placed on the Anti-Elite dimension. Separate exercises are run for East and West Germany. I allow ideal-point distributions to vary across federal states – that is, each socio-demographic group d in each state s has its own ideal-point distribution $f(x \mid s, d)$, while imposing the same weight w_{AE} for all groups within a given year and region. To run this exercise, the grid of possible weights is constructed. The grid is symmetric around 1 and implies relative weights across the two dimensions of $(0.2, 0.3, \ldots, 0.9, 1)$ and the corresponding inverse ratios. See Appendix Figure J3 for full results over the grid. Diamonds denote the point estimates corresponding to the best-fitting weight for each year and region.

The weights are set-identified, since a range of weights implies the same set of available profiles and hence the same feasibility problem. Because of noise, however, the optimization yields a single best-fitting estimate (Appendix D illustrates this with simulations). To reflect set identification, I also report the set of weights whose fit error lies within 1% of the minimum (Appendix Figure J1), and confidence sets calibrated from subsample re-estimations, which replace the fixed tolerance with one derived from the sampling variability of the fit (Appendix Figure J2). Both are narrow in all years except 2013.

The model's ability to discriminate between weights also depends on how candidate positions co-locate along the two dimensions. The correlation and dispersion of candidates' positions determine how many cells a given weight induces in the Minimal Relevant Partition, and hence how much the fit could improve mechanically with the number of cells. Two features of the design address this. First, the estimation holds the number of cells fixed across weights, so differences in fit reflect the composition of the cells, not their number (Appendix C.1). Second, these supply-side patterns do not track the estimated differences in weights: they remain broadly similar between East and West within each year (Appendix Figure I9 and Table I5). The estimated differences in weights are therefore unlikely to be driven by differences in the supply of candidate positions.

In the following analysis, I use the values for the estimated weights to define the partition of the policy space across the two regions and elections unless stated otherwise.

6.2 Distribution of ideal points

6.2.1 Voters' ideal points in fixed areas of the policy space

Where are voters' ideal points located in the policy space? I examine this question for AfD voters in 2013–2025. I partition each policy dimension into intervals of width 0.5 and, for each interval, define the target parameter as the share of baseline AfD voters whose ideal points lie in it.

For a given election year, let $\mathcal{J}$ denote the set of the six parties, let $\bar{Y}_j$ denote the coordinate-wise median of the positions of party j's candidates nationwide, and let w be the estimated weight for the given region and year. The AfD's baseline attraction region is the set of ideal points $\mathcal{A}_{\text{AfD}} := \{x \in \mathcal{P} : dist_w(x, \bar{Y}_{\text{AfD}}) \leq dist_w(x, \bar{Y}_j) \text{ for all } j \neq \text{AfD}\}$. Let $f(x \mid s, d)$ denote the density of ideal points among voters who turned out in state s and socio-demographic group d, and let $F_{S,D}$ denote the joint distribution of states and socio-demographic groups among those voters, constructed using the turnout-adjusted socio-demographic shares. The share of baseline AfD voters whose ideal points lie in $[0.5, 1]$ on the Anti-Elite dimension can then be written as the target parameter

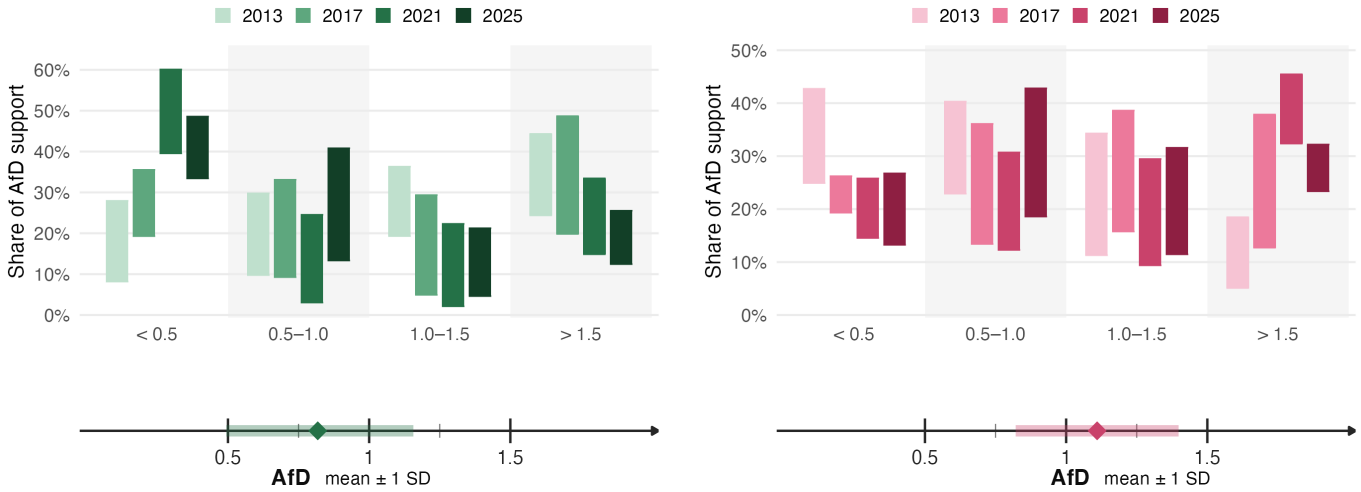
$$\theta_{\text{AE},[0.5,1]} := \frac{\displaystyle\iint_{\mathcal{A}_{\text{AE},[0.5,1]}} f(x \mid s, d)\, dx\, dF_{S,D}(s, d)}{\displaystyle\iint_{\mathcal{A}_{\text{AfD}}} f(x \mid s, d)\, dx\, dF_{S,D}(s, d)}, \qquad \text{where} \quad \mathcal{A}_{\text{AE},[0.5,1]} := \mathcal{A}_{\text{AfD}} \cap \{x \in \mathcal{P} : 0.5 \leq x_{\text{AE}} \leq 1\}.$$

Because $f(x \mid s, d)$ is only partially identified from the municipal vote shares and demographic composition, I report bounds on $\theta_{\text{AE, [0.5, 1]}}$ and other conceptually similar target parameters.[26] Figure 8 plots the results separately for the Anti-Elite and Left–Right dimensions for different dimension bins. Appendix Figure J4 reports sampling uncertainty behind the bins: I resample municipalities, re-estimate the bounds, and construct 95% confidence intervals for the target parameters.

The results suggest a recomposition of AfD support between 2013 and 2025. In 2013, AfD support came predominantly from voters with strongly anti-elite views rather than from the radical right: up to 45% of AfD supporters had an ideal point above 1.5 on the Anti-Elite dimension, whereas at least a quarter had an ideal point below 0.5 on the Left–Right dimension, and at most 19% lay beyond 1.5. Over the following decade this pattern reversed. The share of strongly anti-elite supporters fell to about a fifth by 2025, while the share with radical-right positions rose sharply. Notably, the party's own position moved far less than its electorate: the national AfD median remained between 0.9 and 1.2 on the Left–Right dimension and between 0.7 and 1.0 on the Anti-Elite dimension throughout. This placed the party well short of its most anti-elite supporters in 2013, and to the left of much of its right-wing support from 2021 onward.

Two mechanisms could lie behind this recomposition. First, the AfD mobilized many previous non-voters over this period, which mechanically changes the composition of its electorate: the estimates partly

[26] Because the partition is coarsened to a fixed number of cells, the bounds reported below approximate the *sharp* identified set; re-estimating bounds on the full partition yields similar results (see Appendix C.1).



(a) Anti-Elite dimension

(b) Left–Right dimension

Figure 8. Distribution of AfD supporters' ideal points along each dimension, 2013–2025

Notes: Each panel plots bounds on the share of AfD voters whose ideal points lie in each interval of one policy dimension (the other dimension is left unrestricted): panel (a) the Anti-Elite dimension, panel (b) the Left–Right dimension. Within each interval, bars run from 2013 (lightest) to 2025 (darkest). The line below each panel places the intervals on the dimension; the diamond marks the mean position of AfD district candidates pooled over 2013–2025 and the band one standard deviation around it.

capture different people across years.[27] Second, the meaning of the Left–Right dimension, which combines economic, cultural, and immigration positions, may have changed over time. In Section 9, I decompose it into economic and cultural subdimensions and find that voters increasingly weight the cultural one, while their ideal points on both remain relatively stable, and that AfD voters are more conservative on the cultural subdimension than on the economic one, consistent with Danieli et al. (2026). Together, these patterns suggest that the changing meaning of the Left–Right dimension could explain part of the recomposition.

Much of the political economy literature estimates voters' preferences from stated vote choices and issue positions (Le Pennec and Pons, 2019; Gethin, Martínez-Toledano and Piketty, 2022), including for populist support (Guiso et al., 2017; Algan et al., 2017; Danieli et al., 2026; Bartels, 2023). Appendix E.1 compares such survey-based estimates with the revealed-preference ones. The patterns are broadly similar: AfD voters are further left on the Left–Right dimension in earlier years and further right in later years, and both measures show a decline in anti-elite attitudes over 2017–2025, although the survey measure is less informative on this dimension. The revealed-preference estimates place more voters at extreme positions, consistent with the underreporting of extreme views in surveys. Exact shares are hard to compare, given differences in measurement and possible misreporting of party choice in surveys.

These patterns provide first evidence on how the distribution of AfD voters' ideal points has evolved, pointing to a changing composition of the party's electorate. They do not, however, speak directly to why voters choose the AfD or to how establishment parties might win them back, which the next exercises address.

[27] This was most pronounced in 2017 and 2025, when the AfD drew about 1.5 million or more votes from people who had not voted in the previous election (Infratest dimap, 2017, 2025).

6.2.2 Voters' ideal points and their motivations for AfD support

The AfD combines a radical-right stance with an anti-elite appeal, so a ballot cast for the party does not reveal which of the two motivated it. Recovering that motivation matters because it speaks to the policy demands the vote expresses and could help policymakers craft more appropriate responses. I address this ambiguity with a counterfactual: suppose the AfD split into two parties, one retaining its radical-right position on the Left–Right dimension but moving toward the elite, and the other retaining its anti-elite appeal but moderating its right-wing position. What share of baseline AfD voters would be closer to each?

Using the nationwide median AfD position defined above, the two shifts of size δ partition the baseline AfD attraction region $\mathcal{A}_{\text{AfD}}$ into an anti-elite region $\mathcal{A}_{\text{AE}}$ and a Right region $\mathcal{A}_{\text{Right}}$, shown in green and pink in Figure 9a. Voters in each region are closer to the hypothetical successor that retains, respectively, the AfD's anti-elite or Left–Right position.[28]

The target parameter is the share of baseline AfD voters in the anti-elite region,

$$\theta_{\text{AE}} := \frac{\iint_{\mathcal{A}_{\text{AE}}} f(x \mid s, d)\, dx\, dF_{S,D}(s,d)}{\iint_{\mathcal{A}_{\text{AfD}}} f(x \mid s, d)\, dx\, dF_{S,D}(s,d)}.$$

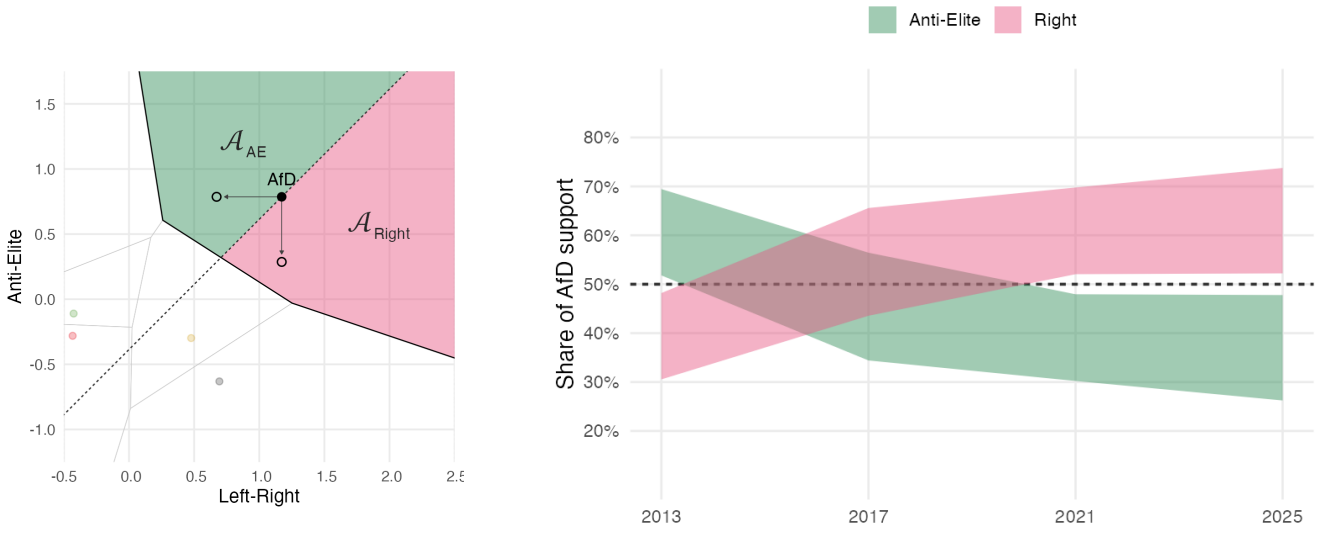
I report bounds on θ_{AE} and on $\theta_{\text{Right}} = 1 - \theta_{\text{AE}}$, the share of baseline AfD voters in the Right region.

Figure 9b provides the bounds for the target parameters over years. The results indicate that AfD support has changed over time. In earlier years, AfD support was driven more by anti-elite motivations; this pattern flipped over time, and in 2021 up to 70% of AfD support comes from individuals with stronger ideological Right motivations. Appendix Figure J6 quantifies sampling uncertainty by re-estimating the bounds on bootstrapped samples of municipalities.

These aggregate results mask substantial regional and subgroup heterogeneity. Appendix Figure J7 reports, separately for eastern and western Germany, the share of AfD voters whose ideal points lie in the Right region of Figure 9a. The pattern for western Germany resembles the national pattern presented above, whereas voters in eastern Germany are more strongly driven by anti-elite attitudes until 2025.

Appendix Figure J8 highlights subgroup heterogeneity among Western voters over time: it plots the same share by education, age, and gender. Education shows a stable pattern: a larger share of voters without higher education lies in the Right area in every election, and the bands separate after 2021. Moreover, AfD voters with higher education are never identified as predominantly right-wing: their Right share is bounded below one half in 2013 and 2017, and the band only barely reaches it in 2025, whereas for voters without higher education it lies entirely above one half from 2021 on. For age the pattern reverses in 2025: the share of young AfD voters located in the Right group is larger than the corresponding share of older AfD voters.

[28] The benchmark uses $\delta = 1$. The results are quantitatively similar for other positive shifts of comparable size.



(a) Illustration

(b) Estimated shares of voters preferring the hypothetical Anti-Elite vs. Right party among AfD voters ($\theta_{AE}, \theta_{Right}$)

Figure 9. Illustration of a target parameter (a) and results (b)

Notes: Panel (a) shows party positions at their national medians; hollow circles mark the counterfactual AfD *shifts*. The shaded areas contain AfD voters motivated more by the party's anti-elite stance (green, $\mathcal{A}_{\text{AE}}$) or by its right-wing stance (pink, $\mathcal{A}_{\text{Right}}$). The illustration uses $w_{AE} = 1$; the estimation uses the weights estimated in Section 6.1. Panel (b) plots bounds on the shares of AfD voters whose ideal points lie in the green and pink regions, by year, relative to the year-specific median AfD position. Appendix Figure J5 shows that treating the weights as set-identified (all values within 1% of the best fit) leaves the results similar, and those for 2017–2025 essentially unchanged.

This is in line with the recent evidence on the radicalization of young voters' attitudes and party choices (van der Brug, Hobolt and Popa, 2026). By gender the bands overlap throughout, but from 2021 on the men's band lies entirely at or above one half whereas the women's straddles it; the weak identification by gender is expected, since the gender composition of municipalities varies the least.

Appendix E.2 presents the same estimates using survey data. The patterns are somewhat similar to the revealed-preference ones, although the survey attributes a larger role to Anti-Elite motivations in 2017–2021, again consistent with underreporting of right-wing views. The comparison concerns ideal points only: the weights are taken from the revealed-preference estimates, since they cannot be recovered from survey data.

Robustness. I now report the four robustness exercises anticipated in Section 5.2, which address the selection-on-observables assumption about candidate positions. The first restricts the sample to municipalities with fewer than 5,000 inhabitants (Appendix Figure J9). The second restricts it to non-contender districts, separately for Eastern and Western Germany (Appendix Figure J10). The third allows preferences to differ between districts with more moderate and more conservative AfD candidates (Appendix Figure J11). The fourth restricts it to municipalities on district borders and requires preferences to be identical only on either side of the same border, rather than across the whole state (Appendix Figure J12). The main qualitative results hold across all four specifications. The bounds are noticeably wider in the last two, which impose weaker equality restrictions on preference distributions within states.

The analysis so far has described the evolution of AfD voters' preferences. The next sections use these preferences as primitives for counterfactuals in which the AfD, its center-right rival, the Christian Democratic Union (CDU), or both parties reposition, yielding new insights into the competition between populist and mainstream parties.

7 Repositioning Counterfactuals

7.1 How many of its voters would the AfD lose by moderating?

The AfD does not exist in isolation but competes with other parties for voters, so assessing whether its success rests on right-wing ideology or on anti-elite sentiment requires knowing how voters respond when party positions change. Combining the estimated ideal-point distributions and relative weights with the median positions of all parties, I measure the AfD's own-position *elasticity*: how many of its baseline supporters it would lose by adopting a more pro-elite or a more left-wing position. Specifically, I move the AfD by 0.2 or 0.5 units along each dimension, relative to its national median in each year.[29] Appendix Figure J14 illustrates these shifts.

I interpret the step sizes in two ways. First, I compare them with the national spread of each party's district-specific candidate positions within a year: over 2013–2025, the standard deviation is about 0.3–0.5 on both dimensions for both the AfD and the CDU/CSU. A step of 0.2 thus corresponds to roughly half a standard deviation of the party's own candidates, and a step of 0.5 to about one; I treat 0.2 as an ambitious but realistic shift.

Second, I translate the step sizes into changes in candidates' survey answers. Each item shifts the dimension score by a different amount per scale point. On the Anti-Elite dimension, the most influential item, *Satisfaction with the way democracy works* (a 1–4 scale), moves the score by almost 0.2 per point, and a five-point item such as *Legislation reflects citizens' interests* by about 0.14. On the Left–Right dimension, potent five-point items such as *Immigration is good for the economy* or *Marriage only for heterosexual couples* move the score by about 0.10 per point. A 0.2 step on the Left–Right dimension thus amounts to shifting the answers to two potent items by one point, or to three or four less potent ones, whereas a 0.5 step requires a substantial repositioning across many items at once. Appendix Table I6 reports item-level sensitivities for 2017 as an illustration.

The target parameter is the share of baseline AfD voters who would abandon the party if it shifted to **AfD'**.[30] Figure 10 plots the results separately for different step sizes and directions of the shift.

[29] Positions are standardized by construction, so a given step is comparable across dimensions and years.

[30] Formally, for a shift along the Anti-Elite dimension,

$$\theta_{\text{Loss, A-E}} := \frac{\iint_{\mathcal{A}_{\text{Loss,AE}}} f(x \mid s,d)\,dx\,dF_{S,D}(s,d)}{\iint_{\mathcal{A}_{\text{AfD}}} f(x \mid s,d)\,dx\,dF_{S,D}(s,d)}, \qquad \mathcal{A}_{\text{Loss,AE}} = \mathcal{A}_{\text{AfD}} \setminus \mathcal{A}_{\text{AfD}'},$$

and analogously for shifts along the Left–Right dimension.

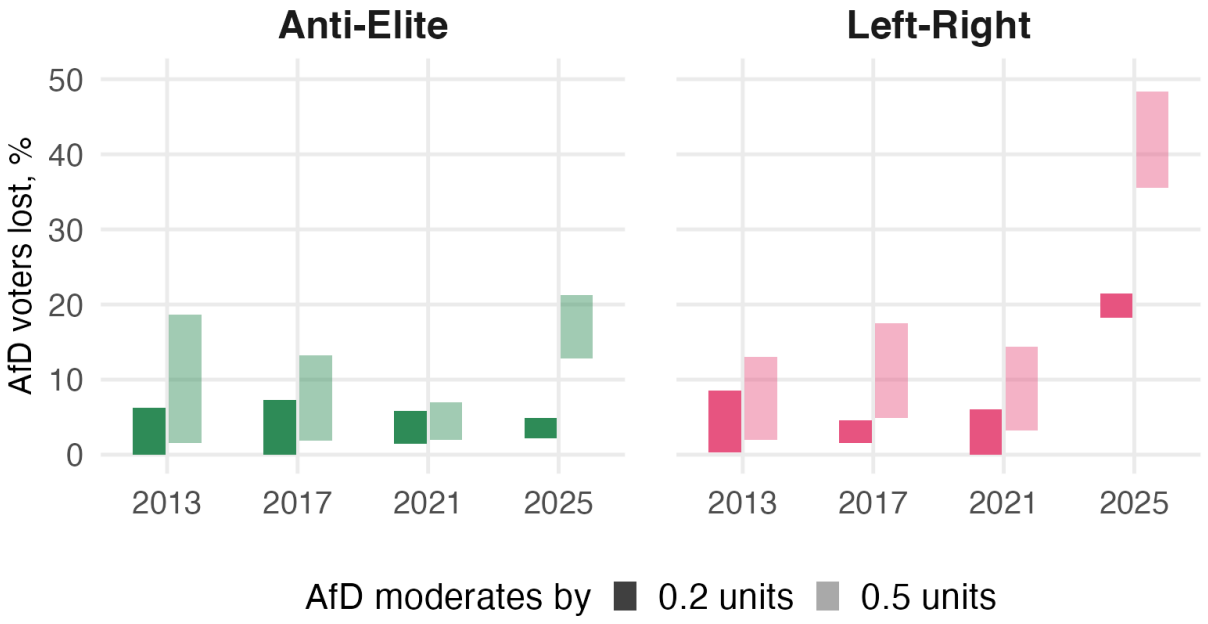


Figure 10. Estimated shares of AfD voters lost under counterfactual AfD shifts

Notes. Each bar is the interval for the share of *baseline* AfD voters who would abandon the party if it moderated by one step along the indicated dimension, holding all other parties at their median observed positions in a given year. In the left panel the AfD moves toward the elite along the Anti-Elite dimension; in the right panel the AfD moderates along the Left–Right dimension. Within each year, shading denotes the size of the moderation step (0.2/0.5 in standardized policy-space units; see the discussion just above for its substantive interpretation).

Before 2025, the AfD's electorate was relatively insensitive to small shifts: a 0.2 step on either dimension would have cost the party at most about 9% of its voters in any year, and a 0.5 step at most about 19%. In 2025, the picture changes along the Left–Right dimension: a 0.2 step towards the center would cost the AfD about 18–22% of its voters, and a 0.5 step 35–48%. Along the Anti-Elite dimension, by contrast, losses remain moderate even in 2025, at most about 5% for a 0.2 step and at most 21% for a 0.5 step. Competition between the AfD and other parties thus became much tighter by 2025, and primarily along the Left–Right dimension: even a small moderation of the party's right-wing positions would now drive many of its voters to other parties. Before 2025, by contrast, the party's support existed in relative isolation. Appendix Figure J15 reports the corresponding 95% confidence intervals.

If competition with the establishment parties has indeed intensified, those parties should be able to capture a large share of AfD voters by shifting their positions only slightly. I now turn to this analysis.

7.2 How many AfD voters could the CDU win?

German establishment parties, which once commanded a large majority of the electorate, are now losing ground to the AfD, which leads the polls nationwide as of October 2026. This development is by no means unique to Germany: across Western democracies, populist and radical parties have risen sharply over the past two decades, in many cases entering national legislatures and governments, as in Italy, France, the

Netherlands, Spain, and Sweden. Beyond Europe, the rise of Donald Trump in the United States and of Javier Milei in Argentina is widely regarded as part of the same populist wave.

How should establishment parties compete with populists? This section examines how effective it would be for them to adopt positions closer to the AfD's in the two-dimensional space defined by the Left–Right and Anti-Elite dimensions. In this and the next section, I focus on the Christian Democratic Union (CDU), the main establishment party on the right of the Left–Right spectrum.

Mirroring the previous section, I measure the AfD's cross-position *elasticity*: how many AfD voters the CDU would capture by moving a given step towards the AfD along the Left–Right or Anti-Elite dimension (Appendix Figure J16 illustrates these scenarios). Figure 11 plots the results by direction of the shift.

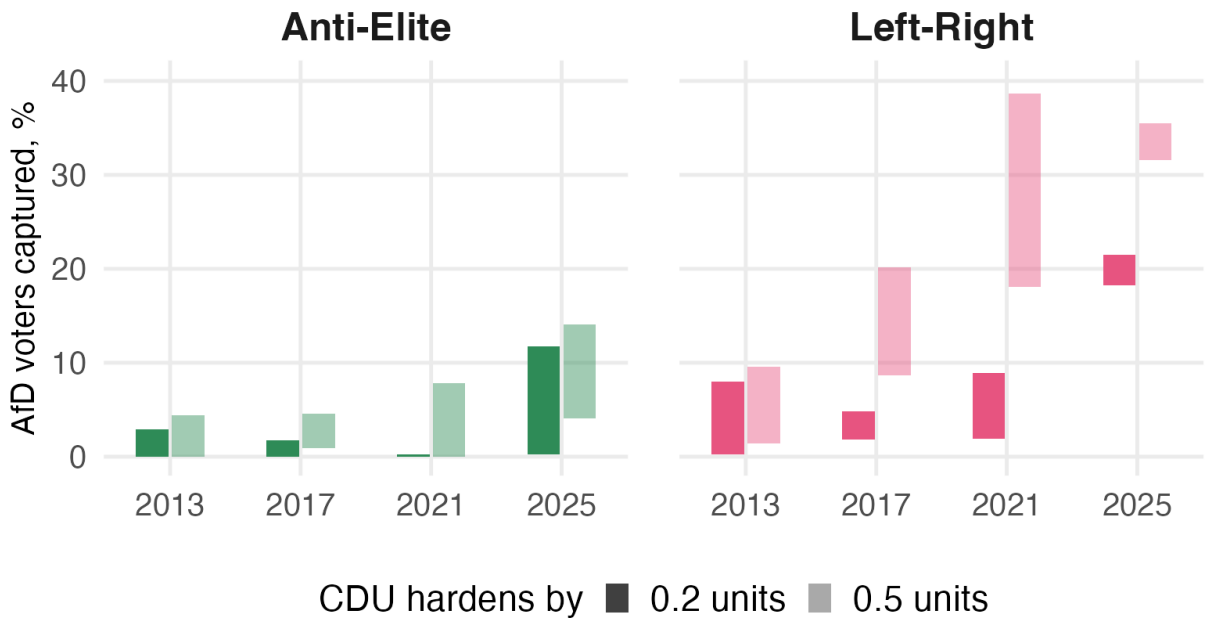


Figure 11. Estimated shares of AfD voters captured under counterfactual CDU shifts

Notes. Each bar is the identified interval for the share of *baseline* AfD voters who would abandon the party if the CDU moved by one step along the indicated dimension, holding all other parties at their observed positions. In the left panel the CDU moves towards more anti-elite positions; in the right panel the CDU moves towards more right-wing positions. Within each year, shading denotes the size of the step (0.2/0.5 in standardized policy-space units; see the discussion just above for its substantive interpretation).

Over 2013–2021, a 0.2 step to the right would allow the CDU to capture up to 9% of AfD support, whereas a 0.2 step towards more anti-elite positions would capture less than 3%. In 2025, the same rightward step would capture about 20% of the AfD electorate, more than any anti-elite step would capture. Larger rightward steps have also become increasingly effective over time. Although this exercise does not solve the CDU's vote-share maximization problem, a fifth of AfD voters is a large share by 2025, and the CDU is unlikely to lose as many voters on its left flank; the next section examines this trade-off in more detail. Appendix Figure J17 reports the corresponding confidence intervals.

Put into perspective, this result implies that had the CDU slightly hardened its position on immigration in the 2025 election, it could have captured about a fifth of the AfD electorate. Concretely, this would mean moving from "Rather agree" to "Neither agree nor disagree" on *"Immigrants are good for the German economy"*, from "Rather agree" to "Strongly agree" on *"Immigrants should adapt to German culture"*, and from 8 to 9 on *"Own position on immigration"*.[31]

These results are in line with the experimental findings of Guenther and Nunnari (2025): perceiving the CDU as one unit more liberal on immigration (on a 0–10 scale) raises the intention to vote for the AfD by 5 percentage points, about a quarter of the AfD vote. On the same scale, my estimates imply that a 2-unit hardening would capture about a fifth of the AfD electorate. They are also consistent with observed voter flows: Infratest dimap estimates that the AfD gained just over one million voters from the CDU on balance between 2021 and 2025 (Schläger, Engels and Loew, 2025). Their previous support for an establishment party suggests that many may not share the AfD's anti-elite outlook, and that a slightly harder Left–Right position could more plausibly win them back.

Overall, at least two factors determine the outcome of these exercises: the relative positions of the parties and the distribution of voters' ideal points. In earlier years the AfD's electorate was more anti-elite and less radical-right than later (Section 6.2.1), so a rightward CDU move could capture few of its voters, and the anti-elite move that might have done so was implausible (or still small) for an establishment party. Between 2017 and 2021, the gap between the CDU and the AfD, if anything, widened (Appendix Figure K2), yet the CDU's reach grew, as the shift in the AfD's electorate brought many of its voters within range of a large step. By 2025, the CDU had moved right under Friedrich Merz, who comes from its conservative wing, narrowing the gap and bringing many AfD voters within reach of even a small step. With both parties competing intensely for the same voters, I turn to the strategic interaction between them in 2025.

8 Equilibrium: Spatial Competition between the Establishment and the Populists

One of the puzzles of the rise of populism in Europe is why establishment parties do not adjust their platforms so as to hinder the entry and growth of populist rivals. Several answers are plausible: the establishment may anticipate that the populists would move in turn, it may fear losses on its other flank, or repositioning may be costly. The framework of this paper can separate the first two, since both are governed by the distribution of voters' preferences.

To do so, I study a spatial competition game between the CDU and the AfD, taking the estimated ideal-point distributions and weights as primitives.[32] I focus on the 2025 Bundestag election, since before

[31] All illustrative answers are given relative to the median answers of CDU candidates in the 2025 survey. See Appendix Table I7 for the contribution of a one-point shift in each survey item.

[32] I use the weights estimated for AfD voters throughout, since the game centers on competition for these voters and the model imposes a common weight.

2025 a plausible rightward step by the CDU would have captured only a small share of the AfD's electorate (Figure 11). A rightward shift would indeed pay off for the CDU if the AfD stayed put; once the AfD can respond, however, the CDU's gains disappear.

The game. All parties other than the CDU and the AfD are held at their observed 2025 positions; the set of players is $N = \{\text{CDU}, \text{AfD}\}$. Each player $n \in N$ may shift its position Y_n by at most 0.2 standardized units on each dimension, in steps of 0.1. On Left–Right both parties may move in either direction (and I separately consider a game of Left–Right moves only below). On Anti-Elite I only admit moves that are plausible for each party: the CDU may become less anti-elite but not more, the AfD more anti-elite but not less. The action sets are therefore 15 moves each, including the status quo $(0, 0)$:

$$A_{\text{CDU}} = \{-0.2, \ldots, 0.2\} \times \{-0.2, -0.1, 0\}, \qquad A_{\text{AfD}} = \{-0.2, \ldots, 0.2\} \times \{0, 0.1, 0.2\}.$$

The payoff of a profile $a = (a_{\text{CDU}}, a_{\text{AfD}}) \in A_{\text{CDU}} \times A_{\text{AfD}}$ is a *target parameter* – the change in n's share of the electorate relative to the status quo,

$$u_n(a) \;=\; \iint_{\mathcal{A}_n(a)} f(x \mid s, d)\, dx\, dF_{S,D}(s, d) \;-\; \iint_{\mathcal{A}_n(0)} f(x \mid s, d)\, dx\, dF_{S,D}(s, d),$$

where $\mathcal{A}_n(a)$ is n's attraction region under the counterfactual positions ($Y_{\text{CDU}}+a_{\text{CDU}}$, $Y_{\text{AfD}}+a_{\text{AfD}}$), and $\mathcal{A}_n(0)$ is the same area under the status quo. The payoff is computed over the 2025 electorate with participation held fixed, and does not model the turnout margin.

Parties are assumed to know the distribution of voters' ideal points, which the econometrician only knows up to the identified set. I therefore draw $K = 100$ distributions from the identified set and solve the game for each.[33,34] For draw k the payoffs $u_n^{(k)}(a)$ define a normal-form game $G^{(k)}$, and I report how outcomes are distributed across draws.

I study the sequential version of the game in which the CDU (the leader) moves first and the AfD (the follower) observes a_{CDU} before choosing a_{AfD}. The follower's best-response correspondence:

$$B_{\text{AfD}}^{(k)}(a_{\text{CDU}}) \;=\; \arg \max_{a_{\text{AfD}} \in A_{\text{AfD}}} u_{\text{AfD}}^{(k)}(a_{\text{CDU}}, a_{\text{AfD}}),$$

[33] Each draw is a random vertex of the identified set: I minimize a linear objective with i.i.d. standard-normal coefficients over the set of voter distributions consistent with the data, which returns an extreme point of the set. Every admissible distribution is a convex combination of vertices, and its payoffs are the same combination of vertex payoffs, so they lie within the range spanned by the vertices, which the draws sample.

[34] Solving the game draw by draw is a simplification, but still keeps the approach free of any functional-form assumption on voters' preferences: each draw is a distribution the data admit, not one I impose. Best replies and the Stackelberg outcome are similar across draws (I report the distributions below), so the equilibrium conclusions do not depend on which admissible distribution is the true one.

and the leader solves

$$a^{*(k)}_{\text{CDU}} \in \arg \max_{a_{\text{CDU}} \in A_{\text{CDU}}} \min_{a_{\text{AfD}} \in B^{(k)}_{\text{AfD}}(a_{\text{CDU}})} u^{(k)}_{\text{CDU}}(a_{\text{CDU}}, a_{\text{AfD}}),$$

the inner minimum resolving ties in the follower's response pessimistically for the leader. A profile $a^{*(k)} = (a^{*(k)}_{\text{CDU}}, a^{*(k)}_{\text{AfD}})$ with $a^{*(k)}_{\text{CDU}}$ solving the leader's problem and $a^{*(k)}_{\text{AfD}} \in B^{(k)}_{\text{AfD}}(a^{*(k)}_{\text{CDU}})$ is a **Stackelberg equilibrium** of $G^{(k)}$. The sequential formulation matches the question being asked, *whether the establishment party should move towards its populist rival*, and it sidesteps a well-known difficulty of simultaneous-move spatial competition: in more than one dimension a pure-strategy Nash equilibrium exists only under restrictive conditions on preferences (Osborne, 2025), and in the present game in practice it does not exist.

When the AfD stands still. I first ask whether moving right pays for the CDU when the AfD does not respond. Figure 12a reports the CDU's median net change in vote share (gray bars) for each Left–Right move, holding its Anti-Elite position and the AfD's position fixed. The answer is yes, and increasingly so: a 0.1 step gains about 0.3 percentage points of the turned-out electorate and a 0.2 step about 3 points, while steps to the left lose votes. The AfD's changes (light blue bars) mirror the CDU's: the AfD loses when the CDU moves closer and gains when it moves away. A rightward step thus gains the CDU more from the AfD than it loses to other parties on its left flank.

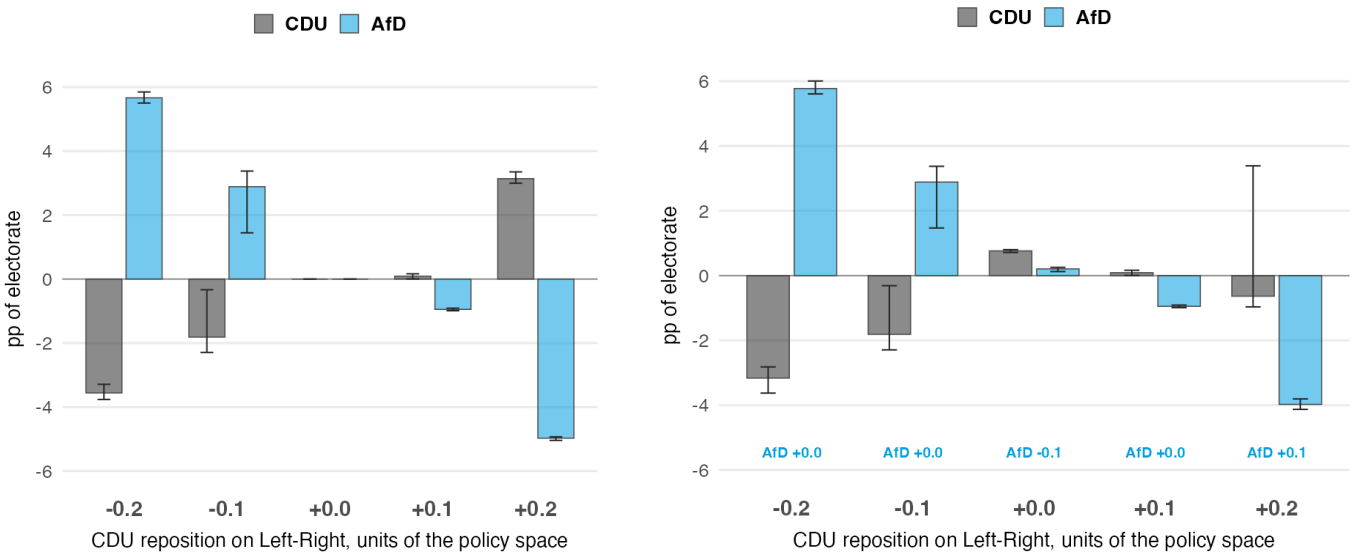
Figure 13a shows what lies behind the 0.2 step: the CDU captures the wedge of the AfD's area closest to it on the Left–Right dimension (dark gray) and loses a thin strip on its left flank to the FDP (hatched). The wedge is large because, as Figure 8 showed, a substantial mass of AfD voters holds relatively pro-elite, radical-right positions, and the CDU's move brings it closer to them than the AfD is. If the AfD cannot respond, accommodation pays, and pays well.

When the AfD responds on Left–Right. The conclusion above rests on the AfD standing still while the CDU takes its voters. I first study the game in which both parties can move only on the Left–Right dimension.[35] Figure 12b reports the median net changes of the CDU and the AfD after the AfD's best reply to each Left–Right move of the CDU.

The AfD's modal reply to the CDU's +0.2 is a further hardening of its own, +0.1 to the right. Figure 13b shows this reply and illustrates why: the AfD wins back the far-right wedge the CDU had captured (purple). Section 6.2.1 explains why this is profitable: in later years, large shares of AfD voters have ideal points far to the right, and many are relatively pro-elite.

The game is not zero-sum, since the other parties also gain and lose voters along the way, but the AfD's reply is aimed at the CDU. This is one possible answer of a populist party to accommodation: when the establishment party moves towards it, it does not concede the contested voters but radicalizes further,

[35] This restriction reflects the view that Left–Right positions are easier to change credibly than anti-elite ones.



(a) The AfD held at its 2025 position

(b) After the AfD's best reply

Figure 12. Estimated net CDU and AfD vote share gain relative to the status quo by CDU Left–Right moves, 2025

Notes. Bars show the median net change in each party's vote share across 100 draws from the identified set, in percentage points of the vote share; whiskers show interquartile ranges. National payoffs are population-weighted sums of the East and West changes. Both parties keep their 2025 Anti-Elite positions and move on Left–Right only, in steps of 0.1 up to ±0.2. In panel (a) the AfD does not move; in panel (b) it plays its best reply to each CDU move, and the labels under the bars give the modal reply across draws.

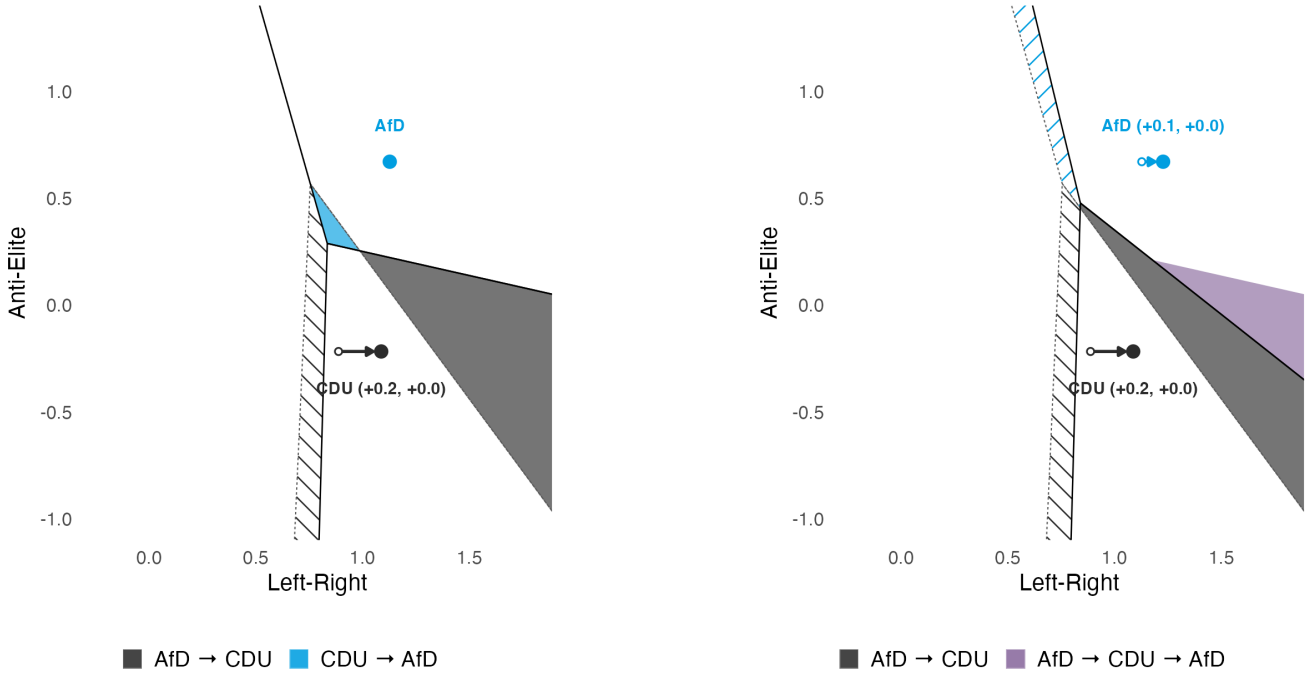
moving away from the challenger along the same dimension. The AfD's net change remains negative: the reply recovers part of its losses, not all of them. For the CDU, the +0.2 step is no longer profitable once the reply is taken into account (−0.6 points at the median).

The Stackelberg equilibrium across preference draws features the CDU staying put in 70% of draws, and moving +0.2 or −0.1 in the other 20% and 10% respectively.

Equilibrium in both dimensions. Appendix Figure J22 reports the CDU's and the AfD's net changes after the AfD's best reply when both parties can move on both dimensions. The CDU's 0.2 step to the right leaves both parties worse off: the CDU is left with −0.5 points and the AfD with −3.0 points at the median. The AfD's best reply to the CDU's $(+0.2, 0)$ is $(+0.1, +0.2)$ or $(+0.1, +0.1)$ in every draw: it hardens its Left–Right position and sharpens its anti-elite profile at the same time. The Left–Right step wins back the far-right wedge, as before; hardening on anti-elite allows the AfD to capture anti-elite voters with relatively center-right positions, at the cost of an anti-elite step that is cheap when the Anti-Elite dimension carries little weight ($w_{AE} = 0.2$ in the West).

Once the AfD's response is anticipated, moving right is no longer the CDU's best move. The Stackelberg equilibrium has the CDU staying put in 64% of draws, moving −0.1 in 34%, and moving +0.2 in only 2%.

The exercises in this section confirm the conventional wisdom that establishment parties can capture



(a) The CDU moves $(+0.2, 0)$, the AfD stands still

(b) The AfD replies $(+0.1, 0)$

Figure 13. Areas of attraction under the counterfactual moves, 2025

Notes. Each panel shows the (Left–Right, Anti-Elite) policy space with the CDU (dark gray) and the AfD (blue) at their 2025 national medians and the other parties fixed; open circles and arrows mark the moves. Dotted gray lines are the boundaries of the two parties' areas of attraction at the status quo, solid black lines the boundaries after the moves. Shaded areas are voters who change parties between the two players: dark gray, from the AfD to the CDU; blue, from the CDU to the AfD; purple, captured by the CDU in panel (a) and won back by the AfD in panel (b). Hatching marks areas lost to the other parties. Partitions are drawn with the estimated West weight $w_{AE} = 0.2$.

some populist voters by moving towards the populists' positions. But this logic overlooks the populists' strategic response. Once that response is taken into account, (i) in equilibrium the establishment party is better off staying put, and (ii) a rightward move by the establishment party could leave the party system as a whole more radicalized than before.

Lessons and caveats of the exercise. First, the multidimensionality of the setup is essential. In one dimension the exercise would be trivial: the populist party's reply to an approaching rival is to move towards it, not away from it – its outer flank is already its own, and hardening further would only cede the contested middle. Hardening becomes a best reply only once the rival overtakes it. In a weighted two-dimensional space this no longer holds: hardening can be a best reply even without overtaking.

Second, what matters for the results is the distribution of voters' ideal points in the policy space: parties' best replies are determined by where the mass of voters sits relative to the boundaries they can move.

Third, the equilibrium analysis holds the positions of all other parties fixed. Some features of the equilibria, such as the AfD moving to the left, partly reflect this restriction: competition on each player's

outer flank is one-sided, whereas competition between the two players is a genuine game. The main answer to the question this section asks is less sensitive to it: once the AfD can respond, moving right is not the CDU's best move. Considering the populists' reaction is thus essential when assessing the establishment's options, a channel long emphasized in the theoretical literature on party competition and entry (Palfrey, 1984) but rarely quantified.

9 Discussion

The framework developed in this paper rests on several modeling choices. I now discuss alternatives to these choices, as well as possible extensions.

Probabilistic voting model. The model considered in the main text is Downsian: voters care only about the policy platforms of the candidates. Probabilistic voting models augment this framework with another component, interpreted as an individual-specific taste for a candidate (Persson and Tabellini, 2002). I now show that the estimated model can be viewed as a probabilistic voting model with particular restrictions on the distribution of the error term.

The estimated model is a model with *random parameters*: the object of estimation is the distribution of a parameter (here, the ideal point). Let X_i be the average ideal point of type i, ϵ_i an additive shock defining the deviation from it, and Y_j the position of candidate j:

$$U_{ij}^{RPM} = -\,dist(X_i + \epsilon_i,\ Y_j)\,.$$

The probabilistic voting model is instead a *random utility* model,

$$U_{ij}^{RUM} = -\,dist(X_i,\ Y_j) + \bar{\epsilon}_{ij},$$

with $\bar{\epsilon}_{ij}$ usually assumed to be i.i.d. across candidates and independent of X_i. For ease of exposition here and below in this section I assume the quadratic distance function and omit subscripts. Expanding the utility,

$$U_j^{RPM} = -(X + \epsilon - Y_j)^2 = -(X - Y_j)^2 + \underbrace{2\epsilon(Y_j - X) - \epsilon^2}_{\nu_j}\,.$$

The random-parameter model is therefore itself a random utility model, and a probabilistic voting model, $U_j^{RPM} = -(X - Y_j)^2 + \nu_j$, in which the distribution of the error is induced by the distribution of the ideal-point shock rather than assumed directly. The two models differ only in the restrictions placed on the joint distribution of the errors.

I assume the ideal-point shock is independent of both voter type and candidate positions, $\epsilon_i \perp (X_i, Y)$. Independence from X_i requires that the distribution of deviations does not vary systematically with the

type's position; independence from Y rules out candidates positioning in response to the realization of the shock. The former just states that given the type, the deviations are pure noise. The latter is the stronger identification restriction, discussed above: it rules out that, within a type, deviations are shifted towards the candidates' positions.[36]

Note that this treatment of the framework also clarifies the consequences of potential misspecifications, such as dropping the probabilistic voting term, omitting important dimensions, or measurement error in the candidate positions. I briefly discuss the effect of the first below, as it belongs to the probabilistic voting model, and turn to the others in the following paragraphs.

If there is a non-spatial *taste component* Z_{ij} orthogonal to both X_i and Y_j, dropping it does not bias the recovered ideal points, as orthogonality rules out any systematic tilt towards particular positions. It could nonetheless inflate the dispersion of the estimated distributions for the ideal points, as choices driven by *taste* are rationalized *spatially*.

More than two dimensions. The framework restricts attention to two dimensions, and one could argue that others matter: nationalism, anti-globalism, or attitudes towards the European Union. In principle the framework accommodates more dimensions; in practice this is infeasible with the data at hand. Moreover, the candidate surveys contain no items suitable for estimating positions on nationalism or anti-globalism, and only two items for the EU dimension.

How would my results be affected if there is another dimension which my estimates miss? Consider the cleanest case: one studied dimension with weight w and positions Y_j, and one omitted dimension with weight w^o, positions $Y_j^o = \gamma Y_j + e_j$, $\gamma > 0$ (dimensions bundled at the candidate level), and voter component X^o. Set $e_j = 0$.[37] True utility is

$$U_j = -w\,(X - Y_j)^2 - w^o\,(X^o - Y_j^o)^2.$$

Collecting the terms that depend on j and dropping alternative-invariant ones,

$$U_j \;\sim\; -(w + w^o\gamma^2)\,Y_j^2 + 2\,(wX + w^o\gamma X^o)\,Y_j \;=\; -(w + w^o\gamma^2)\left(Y_j - \tilde{X}\right)^2 + \text{const},$$

where

$$\tilde{X} = \frac{w\,X + w^o\gamma\,X^o}{w + w^o\gamma^2}, \qquad \tilde{w} = w + w^o\gamma^2.$$

[36] Two properties of ν_j follow from ϵ_i being a single voter-level shock rather than a voter-candidate-level shock. First, the term ϵ^2 is common to all candidates and cancels from every utility comparison, so effectively $\nu_j = 2\epsilon(Y_j - X)$ up to a candidate-invariant constant. Second, because all ν_j are driven by the same draw of ϵ, the errors are correlated across candidates, the more strongly the closer the candidates' positions in the policy space. Consequently, whenever the number of candidates exceeds the dimension of the policy space plus one, no distribution of the ideal-point shock reproduces an i.i.d. error vector: the induced errors are confined to a lower-dimensional subspace. Exact equivalence obtains, e.g., in the case of two candidates.

[37] With $e_j \neq 0$, further terms arise: a voter-level factor error $2w^oX^oe_j$, absorbed into the ideal-point shock when X^o is independent of (X, Y) within type; and candidate-specific intercepts $-2w^o\gamma\,e_jY_j - w^oe_j^2$, which a pure characteristics model rationalizes spatially. All are proportional to w^o and vanish in the same limit.

The misspecified one-dimensional model thus retains the correct form but has distorted primitives: the recovered ideal point is $\tilde{X}$, a weighted average of the true position on the studied dimension and the omitted-dimension component, and the recovered weight is $\tilde{w} = w + w^o\gamma^2$, inflated by the omitted dimension's importance and its correlation with the studied one. Both distortions vanish as (1) $\gamma \to 0$: dimensions along which candidates' positions are unrelated to the studied ones can be omitted without bias; or (2) $w^o \to 0$: unimportant dimensions can be omitted without bias, regardless of bundling.

Hence, my approach is not robust to omitting relevant dimensions correlated with the included ones. An omitted dimension is projected onto the included ones, and how it is projected depends on its correlation with each of them. For instance, if the correct model also included a Nationalism dimension, the Left–Right estimates in my model would partly capture it, since nationalist positions are typically correlated with Left–Right positions (see, e.g., the correlation across European parties in the Chapel Hill Expert Survey, Appendix Figure K3).

By contrast, the Anti-Elite and Left–Right dimensions are essentially uncorrelated, both across parties in the same expert-survey data and in my sample. This makes them well suited for a two-dimensional representation: in the expert-survey data, the two dimensions jointly explain 72% of the total variance of party positions across eight policy dimensions.

Left–Right as a combination of Economics and Culture. So far the paper has treated the Left–Right dimension as a bundle of economic and cultural policies without unpacking the two. Estimating the model in three dimensions is infeasible given the data, so I instead re-estimate it in two: Economics and Culture. Omitting the Anti-Elite dimension here is relatively innocuous: it is only weakly correlated with both (Appendix Figure K3), so, following the discussion above, its omission is unlikely to affect the recovered preferences.

To do so, I first split the Left–Right items into those concerning culture and those concerning economics; see Appendix Table K2 for the classification.[38] I then follow exactly the procedure of Section 6, applied to these two dimensions.

Appendix Figure J19 reports the estimated weights by year and part of the country, Appendix Figure J20 reports the evolution of preferences across the dimension bins, and Appendix Figure J21 reports the counterfactual motivation exercise, which splits the AfD into two hypothetical parties: one identical to the AfD on culture but economically liberal, the other identical on economics but culturally liberal. The decomposition shows a gradual realignment of voting motives, from economic ones, which clearly dominate in 2013, to cultural ones, which clearly dominate in 2025. In the West, voters attach a much larger weight to the cultural than to the economic dimension from 2017 onward; in the East the two weights are closer, although imprecisely estimated. AfD voters' ideal points are relatively stable over time on the cultural dimension, but

[38] Two coding decisions are worth noting: immigration items are classified as cultural (Alesina and Tabellini, 2024), and the Left–Right self-placement item is dropped.

shift slightly to the left on the economic one.[39] These results are partly in line with Danieli et al. (2026), who attribute the rise of radical-right parties largely to an increase in the weight voters attach to cultural relative to economic issues.

Candidate positions estimated with an error. Assume the candidate positions are measured with error: the estimation uses $\bar{Y}_j = Y_j + \xi_j$ in place of the true Y_j. The effective utility is then

$$U_j^{RPM} = -(X + \epsilon - Y_j - \xi_j)^2,$$

so the error enters as a perturbation of the candidate's location rather than as an additional disturbance: unlike ϵ_i, which varies at the voter level, ξ_j is common across voters, and the recovered ideal points are located relative to the perturbed positions. The relevant question is therefore whether the objects of interest are stable under small perturbations of the positions. Appendix Figure J18 assesses this for the West 2025 results: I add an i.i.d. $N(0, 0.01^2)$ shock to each coordinate of every estimated district-level party position (and of the national anchor of the counterfactual parties) and re-estimate the target parameter of Section 6.2.2 for each of 10 draws. A standard deviation of 0.01 is not small relative to the cell partition: about 35% of the cells are thinner than 0.02 (two s.d.). The shocks move the bounds without widening them, and on average pull them toward the uninformative composition, while preserving the overall results.[40] Moderate error in the measured positions therefore weakens but does not systematically overturn the conclusions. In the limit where positions carry no district-level information, with every district in a state assigned its state-median party positions, the bounds on the target parameters become uninformative.

The Role of the supply side. To what extent are the results presented above driven by supply-side repositioning rather than by demand-side changes (genuine evolution of preferences, or compositional changes)? To answer this question, I decompose the change in the target parameter of Section 6.2.2 into three channels (party positions, the composition of the AfD electorate, and within-group preferences) using a Shapley-value approach (Appendix G provides details).[41]

The results suggest that the evolution of the share of Right-motivated voters documented in this paper is driven primarily not by parties changing their positions but by demand-side changes: the evolution of preferences within demographic groups, the changing demographic composition of the AfD electorate, or both. The data do not allow me to separate these two demand-side channels, as the bounds on their Shapley values remain uninformative (see Appendix Table G2).

[39] These results also speak to the AfD's transformation around 2015. In 2013, the party, then dubbed the "professors' party" after its economist founders and centered on opposition to the euro, drew a different electorate: up to half of its voters were economically conservative, and they attached little weight to the cultural dimension.

[40] The mechanism resembles attenuation bias from measurement error in OLS: the shocks scramble the cross-election consistency of the cell labels that the *correct* distributions exploit.

[41] See Shapley (1953) for the original reference, or Guriev et al. (2026); Longuet-Marx (2026) for recent applications in political economy.

Firewall against the far right. Germany's mainstream parties maintain a firewall (*Brandmauer*) against the AfD, a commitment not to cooperate with the party, including in government. Appendix H applies the estimated 2025 preferences to this firewall, comparing the two coalitions that were numerically feasible after the 2025 election, CDU–SPD and CDU–AfD, with each coalition assumed to implement the midpoint of its members' positions. I evaluate them under two criteria: utilitarian welfare, measured by the average weighted distance from voters' ideal points to each coalition policy, and majoritarian choice, measured by the share of voters closer to one policy than the other. The two criteria disagree. The national welfare bounds straddle zero ($[-0.18,\ 0.03]$), so neither coalition dominates, although the bounds are strictly negative in the East, favoring CDU–AfD. Yet for every distribution in the identified set a clear majority is closer to the CDU–AfD policy (61.9–72.8% nationally; at least 69.9% in the East and 58.5% in the West). The gap between the two criteria reflects intense opposition to a CDU–AfD coalition among voters on the left, consistent with survey evidence on coalition desirability. Whether the firewall leaves voters underrepresented therefore depends on the aggregation rule, and the tension is sharpest in Eastern Germany.

Approach of Merlo and De Paula (2017). I compare the estimates obtained with the method developed in this paper with those from the approach of Merlo and De Paula (2017), who obtain point estimates of the same spatial model by restricting the ideal-point distribution to a semi-nonparametric class and estimating it by sieve minimum distance. First, Appendix D reports simulations comparing the two approaches to estimating the weight. Second, Appendix F re-estimates the results of Section 6.2.2 and the repositioning game of Section 8 with their estimator, applied to the same data, demographic groups, and counterfactual masks. Their point estimates lie inside or close to my bounds in 2017–2025 but not in 2013, and display no systematic trend: the dominant motive switches between Anti-Elite and Right from one election to the next. In the game, the AfD's best reply under their specification is always a step to the right, so in the Stackelberg equilibrium both parties move right. These comparisons highlight the role that functional-form assumptions can play in this setting.

10 Conclusion

Is populist support rooted in Left–Right ideology or in protest against the elites? How could parties of the establishment respond to the rise of populism? This paper develops a novel approach to recovering voters' preferences from aggregate electoral data in a multidimensional spatial voting model, and uses it to answer these questions.

The paper offers a framework for unpacking preferences over the bundle of characteristics embodied in any political option, building on recent advances in the nonparametric discrete choice literature (Tebaldi, Torgovitsky and Yang, 2023). Any voting option, whether a candidate, a party, or a referendum answer, can be seen as a product with multiple attributes. Voting departs from most settings in empirical industrial organization: in a spatial model, individuals hold unobserved ideal points on each dimension, so one cannot impose the linear characteristic-utility relationship standard in that literature. The question *"Do you prefer*

sweeter or less sweet cereal?" does not translate into *"Do you prefer a more or less right-wing candidate?"* since voters differ not in how much of an attribute they want, but in where their ideal point lies.

The approach is broadly applicable, requiring only that voters face different sets of options, as arises naturally across districts, in simultaneous elections for several offices, or under ranked-choice voting.

Applying the approach to the rise of the radical-right populist party Alternative for Germany (AfD), I find that the party's support became increasingly grounded in its right-wing ideology rather than its anti-elite stance over 2013–2025. Counterfactuals suggest that in 2025, if the AfD stood still, the main center-right establishment party, the CDU, could capture a substantial share of AfD voters by modestly hardening its Left–Right position, for instance on immigration. This opening is illusory once the AfD can respond: in a repositioning game between the two parties, the AfD answers a rightward CDU move by moving further right and wins these voters back, so in equilibrium the CDU does not move right.

The insights of this paper are therefore uncomfortable: accommodation of populist policies may backfire and give rise to an even more radicalized politics. What, then, could establishment parties do? The anti-elite dimension captures trust in the institutions of representative democracy: anti-elite voters see a self-serving establishment opposed to ordinary people and favor direct votes over decisions by elected representatives, while "pro-elite" voters trust politicians, parties and institutions to act broadly in the public interest. This suggests two avenues. First, establishment parties could raise the salience of this dimension. In recent elections a non-negligible share of AfD voters holds relatively pro-elite attitudes yet still chooses the AfD, because it values the bundle of right-wing and anti-elite rhetoric more than the bundle offered by the CDU. Stressing what is at stake for democratic norms and institutions would make the AfD's anti-elite stance costlier for these voters. Second, parties, especially while in government, could make governing more transparent and participatory, through measures such as deliberative democracy or inclusive policy-making; both could shift voters' ideal points towards the pro-elite end. Although derived from entirely different premises and frameworks, these avenues mirror those in Guriev and Papaioannou (2022) and Guriev (2024).

References

Abou-Chadi, Tarik and Werner Krause. 2020. "The causal effect of radical right success on mainstream parties' policy positions: A regression discontinuity approach." *British Journal of Political Science* 50(3):829–847.

Acemoglu, Daron, Georgy Egorov and Konstantin Sonin. 2013. "A political theory of populism." *The Quarterly Journal of Economics* 128(2):771–805.

AfD. 2013. "Wahlprogramm: Parteitagsbeschluss vom 14.04.2013." Party manifesto. Accessed: 2026-02-15. **URL:** *https://www.bundestagswahl-2021.de/wp-content/uploads/sites/4/2021/04/Wahlprogramm-der-AfD-Bundestagswahl-2013.pdf*

Akkerman, Agnes, Cas Mudde and Andrej Zaslove. 2014. "How populist are the people? Measuring populist attitudes in voters." *Comparative Political Studies* 47(9):1324–1353.

Alesina, Alberto and Marco Tabellini. 2024. "The political effects of immigration: Culture or economics?" *Journal of Economic Literature* 62(1):5–46.

Algan, Yann, Elizabeth Beasley, Daniel Cohen and Martial Foucault. 2018. "The rise of populism and the collapse of the left-right paradigm: lessons from the 2017 French presidential election." CEPR Discussion Paper No. DP13103.

Algan, Yann, Sergei Guriev, Elias Papaioannou and Evgenia Passari. 2017. "The European trust crisis and the rise of populism." *Brookings Papers on Economic Activity* 2017(2):309–400.

Autor, David, David Dorn, Gordon Hanson and Kaveh Majlesi. 2020. "Importing political polarization? The electoral consequences of rising trade exposure." *American Economic Review* 110(10):3139–3183.

Bakker, Ryan, Catherine De Vries, Erica Edwards, Liesbet Hooghe, Seth Jolly, Gary Marks, Jonathan Polk, Jan Rovny, Marco Steenbergen and Milada Anna Vachudova. 2015. "Measuring party positions in Europe: The Chapel Hill expert survey trend file, 1999–2010." *Party Politics* 21(1):143–152.

Bartels, Larry M. 2023. *Democracy erodes from the top: Leaders, citizens, and the challenge of populism in Europe*. Princeton, NJ: Princeton University Press.

Berry, Steven T, Christian Cox and Philip Haile. 2025. Selective turnout, voting policy, and partisan bias: Evidence from multi-level data. Technical report National Bureau of Economic Research.

Bonomi, Giampaolo, Nicola Gennaioli and Guido Tabellini. 2021. "Identity, beliefs, and political conflict." *The Quarterly Journal of Economics* 136(4):2371–2411.

Bundesministerium für Wirtschaft und Energie. 2025. Jahresbericht der Bundesregierung zum Stand der Deutschen Einheit 2025 – 35 Jahre: Aufgewachsen in Einheit? Bericht Bundesministerium für Wirtschaft und Energie (BMWE), Beauftragte der Bundesregierung für Ostdeutschland.

Callander, Steven and Juan Carlos Carbajal. 2022. "Cause and effect in political polarization: A dynamic analysis." *Journal of Political Economy* 130(4):825–880.

Colantone, Italo and Piero Stanig. 2018. "The trade origins of economic nationalism: Import competition and voting behavior in Western Europe." *American Journal of Political Science* 62(4):936–953.

Cox, Karen E. and Leonard J. Schoppa. 2002. "Interaction effects in mixed-member electoral systems: Theory and evidence from Germany, Japan, and Italy." *Comparative Political Studies* 35(9):1027–1053.

Danieli, Oren, Noam Gidron, Shinnosuke Kikuchi and Ro'ee Levy. 2026. "Decomposing the Rise of the Populist Radical Right." Working paper.

Degan, Arianna and Antonio Merlo. 2009. "Do voters vote ideologically?" *Journal of Economic Theory* 144(5):1868–1894.

Desmet, Klaus, Ignacio Ortuño-Ortín and Romain Wacziarg. 2024. Latent Polarization. Technical report National Bureau of Economic Research.

Detterbeck, Klaus. 2016. "Candidate selection in Germany: Local and regional party elites still in control?"

American Behavioral Scientist 60(7):837–852.

Deutscher Bundestag, Wissenschaftliche Dienste. 2022. Binnenwanderung in Deutschland. Technical Report WD 1 – 3000 – 040/22 Deutscher Bundestag.
URL: *https://www.bundestag.de/resource/blob/933298/f5ee8aa662372c5210e00b14515fe435/WD-1-040-22-pdf-data.pdf*

Die Bundeswahlleiterin. 2025. "Representative Electoral Statistics." https://www.bundeswahlleiterin.de/en/bundestagswahlen/2025/informationen-waehler/rws.html. Accessed October 2026.

Die Bundeswahlleiterin. n.d. "Wahlkreiseinteilung." Web page.
URL: *https://www.bundeswahlleiterin.de/bundestagswahlen/2025/wahlkreiseinteilung.html*

Dippel, Christian, Robert Gold, Stephan Heblich and Rodrigo Pinto. 2022. "The effect of trade on workers and voters." *The Economic Journal* 132(641):199–217.

Downs, Anthony. 1957. "An economic theory of political action in a democracy." *Journal of Political Economy* 65(2):135–150.

Draca, Mirko and Carlo Schwarz. 2024. "How polarised are citizens? Measuring ideology from the ground up." *The Economic Journal* 134(661):1950–1984.

Dustmann, Christian, Kristine Vasiljeva and Anna Piil Damm. 2019. "Refugee migration and electoral outcomes." *The Review of Economic Studies* 86(5):2035–2091.

Ferrara, Federico, Erik Herron and Misa Nishikawa. 2005. *Mixed electoral systems: Contamination and its consequences*. Springer.

Fiorina, Morris P and Samuel J Abrams. 2008. "Political polarization in the American public." *Annual Review of Political Science* 11(1):563–588.

Gallant, A Ronald and Douglas W Nychka. 1987. "Semi-nonparametric maximum likelihood estimation." *Econometrica* 55(2):363–390.

Gennaioli, Nicola and Guido Tabellini. 2025. "Presidential Address: Identity Politics." *Econometrica* 93(6):1937–1967.

Gerber, Elisabeth R. and Jeffrey B. Lewis. 2004. "Beyond the Median: Voter Preferences, District Heterogeneity, and Political Representation." *Journal of Political Economy* 112(6):1364–1383.

Gethin, Amory, Clara Martínez-Toledano and Thomas Piketty. 2022. "Brahmin left versus merchant right: Changing political cleavages in 21 Western Democracies, 1948–2020." *The Quarterly Journal of Economics* 137(1):1–48.

GLES. 2014. "Kandidatenstudie 2013, Befragung, Wahlergebnisse und Strukturdaten (GLES)." (ZA5716; Version 1.0.0.) [Data set]. GESIS, Köln.

GLES. 2018. "Kandidatenstudie (GLES 2017)." (ZA6814; Version 3.0.0) [Data set]. GESIS, Köln.

GLES. 2023. "GLES Kandidierendenstudie 2021." (ZA7704; Version 2.0.0) [Data set]. GESIS, Köln.

GLES. 2025. "GLES Kandidierendenstudie 2025." (ZA10102; Version 1.0.0) [Data set]. GESIS, Köln.

Gschwend, Thomas and Thomas Zittel. 2015. "Do constituency candidates matter in German Federal Elections? The personal vote as an interactive process." *Electoral Studies* 39:338–349.

Gu, Jiaying, Thomas M Russell and Thomas Stringham. 2026. "Counterfactual identification and latent space enumeration in discrete outcome models." *The Review of Economic Studies* 93(3):1847–1888.

Guenther, Laurenz and Salvatore Nunnari. 2025. Do Political Representation Gaps Cause Populism?: Evidence from the 2025 German Election. Cepr discussion paper Centre for Economic Policy Research.

Guiso, Luigi, Helios Herrera, Massimo Morelli and Tommaso Sonno. 2017. Demand and supply of populism. CEPR Discussion Paper 11871 Centre for Economic Policy Research.

Guiso, Luigi, Massimo Morelli, Tommaso Sonno and Helios Herrera. 2021. "The financial drivers of populism in Europe." *BAFFI CAREFIN Centre Research Paper* (2021-166).

Guriev, Sergei. 2024. "Causes and consequences of spreading populism: How to deal with this challenge." *EconPol Forum* 25(2):5–8.

Guriev, Sergei and Elias Papaioannou. 2022. "The political economy of populism." *Journal of Economic Literature* 60(3):753–832.

Guriev, Sergei, Emeric Henry, Théo Marquis and Ekaterina Zhuravskaya. 2026. "Curtailing False News, Amplifying Truth." *Econometrica* 94(5):1563–1601.

Hainmueller, Jens and Holger Lutz Kern. 2008. "Incumbency as a source of spillover effects in mixed electoral systems: Evidence from a regression-discontinuity design." *Electoral Studies* 27(2):213–227.

Hangartner, Dominik, Elias Dinas, Moritz Marbach, Konstantinos Matakos and Dimitrios Xefteris. 2019. "Does exposure to the refugee crisis make natives more hostile?" *American Political Science Review* 113(2):442–455.

Herron, Erik S. and Misa Nishikawa. 2001. "Contamination effects and the number of parties in mixed-superposition electoral systems." *Electoral Studies* 20(1):63–86.

Hinich, Melvin J and Michael C Munger. 1996. *Ideology and the theory of political choice*. University of Michigan Press.

Hofmann, Elias and Heiko Rüger. 2024. "Umzugsmotive in Deutschland: Wie unterscheiden sich die Motive mit der Umzugsdistanz?" *Bevölkerungsforschung Aktuell* 45(1):10–13.
URL: *https://nbn-resolving.org/urn:nbn:de:bib-bfa0120241*

Höhne, Benjamin. 2023. "How democracy works within a populist party: candidate selection in the Alternative for Germany." *Government and Opposition* 58(3):478–496.

Iaryczower, Matias, Sergio Montero and Galileu Kim. 2024. "Representation failure." Working paper.

Ichimura, Hidehiko and T Scott Thompson. 1998. "Maximum likelihood estimation of a binary choice model with random coefficients of unknown distribution." *Journal of Econometrics* 86(2):269–295.

Imbens, Guido W and Charles F Manski. 2004. "Confidence intervals for partially identified parameters." *Econometrica* 72(6):1845–1857.

Infratest dimap. 2017. "Wählerwanderung, Bundestagswahl 2017." Wahlberichterstattung im Auftrag der ARD, reported in Konrad-Adenauer-Stiftung, Tabellenanhang zur Bundestagswahl 2017, https://www.kas.de/c/document_library/get_file?uuid=4b5bf86d-991e-b23c-8029-7932d03cb8aa&groupId=268877.

Infratest dimap. 2025. "Wählerwanderung, Bundestagswahl 2025." ARD-DeutschlandTrend / tagesschau.de.

Jolly, Seth, Ryan Bakker, Liesbet Hooghe, Gary Marks, Jonathan Polk, Jan Rovny, Marco Steenbergen and Milada Anna Vachudova. 2022. "Chapel Hill expert survey trend file, 1999–2019." *Electoral Studies* 75:102420.

Kawai, Kei and Takeaki Sunada. 2025. "Estimating candidate valence." *Econometrica* 93(2):463–501.

Kawai, Kei and Yasutora Watanabe. 2013. "Inferring strategic voting." *American Economic Review* 103(2):624–662.

Klingemann, Hans-Dieter and Bernhard Wessels. 2001. The political consequences of Germany's mixed-member system: Personalization at the grass roots? In *Mixed-Member Electoral Systems: The Best of Both Worlds?*, ed. Matthew S. Shugart and Martin P. Wattenberg. Oxford: Oxford University Press pp. 279–296.

Krause, Werner, Denis Cohen and Tarik Abou-Chadi. 2023. "Does accommodation work? Mainstream party strategies and the success of radical right parties." *Political Science Research and Methods* 11(1):172–179.

Le Pennec, Caroline and Vincent Pons. 2019. How do campaigns shape vote choice? Multi-country evidence from 62 elections and 56 TV debates. Technical report National Bureau of Economic Research.

Lieber, Jonas, Pietro Tebaldi and Alexander Torgovitsky. 2024. "Semiparametric Equilibrium Counterfactuals: Emission Policies in the New and Used Car Market." Working Paper.

Longuet-Marx, Nicolas. 2026. "Party lines or voter preferences? Explaining political realignment." Working paper.

Margalit, Yotam. 2019. "Economic insecurity and the causes of populism, reconsidered." *Journal of Economic Perspectives* 33(4):152–170.

Meguid, Bonnie M. 2005. "Competition between unequals: The role of mainstream party strategy in niche party success." *American Political Science Review* 99(3):347–359.

Merlo, Antonio and Áureo De Paula. 2017. "Identification and estimation of preference distributions when voters are ideological." *The Review of Economic Studies* 84(3):1238–1263.

Mudde, Cas. 2004. "The populist zeitgeist." *Government and Opposition* 39(4):541–563.

Mudde, Cas. 2013. "Three decades of populist radical right parties in Western Europe: So what?" *European Journal of Political Research* 52(1):1–19.

Neuner, Fabian G and Christopher Wratil. 2022. "The populist marketplace: Unpacking the role of "thin" and "thick" ideology." *Political Behavior* 44(2):551–574.

Neuner, Fabian G, Kirill Zhirkov, Bruno Castanho Silva and Christopher Wratil. 2025. "Comparing Revealed and Expressed Populism: The Case of Voting for the 'Alternative für Deutschland' and 'Die Linke'." *German Politics* pp. 1–23.

Nikiforova, Nina. 2026. "What matters to voters? Identifying important ideological dimensions with partial rankings of parties." Working paper.

Norris, Pippa and Ronald Inglehart. 2019. *Cultural backlash: Trump, Brexit, and authoritarian populism*. Cambridge University Press.

Osborne, Martin. 2025. *Models in Political Economy: Collective Choice, Voting, Elections, Bargaining, and Rebellion*. Open Book Publishers.

Osborne, Martin J and Al Slivinski. 1996. "A model of political competition with citizen-candidates." *The Quarterly Journal of Economics* 111(1):65–96.

Palfrey, Thomas R. 1984. "Spatial equilibrium with entry." *The Review of Economic Studies* 51(1):139–156.

Persson, Torsten and Guido Tabellini. 2002. *Political economics: explaining economic policy*. MIT Press.

Poole, Keith T and Howard Rosenthal. 1985. "A spatial model for legislative roll call analysis." *American Journal of Political Science* pp. 357–384.

Rodrik, Dani. 2021. "Why does globalization fuel populism? Economics, culture, and the rise of right-wing populism." *Annual Review of Economics* 13:133–170.

Schläger, Catrina, Jan Niklas Engels and Nicole Loew. 2025. Analyse der Bundestagswahl 2025: Eine harte Niederlage mit einer doppelten Herausforderung für die Sozialdemokratie. Analyse Friedrich-Ebert-Stiftung. **URL:** *https://epflicht.ulb.uni-bonn.de/download/pdf/786158*

Schroeder, Wolfgang, Daniel Ziblatt and Florian Bochert. 2025. "Is the Brandmauer ("firewall") breaking from below? Party cooperation with the "Alternative for Germany (AfD)" in all German local government councils (2019-2024)." Working paper.

Schüttemeyer, Suzanne S, Pia Berkhoff, Malte Cordes, Oscar W Gabriel, Daniel Hellmann, Benjamin Höhne, Oliver Kannenberg, Sophie Kopsch, Oskar Niedermayer, Anastasia Pyschny et al. 2024. *Die Aufstellung der Kandidaten für den Deutschen Bundestag*. Baden-Baden: Nomos.

Shapley, Lloyd S. 1953. A value for n-person games. In *Contributions to the Theory of Games, Volume II*, ed. Harold W. Kuhn and Albert W. Tucker. Princeton, NJ: Princeton University Press pp. 307–317.

Silva, Bruno Castanho, Mario Fuks and Eduardo Ryô Tamaki. 2022. "So thin it's almost invisible: Populist attitudes and voting behavior in Brazil." *Electoral Studies* 75:102434.

Snyder, James M., Jr. 1996. "Constituency Preferences: California Ballot Propositions, 1974–1990." *Legislative Studies Quarterly* 21(4):463–488.

Spoon, Jae-Jae and Heike Klüver. 2020. "Responding to far right challengers: does accommodation pay off?" *Journal of European Public Policy* 27(2):273–291.

Statistisches Bundesamt. 2025. "Wanderungen: Statistischer Bericht 2024." Wanderungsstatistik, Wiesbaden. https://www.destatis.de/DE/Themen/Gesellschaft-Umwelt/Bevoelkerung/Wanderungen/_inhalt.html.

Tebaldi, Pietro, Alexander Torgovitsky and Hanbin Yang. 2023. "Nonparametric estimates of demand in the California health insurance exchange." *Econometrica* 91(1):107–146.

Ujhelyi, Gergely, Somdeep Chatterjee and Andrea Szabó. 2021. "None of the Above: Protest Voting in the World's Largest Democracy." *Journal of the European Economic Association* 19(3):1936–1979.

van der Brug, Wouter, Sara B Hobolt and Sebastian Adrian Popa. 2026. "The kids are Alt right? Age, authoritarian attitudes and far-right support in Europe." *Journal of European Public Policy* 33(2):469–494.

Van Kessel, Stijn. 2015. *Populist parties in Europe: Agents of discontent?* Springer.

Appendix

A Finite-dimensional problem: bounds on a target parameter **54**

B Identification of dimension weights **55**

B.1 Setting and notation 55

B.2 Set identification of weights 57

B.3 Target parameter bounds with weights 58

C Model Estimation **59**

C.1 Setup and notation 59

C.2 Weights 61

C.3 Bounds on target parameters 62

C.4 Model fit 63

D Weights estimation: Simulation **66**

E Stated Preferences **69**

E.1 Preference bins 69

E.2 Motivations exercise 70

F Comparison with the Merlo–De Paula Point Estimates **71**

G The Role of the Supply Side: Shapley Values **73**

H Coalition-Building and the Firewall Against Cooperation with the Far Right **75**

I Tables & Figures **78**

J Additional Results **91**

K Candidate Data **104**

K.1 Candidate nominations across state lists and single-member districts 104

K.2 Estimation of candidates' positions 105

A Finite-dimensional problem: bounds on a target parameter

This appendix adapts Propositions 1 and 2 of Tebaldi, Torgovitsky and Yang (2023), and their proofs, to the election model of Section 2. The weight w is fixed, and $\mathbb{X}$ is the MRP of Section 2.3, augmented by any hypothetical elections and refined by any regions of the policy space that enter the target parameter; its cells are the sets $\mathcal{X}_\mathbf{j}$ of positive Lebesgue measure. Throughout, f ranges over $\mathcal{F}_{1\mathrm{a}}$, the elements of $\mathcal{F}$ satisfying Assumption 1a, each identified with its collection of state–type densities $\{f(\cdot \mid s,d)\}$.

Cell masses. Let Φ denote the set of conditional mass functions on $\mathbb{X}$,

$$\Phi = \Big\{\phi \in \mathbb{R}_+^{d_\phi} : \sum_{\mathcal{X}\in\mathbb{X}} \phi(\mathcal{X} \mid s,d) = 1 \ \text{ for all } (s,d) \in \operatorname{supp}(S_i, D_i)\Big\},$$

where d_ϕ is the cardinality of $\mathbb{X} \times \operatorname{supp}(S_i, D_i)$. For $f \in \mathcal{F}$, define $\phi_f \in \Phi$ by

$$\phi_f(\mathcal{X} \mid s,d) = \int_{\mathcal{X}} f(x \mid s,d)\, dx \qquad \text{for all } \mathcal{X} \in \mathbb{X},\ (s,d) \in \operatorname{supp}(S_i, D_i). \tag{A1}$$

Recall that $\mathbb{X}_j(e)$ denotes the cells contained in $\mathcal{A}_j(e)$. Since $\mathcal{A}_j(e)$ is the union of these cells up to a set of measure zero, the predicted shares depend on f only through ϕ_f. For $\phi \in \Phi$, consider the condition

$$\mathcal{s}_j(m) = \sum_d q_{m,d} \sum_{\mathcal{X}\in\mathbb{X}_j(e(m))} \phi(\mathcal{X} \mid s(m),d) \qquad \text{for all } j \in \mathcal{J}_{e(m)} \text{ and all municipalities } m. \tag{A2}$$

Then f satisfies (5) if and only if ϕ_f satisfies (A2).

Representability of the target parameter. Assume there exists a known function $\bar\theta : \Phi \to \mathbb{R}^{d_\theta}$ such that

$$\theta(f) = \bar\theta(\phi_f) \qquad \text{for all } f \in \mathcal{F}. \tag{A3}$$

Since Φ depends on the MRP, and the MRP on the elections and regions used to build it, (A3) can always be satisfied by including in $\mathcal{E}$ every hypothetical election relevant for the target parameter and refining $\mathbb{X}$ by every region that enters it. In the leading example, $\bar\theta(\phi) = \phi(\mathcal{X}_1 \mid s,d) + \phi(\mathcal{X}_2 \mid s,d) + \phi(\mathcal{X}_3 \mid s,d)$.

Let $\Phi^* = \{\phi \in \Phi : \phi \text{ satisfies (A2)}\}$ denote the cell-mass functions consistent with the returns.

Proposition 1. *Suppose (A3) holds. For $t \in \mathbb{R}^{d_\theta}$, let $\Phi^*(t) = \Phi^* \cap \{\phi \in \Phi : \bar\theta(\phi) = t\}$. Then $t \in \Theta^*$ if and only if $\Phi^*(t) \neq \emptyset$.*

Proof. Suppose $t \in \Theta^*$, so that $\theta(f) = t$ for some $f \in \mathcal{F}^*$. Then $\phi_f \in \Phi$, since $\mathbb{X}$ is a partition of $\mathcal{P}$ up to a set of measure zero and $f(\cdot \mid s,d)$ is a density on $\mathcal{P}$; ϕ_f satisfies (A2), since f satisfies (5); and $\bar\theta(\phi_f) = \theta(f) = t$ by (A3). Hence $\phi_f \in \Phi^*(t)$.

Conversely, let $\phi \in \Phi^*(t)$ and define

$$\tilde f(x \mid s,d) = \sum_{\mathcal{X}\in\mathbb{X}} \frac{\mathbf{1}\{x \in \mathcal{X}\}}{\lambda(\mathcal{X})}\, \phi(\mathcal{X} \mid s,d),$$

where λ is Lebesgue measure on $\mathbb{R}^2$; every cell has positive measure by construction. The function $\tilde f(\cdot \mid s,d)$ distributes the mass $\phi(\mathcal{X} \mid s,d)$ uniformly over each cell. Because the cells are disjoint up to measure zero, $\int_{\mathcal{X}} \tilde f(x \mid s,d)\, dx = \phi(\mathcal{X} \mid s,d)$ for every $\mathcal{X} \in \mathbb{X}$, so $\phi_{\tilde f} = \phi$ and $\int_{\mathcal{P}} \tilde f(x \mid s,d)\, dx = \sum_{\mathcal{X}\in\mathbb{X}} \phi(\mathcal{X} \mid s,d) = 1$;

together with non-negativity, $\tilde{f}(\cdot \mid s,d)$ is a density on $\mathcal{P}$ for every (s,d); setting $\tilde{f}(x \mid m,d) = \tilde{f}(x \mid s(m),d)$ gives $\tilde{f} \in \mathcal{F}_{1a}$. Since $\mathcal{A}_j(e(m))$ is the union of the cells in $\mathbb{X}_j(e(m))$ up to a set of measure zero,

$$\sigma_j(m;\tilde{f}) = \sum_d q_{m,d} \sum_{\mathcal{X} \in \mathbb{X}_j(e(m))} \phi(\mathcal{X} \mid s(m),d) = \wp_j(m)$$

by (A2), so $\tilde{f}$ satisfies (5) and $\tilde{f} \in \mathcal{F}^*$. Finally, $\theta(\tilde{f}) = \bar{\theta}(\phi_{\tilde{f}}) = \bar{\theta}(\phi) = t$ by (A3), so $t \in \Theta^*$. □

Proposition 2. *Suppose the model is correctly specified, so that $\Phi^* \neq \emptyset$. If $\bar{\theta}$ is continuous on Φ, then Θ^* is compact and connected. In particular, if $d_\theta = 1$, then $\Theta^* = [t^*_{lb}, t^*_{ub}]$, where*

$$t^*_{lb} = \min_{\phi \in \Phi^*} \bar{\theta}(\phi), \qquad t^*_{ub} = \max_{\phi \in \Phi^*} \bar{\theta}(\phi).$$

Proof. Φ is a product of simplices, hence compact and convex. The equations (A2) are linear in ϕ, so Φ^* is a nonempty compact convex polytope, hence connected. By Proposition 1, $\Theta^* = \bar{\theta}(\Phi^*)$, the image of a compact connected set under a continuous map, hence compact and connected. If $d_\theta = 1$, a compact connected subset of $\mathbb{R}$ is a closed interval, whose endpoints are the minimum and maximum of $\bar{\theta}$ over Φ^*. □

When $\bar{\theta}$ is linear in ϕ, as for every target parameter in the application, both bounds are solutions of linear programs.

B Identification of dimension weights

This section characterizes the identified set for the relative weight in the two-dimensional weighted spatial voting model of Section 2. The argument adapts the geometric approach of Lieber, Tebaldi and Torgovitsky (2024): group parameter values according to the voting profiles they permit, and check whether masses on those profiles can reproduce the data. In the spatial voting model, changing the weight rotates the boundaries of the attraction regions. A change of variables turns these rotations into affine changes in the thresholds of a system of linear inequalities, making the same approach applicable.

The argument has two steps. First, each voting profile is available only for weights in an interval. Second, these intervals divide the weight space into finitely many regions within which the available profiles, and hence the equations matching cell masses to the returns, are unchanged. Checking the feasibility of these equations characterizes the sharp identified set. The intersection of the availability intervals of profiles that the data require gives an interval outer bound; the sharp set itself need not be a single interval.

B.1 Setting and notation

I use the model of Section 2, in which the single parameter is the weight $w > 0$ of the second dimension in $d_w(x,y)^2 = (x_1 - y_1)^2 + w\,(x_2 - y_2)^2$. For this appendix I take the policy space to be the nondegenerate box $\mathcal{P} = [a_1, b_1] \times [a_2, b_2]$, with $a_\ell < b_\ell$. There are finitely many states, municipalities, types, elections and candidates, and candidate positions within each election are distinct. Under Assumption 1a, each state–type group has an otherwise unrestricted ideal-point density $f(x \mid s,d)$, common across municipalities in the state.

The attraction region of candidate j in election e at weight w is

$$\mathcal{A}_{j;w}(e) = \{x \in \mathcal{P} : d_w(x, Y_{je}) < d_w(x, Y_{ke}) \ \forall k \neq j\}.$$

The strict inequalities omit ties, which lie on finitely many lines and have probability zero under the density assumption. As in the previous section, let λ denote Lebesgue measure on $\mathbb{R}^2$; a cell is available only if it has

positive area, since a zero-area cell cannot carry positive probability.

Superimposing the attraction regions of every $e \in \mathcal{E}(s)$ yields the minimal relevant partition $\mathbb{X}_w(s)$ of Section 2.3, with its state dependence made explicit. A *voting profile*

$$\mathbf{j} = (j_e)_{e \in \mathcal{E}'}, \qquad j_e \in \mathcal{J}_e, \qquad \mathcal{E}' \subseteq \mathcal{E}(s),$$

records the candidate chosen in each election in $\mathcal{E}'$ and defines the intersection

$$\mathcal{X}_w(\mathbf{j}) \;=\; \bigcap_{e \in \mathcal{E}'} \mathcal{A}_{j_e;w}(e). \tag{B1}$$

Taking $\mathcal{E}' = \mathcal{E}(s)$ gives the cells of $\mathbb{X}_w(s)$; a proper subset gives a *partial* profile, whose region is a union of full-profile cells up to tie boundaries. The geometric argument below applies to full and partial profiles alike. The set of available full profiles in state s is

$$\mathcal{C}_s(w) = \{\mathbf{j} = (j_e)_{e \in \mathcal{E}(s)} : \lambda(\mathcal{X}_w(\mathbf{j})) > 0\}.$$

Two weights are observationally equivalent whenever they induce the same $\mathcal{C}_s(w)$ in every state, even if their cell boundaries differ.

The cell as a system of linear constraints. By (B1), a cell is an intersection of half-planes with $\mathcal{P}$: one half-plane per election $e \in \mathcal{E}'$ and rival $k \neq j_e$, for a total of $R = \sum_{e \in \mathcal{E}'}(|\mathcal{J}_e| - 1)$ comparisons. The set of ideal points closer to an own position p than to a rival p' satisfies

$$2(p'_1 - p_1)x_1 + 2w(p'_2 - p_2)x_2 < (p'^2_1 - p^2_1) + w(p'^2_2 - p^2_2).$$

Indexing the comparisons by $r = 1, \ldots, R$, each takes the form

$$\alpha_r x_1 + w\,\beta_r x_2 < \gamma_r + w\,\delta_r, \qquad r = 1, \ldots, R, \tag{B2}$$

with fixed constants $\alpha_r = 2(p'_1 - p_1)$, $\beta_r = 2(p'_2 - p_2)$, $\gamma_r = p'^2_1 - p^2_1$ and $\delta_r = p'^2_2 - p^2_2$. For the comparison arising from election e and rival k, the own position is $p = Y_{j_e e}$ and the rival is $p' = Y_{ke}$.

The *availability set* of a profile $\mathbf{j}$ is

$$\mathcal{W}(\mathbf{j}) = \{w > 0 : \lambda(\mathcal{X}_w(\mathbf{j})) > 0\} = \{w > 0 : \exists x \in \mathrm{int}(\mathcal{P}) \text{ satisfying (B2) for all } r\}. \tag{B3}$$

The second equality holds because a solution in the interior of $\mathcal{P}$ has a neighborhood satisfying the strict comparisons, while a positive-area cell must contain such a solution.

From rotations to affine thresholds. In (B2), the term $w\beta_r x_2$ makes the boundary's orientation depend on w. Apply the change of variables $(z_1, z_2) = (x_1, wx_2)$, which is a bijection for each $w > 0$. The system defining availability becomes

$$\begin{aligned} \alpha_r z_1 + \beta_r z_2 &< \gamma_r + w\,\delta_r, \qquad r = 1, \ldots, R, \\ a_1 < z_1 &< b_1, \qquad wa_2 < z_2 < wb_2. \end{aligned} \tag{B4}$$

The coefficients on z are now fixed, and w enters only through affine thresholds. This is the geometric feature used by Lieber, Tebaldi and Torgovitsky (2024) for preference coefficients that enter utility linearly. Here the transformed support also depends on w, but its constraints are affine in (z_1, z_2, w) as well. Thus the

substitution preserves the linear geometry needed to characterize profile availability.[42]

B.2 Set identification of weights

Lemma 1. *For every voting profile* $\mathbf{j}$*, the availability set* $\mathcal{W}(\mathbf{j})$ *is an open interval in* $(0,\infty)$*, possibly empty or unbounded.*

Proof. Define the augmented feasible set

$$\Omega(\mathbf{j}) = \left\{ (z_1, z_2, w) \in \mathbb{R}^3 : \begin{array}{ll} w > 0, & a_1 < z_1 < b_1, \quad wa_2 < z_2 < wb_2, \\ \alpha_r z_1 + \beta_r z_2 - w\delta_r < \gamma_r, & r = 1, \ldots, R \end{array} \right\}. \tag{B5}$$

Every inequality is strict and affine in the joint variable (z_1, z_2, w), so $\Omega(\mathbf{j})$ is open and convex. By (B4), its projection onto the w coordinate is exactly $\mathcal{W}(\mathbf{j})$. A coordinate projection preserves openness and convexity, and an open convex subset of $\mathbb{R}$ is an open interval, possibly empty or unbounded. □

Proposition 3. *Let* $\mathcal{C}^+$ *be any collection of full or partial profiles whose regions must carry positive mass under every model consistent with the data. Then*

$$\mathcal{W}(\mathcal{C}^+) \equiv \bigcap_{\mathbf{j} \in \mathcal{C}^+} \mathcal{W}(\mathbf{j})$$

is an open interval in $(0,\infty)$*, possibly empty or unbounded. If the data are generated at the true weight* w_0*, then*

$$w_0 \in \mathcal{W}^* \subseteq \mathcal{W}(\mathcal{C}^+),$$

where $\mathcal{W}^*$ *is the sharp identified set of weights.*

Proof. There are finitely many profiles, and a finite intersection of open intervals is an open interval. Any weight consistent with the data must allow every profile in $\mathcal{C}^+$ to carry positive mass. Absolute continuity requires each of these regions to have positive area, so the weight belongs to every $\mathcal{W}(\mathbf{j})$ in the intersection. The true weight is consistent with the data and hence belongs to this intersection as well. □

Outer versus sharp set. $\mathcal{W}(\mathcal{C}^+)$ is an outer bound: availability of the required profiles does not guarantee that all observed shares can be matched simultaneously. Conversely, let $\mathcal{C}^*$ collect the full profiles carrying positive mass for at least one state–type group under the true model, and let $I^* = \bigcap_{\mathbf{j} \in \mathcal{C}^*} \mathcal{W}(\mathbf{j})$. The true profile masses reproduce the returns at every $w \in I^*$, because every profile they use remains available. Consequently,

$$I^* \subseteq \mathcal{W}^* \subseteq \mathcal{W}(\mathcal{C}^+).$$

The outer bound is sharp if $\mathcal{W}(\mathcal{C}^+) \subseteq I^*$. In general, however, the data can be rationalized by different allocations of mass to different collections of profiles. Sharp identification therefore requires checking the full system of equations matching masses to the returns.

Let Φ_w denote the nonnegative conditional mass functions on the positive-area cells of $\mathbb{X}_w(s)$, with $\sum_{\mathcal{X} \in \mathbb{X}_w(s)} \phi(\mathcal{X} \mid s, d) = 1$ for each (s, d) in the support of (S_i, D_i). Write

$$\Phi_w^* = \{\phi \in \Phi_w : \phi \text{ satisfies (B8)}\},$$

[42] The transformation is applied separately at each candidate weight. It maps a density $f(x \mid s, d)$ on $\mathcal{P}$ to $g_w(z \mid s, d) = w^{-1} f(z_1, z_2/w \mid s, d)$ on $[a_1, b_1] \times [wa_2, wb_2]$, preserving its common distribution across municipalities within each state–type group. No restriction equating these transformed densities across candidate weights is imposed.

where (B8) below is the cell-mass representation of the observed returns.

Proposition 4. *The sharp identified set of weights is*

$$\mathcal{W}^* = \{w > 0 : \Phi_w^* \neq \varnothing\}. \tag{B6}$$

Feasibility is unchanged whenever the sets $\mathcal{C}_s(w)$ are unchanged in every state. The weight space can be divided into finitely many open intervals and singleton boundary points on each of which feasibility is constant. Moreover, $\mathcal{W}^*$ is a finite union of open intervals.

Proof. Any density consistent with the returns induces nonnegative cell masses satisfying (B8), so membership in (B6) is necessary. Conversely, given $\phi \in \Phi_w^*$, define

$$f_w(x \mid s, d) = \sum_{\mathcal{X} \in \mathbb{X}_w(s)} \frac{\phi(\mathcal{X} \mid s, d)}{\lambda(\mathcal{X})} \mathbf{1}\{x \in \mathcal{X}\}.$$

This is a density, common across municipalities within each state–type group, that assigns the prescribed mass to each cell and reproduces the returns. This proves sufficiency. Since the equations depend on the profile labels and their masses, two weights permitting the same profiles in every state have the same feasibility problem.

By Lemma 1, each full profile has at most two finite availability endpoints. There are finitely many full profiles across all states, so the set B of their finite positive endpoints is finite. On each component of $(0, \infty) \setminus B$, all sets $\mathcal{C}_s(w)$ are constant. These components and the singleton sets $\{b\}$, $b \in B$, give the claimed finite division of the weight space.

Finally, take any $w \in \mathcal{W}^*$ and a feasible mass function. Each profile assigned positive mass remains available on an open interval containing w. The intersection of these finitely many intervals is a neighborhood of w on which the same masses remain feasible. Hence $\mathcal{W}^*$ is open. Together with the finite division just established, this implies that it is a finite union of open intervals. □

Note that since $\mathcal{W}^*$ is open, every admissible weight has an admissible neighborhood. With finitely many elections and unrestricted densities, the weight is therefore never point-identified.

B.3 Target parameter bounds with weights

This section combines the weights and cell masses consistent with the data to derive sharp bounds on the target parameters.[43]

The observational-equivalence conditions for densities and cell masses can now be written with the

[43] The closest argument in the literature is Merlo and De Paula (2017), who prove nonparametric point identification of weights and distributions of ideal points in a spatial voting model. Their analysis relies on a large support assumption for candidate positions. In particular, to establish joint identification of ideal points and weights, they rely on the results of Ichimura and Thompson (1998), who prove nonparametric identification in a binary discrete choice model using the Cramér–Wold device: there exists a two-candidate configuration such that the affine hyperplane separating the two candidates delivers different election outcomes under two distinct preference distributions. Building on these results, they extend them to 1) include multiple candidate positions, and 2) introduce weights. Both 1) and 2) are based on the large support assumption, i.e. that one could sample arbitrary candidate configurations from $\mathbb{R}^2$. This assumption is infeasible in political-economy environments, where the set of candidate positions is typically finite and confined to a bounded subset of $\mathbb{R}^2$. Hence, the setup of this paper accommodates the fact that there are just finitely many observed configurations in a bounded policy space.

weight made explicit. The predicted vote share of candidate j in municipality m is

$$\sigma_j(m; f, w) = \sum_d q_{m,d} \int_{\mathcal{A}_{j;w}(e(m))} f(x \mid s(m), d)\, dx.$$

It matches the data if

$$\sigma_j(m; f, w) = \mathcal{s}_j(m) \quad \text{for all } j \in \mathcal{J}_{e(m)} \text{ and all municipalities } m. \tag{B7}$$

For the equivalent finite-dimensional problem, let $\mathbb{X}_{j;w}(e)$ denote the cells of the relevant state's MRP contained in $\mathcal{A}_{j;w}(e)$. The corresponding condition is

$$\sum_d q_{m,d} \sum_{\mathcal{X} \in \mathbb{X}_{j;w}(e(m))} \phi(\mathcal{X} \mid s(m), d) = \mathcal{s}_j(m) \quad \text{for all } j \in \mathcal{J}_{e(m)} \text{ and all municipalities } m. \tag{B8}$$

First, obtain $\mathcal{W}^*$ using (B6). If a candidate weight space $\mathcal{W} \subseteq (0, \infty)$ is imposed, restrict this set to $\mathcal{W}$. Second, for each admissible weight obtain $\Theta^*(w)$ by applying Propositions 1–2 to $\mathbb{X}_w(s)$, refined by any counterfactual elections or policy-space regions entering the target parameter. The sharp identified set of the target parameter is

$$\Theta^* = \bigcup_{w \in \mathcal{W}^*} \Theta^*(w).$$

For the linear scalar target parameters considered in the application, each $\Theta^*(w)$ is an interval with endpoints $t^*_{lb}(w)$ and $t^*_{ub}(w)$. The sharp lower and upper bounds allowing for the unknown weight are therefore

$$\inf_{w \in \mathcal{W}^*} t^*_{lb}(w) \qquad \text{and} \qquad \sup_{w \in \mathcal{W}^*} t^*_{ub}(w).$$

These are bounds on Θ^*; its union representation does not by itself imply that every value between them is attainable.

To compute target bounds from a finite set of representative weights, the division of the weight space must also account for the target parameter. For fixed counterfactual platforms, add their elections before determining profile availability. For targets defined by thresholds on individual policy coordinates, include the corresponding region boundaries; after the same substitution these also give affine inequalities in (z_1, z_2, w). When the target is linear in the resulting masses with coefficients that do not depend directly on w, both the feasible mass set and the objective are constant within each resulting region of the weight space. The bounds can then be computed by solving the relevant linear programs at one representative weight per region, including boundary points. The observed voting profiles alone do not suffice for this reduction when the target requires a finer partition. Appendix C describes the estimation procedure in the application.

C Model Estimation

This section translates the empirical restrictions of the spatial voting model of Section 2 into finite-dimensional linear programs. Appendix Section C.1 introduces the objects common to all programs, Appendix Section C.2 covers the estimation of the weight, and Appendix Section C.3 covers the estimation of the bounds on the target parameters given that weight.

C.1 Setup and notation

Estimation is carried out separately for each region (East, West) and year, within which the weight is fixed across all voters, while ideal-point distributions differ across states and types. Let $\mathcal{S}$ be the states of

the region and $\mathcal{E} = \bigcup_{s \in \mathcal{S}} \mathcal{E}(s)$ its elections. As discussed in Appendix D, identification of the weight requires substantial variation, which pooling the elections of a region delivers.

Elections, municipalities and data. Each state s is divided into electoral districts, each holding one election $e \in \mathcal{E}(s)$. Municipalities are indexed by $m \in \mathcal{M}$; $\mathcal{M}(s) \subset \mathcal{M}$ collects the municipalities of state s, $s(m)$ denotes the state of m, and $e(m)$ its election. The candidates are parties: $\mathcal{J}_e = \mathcal{J}$ is the same set of six parties in every election, with positions Y_{je} varying across elections. Demographic types $d \in \mathcal{D}$ are the 24 age×gender×education bins. For each municipality the data give the number of valid votes cast N_m, the turnout-adjusted share $q_{m,d}$ of type d among the voters who cast them ($\sum_d q_{m,d} = 1$), and the observed vote shares $s_j(m)$, $j \in \mathcal{J}$.[44] Depending on the exercise, estimation fits the shares of all parties or of the AfD alone; to keep notation light, the exposition below treats the AfD-only case. Throughout, the population is the voting population: since the model of Section 2 does not include a turnout decision, I reweight the demographic composition of each municipality by administrative data on turnout rates by state, age and gender (footnote 44), so that $f(x \mid s, d)$ and the cell masses below describe the ideal points of type-d voters in state s who cast a vote, and all counts and shares refer to these voters.

Weight and mesh. The weight on the Left–Right dimension is normalized to one, so the weight of Section 2 reduces to $w = w_{AE} \in \mathbb{R}_+$, the only free parameter, which reshapes the attraction regions $\mathcal{A}_j(e)$. Overlaying the attraction regions of all elections $e \in \mathcal{E}$ in the states of a region yields the Minimal Relevant Partition $\mathbb{X}_w$. I restrict attention to the AfD attraction region and reduce the partition to $n = 1{,}000$ cells by merging the smallest polygons, or slivers, into the neighboring cell with which they share the longest boundary. The resulting mesh, $\{\mathcal{X}_1, \ldots, \mathcal{X}_n\}$, is indexed by ℓ and rebuilt at every w_{AE}. Merging serves three purposes. First, it keeps noise in candidate positions from propagating into the estimation: slivers arise from small displacements of the boundaries between attraction regions and have negligible measure. Second, it fixes the number of cells across weights, so that differences in fit across weights reflect the *composition* of the cells, not their *number*. Third, it keeps the linear programs tractable, which matters especially for the optimization problems of the target parameters in the next step. The merging affect less than 90% of policy space area in all region-years.

For municipality m and cell ℓ, let $a_{m,\ell}(w) \in \{0, 1\}$ indicate that cell $\mathcal{X}_\ell$ lies in the AfD attraction region in election $e(m)$, and $\boldsymbol{a}_m(w) \in \{0, 1\}^n$ the corresponding vector.

Unknowns and fit. The unknowns are the cell masses $\phi_{s,d,\ell} = \phi(\mathcal{X}_\ell \mid s, d)$, collected in

$$\boldsymbol{\phi}_{s,d} = (\phi_{s,d,1}, \ldots, \phi_{s,d,n})^\top \in \mathbb{R}^n_+, \qquad \mathbf{1}^\top \boldsymbol{\phi}_{s,d} = 1, \qquad s \in \mathcal{S},\ d \in \mathcal{D},$$

so each state has its own ideal-point distribution over the shared mesh; $\boldsymbol{\phi}_{s,d}$ is the distribution of ideal points among the type-d voters of state s who cast a vote. $\boldsymbol{\phi}$ denotes the collection $\{\boldsymbol{\phi}_{s,d}\}$ and Φ the set of such collections. Given $\boldsymbol{\phi}$ and w_{AE}, the predicted AfD share in municipality m is

$$\sigma_{\text{AfD}}(m; \boldsymbol{\phi}, w) = \sum_{d \in \mathcal{D}} q_{m,d}\, \boldsymbol{a}_m(w)^\top \boldsymbol{\phi}_{s(m),d}. \tag{C1}$$

[44] Turnout adjustment. For municipality m the census gives the number of eligible voters N^0_m and the shares $q^0_{m,d}$ of the demographic types $d \in \mathcal{D}$ among them, $\sum_d q^0_{m,d} = 1$. The Representative Electoral Statistics report, for each state s and election year, turnout rates by gender and ten age brackets; I collapse the brackets to the six age groups of the model by averaging the bracket rates and apply the same rate to the education and non-education halves of a type, since the statistic has no education split, which gives $r_{s,d}$. The turnout-adjusted share of type d in municipality m is $q_{m,d} = q^0_{m,d}\, r_{s(m),d} \big/ \sum_{d'} q^0_{m,d'}\, r_{s(m),d'}$, i.e. the composition of those eligible voters in m who actually vote, with $\sum_d q_{m,d} = 1$; $N_m\, q_{m,d}$ is then the number of type-d voters in m.

In finite samples the exact match-the-data restriction (5) generically has no solution. I therefore introduce slacks $u_m \geq 0$ and impose

$$-u_m \;\leq\; \sigma_{\text{AfD}}(m;\boldsymbol{\phi},w) - s_{\text{AfD}}(m) \;\leq\; u_m \qquad \forall\, m \in \mathcal{M}. \tag{C2}$$

Let $\mathcal{K}_0(w)$ be the set of pairs $(\boldsymbol{\phi},\boldsymbol{u})$ with $\boldsymbol{\phi} \in \Phi$ and $\boldsymbol{u} \geq 0$ satisfying (C2). All programs below are linear in $(\boldsymbol{\phi},\boldsymbol{u})$ and measure fit by the mean absolute error $\frac{1}{|\mathcal{M}|}\sum_m u_m$.

All-party estimation. Some exercises fit the shares of all six parties. The mesh is then the minimal relevant partition of the whole policy space rather than of the AfD attraction region, with $n = 2,000$ cells; the indicator $a_{m,j,\ell}(w)$ records whether cell $\mathcal{X}_\ell$ lies in the attraction region of party j in election $e(m)$; and (C1)–(C2) are imposed for every $j \in \mathcal{J}$, with the mean absolute error taken over the $|\mathcal{M}|\,|\mathcal{J}|$ fitted shares.

Full partition. I repeat the estimation of the weights and of the bounds on the full Minimal Relevant Partition, without merging. The estimated weights are close to those from the merged mesh across regions and years. As expected, within a region and year, the fit on the full partition improves with the number of cells a weight induces. The fit profiles of the full and merged meshes, however, track each other closely, so the ranking of weights reflects the composition of the cells rather than their number. This mirrors the simulations in Appendix D: (i) across weights, the fit improves with the number of cells, yet the loss profiles with fixed and with all cells track each other closely, so the link reflects more the weight than the count; (ii) the number of cells matters close to the optimum, where near-ties in fit are common and the lowest-loss weight induces slightly more cells than its neighbours; (iii) fixing the number of cells removes this edge without making the estimated weight less accurate. Merging therefore ensures that differences in fit across weights reflect the composition of the cells alone, especially near the optimum, at a much lower computational cost.

The bounds on the target parameters computed on the full partition are sharp; they are slightly narrower than those from the merged mesh in all years except 2013 (Figure J13), at a much higher computational cost. The paper therefore reports the merged-mesh estimates throughout.

C.2 Weights

For each candidate weight, the best attainable fit is the mean absolute municipal AfD-share error,

$$\bar{Q}(w_{AE}) = \frac{1}{|\mathcal{M}|} \min_{(\boldsymbol{\phi},\boldsymbol{u})\in\mathcal{K}_0(w)} \mathbf{1}^\top \boldsymbol{u},$$

where, at each weight, a linear program chooses the masses that fit the data best, so $\bar{Q}$ records only the resulting error and is a function of w_{AE} alone. This program is solved for every weight on the grid $\mathcal{W} = \{0.2, 0.3, \ldots, 0.9, 1.0\} \cup \{1/0.9, \ldots, 1/0.2\}$, symmetric around 1 in log scale, and the best-fitting weight is

$$\hat{w}_{AE} = \arg\min_{w_{AE}\in\mathcal{W}} \bar{Q}(w_{AE}).$$

Values $\hat{w}_{AE} < 1$ mean voters tolerate larger distances on the anti-elite dimension (the Left–Right dimension weighs more); $\hat{w}_{AE} > 1$ the reverse.

A grid search is sufficient here for two reasons. The weight is scalar, so the search is exhaustive. And by Proposition 4 the population criterion depends on w_{AE} only through the set of available cells, which changes at finitely many breakpoints, so the criterion is constant on finitely many intervals and a grid finer than the shortest of them evaluates it exactly.[45] The reported sets are unions of grid intervals, resolved to the grid

[45] This replaces the enumeration algorithm that Lieber, Tebaldi and Torgovitsky (2024) require for a multidimensional preference vector.

step.

The weight is set-identified: in the population, $\bar{Q}$ is zero on every weight in the identified set and strictly positive outside it. Because the data are noisy, the sample criterion has a unique minimizer, but which weight within the set of near-minimizers it picks out is a matter of sampling noise; $\hat{w}_{AE}$ serves as a representative value. Set estimates are reported in two forms: the grid weights whose criterion lies within 1% of the minimum (Figure J1), and a confidence set whose tolerance is calibrated from subsample re-estimations of the profiled criterion (Figure J2).

C.3 Bounds on target parameters

This section describes the bounds used for the motivations exercise of Section 6.2.2; the other exercises follow the same logic. The weight is fixed at $\hat{w}_{AE}$ from Section C.2; the mesh and the indicators $\boldsymbol{a}_m$ are those of Section C.1 evaluated at $\hat{w}_{AE}$.

Target parameter. Let $\bar{Y} = (\bar{Y}_{LR}, \bar{Y}_{AE})$ be the AfD position (national median value in the year) on the Left–Right and Anti-Elite dimensions, as in Section 6.2.2, and define the reference points

$$C^{AE} = (\bar{Y}_{LR} - 1,\ \bar{Y}_{AE}), \qquad C^{Right} = (\bar{Y}_{LR},\ \bar{Y}_{AE} - 1).$$

The AfD attraction region is split into two target areas by which reference point is nearer under $w = \hat{w}_{AE}$: the *anti-elite* area, nearer C^{AE}, and the *right* area, nearer C^{Right}. Let $\mathcal{C}^{AE} \subset \{1, \dots, n\}$ collect the cells that intersect the anti-elite area and $\mathcal{C}^{Right}$ those that intersect the right area; a cell cut by the boundary between the two areas belongs to both sets. The mass of AfD voters in the cells of $\mathcal{C}$ is

$$W(\boldsymbol{\phi}; \mathcal{C}) = \sum_{s \in \mathcal{S}} \sum_{d \in \mathcal{D}} \big(\boldsymbol{b}_{s,d}(\mathcal{C})\big)^{\top} \boldsymbol{\phi}_{s,d}, \qquad b_{s,d,\ell}(\mathcal{C}) = \mathbf{1}\{\ell \in \mathcal{C}\} \sum_{m \in \mathcal{M}(s)} N_m\, q_{m,d}\, a_{m,\ell}, \tag{C3}$$

so that $(\boldsymbol{b}_{s,d}(\mathcal{C}))^{\top}\boldsymbol{\phi}_{s,d}$ counts the state-s, type-d voters who vote AfD from an ideal point in a cell of $\mathcal{C}$. The target parameter is the Right share of AfD support,

$$\theta_{Right}(\boldsymbol{\phi}) = \frac{W(\boldsymbol{\phi}; \mathcal{C}^{Right})}{W(\boldsymbol{\phi}; \mathcal{C}^{Right}) + W(\boldsymbol{\phi}; \mathcal{C}^{AE})}, \qquad \theta_{AE}(\boldsymbol{\phi}) = 1 - \theta_{Right}(\boldsymbol{\phi}),$$

a ratio of two linear functions of $\boldsymbol{\phi}$.

Pooling the states of a region onto one mesh is informative only if their distributions are related, and the coupling restriction imposes this loosely. The idea is that young voters in one state of a region should not hold ideological positions wildly different from young voters in another state of the same region, and likewise for the other broad demographic groups. States are therefore compared only along six coarse marginals, $g \in \{\text{with higher education},\ \text{without higher education},\ \text{young},\ \text{old},\ \text{male},\ \text{female}\}$, each a union of types $d \in g$, and only on a coarse partition of the policy space into blocks B on a 0.5 grid. Let $T_{s,d} = \sum_{m \in \mathcal{M}(s)} N_m\, q_{m,d}$ be the number of type-d voters in state s, and let $\rho_s(g, B)$ be the share of marginal g in state s with ideal points in block B, i.e. the mass $\sum_{d \in g} T_{s,d} \sum_{\ell \in B} \phi_{s,d,\ell}$ divided by $T_{s,g} = \sum_{d \in g} T_{s,d}$. Let $\bar{\rho}(g, B)$ be the population-weighted regional mean of these shares, an auxiliary variable of the program. Every state's share must lie within a factor $1 + \tau_{\text{state}}$ of the regional mean,

$$\frac{\bar{\rho}(g, B)}{1 + \tau_{\text{state}}} \ \le\ \rho_s(g, B) \ \le\ (1 + \tau_{\text{state}})\, \bar{\rho}(g, B) \qquad \forall\, s \in \mathcal{S},\ g,\ B, \tag{C4}$$

which is linear in $(\boldsymbol{\phi}, \bar{\rho})$. The benchmark uses $\tau_{\text{state}} = 2$: each state's block share may be up to three times the regional mean or as low as a third of it, so the coarse moments between two states of the same region may

differ by up to a factor of nine. The coupling is loose: it operates only at the level of coarse marginals and large blocks, and allows substantial differences across states within each marginal–block cell, while keeping the problem computationally tractable.

Bounds. Let $\mathcal{K}$ be the set of $(\boldsymbol{\phi}, \boldsymbol{u}) \in \mathcal{K}_0(\hat{w}_{AE})$ that satisfy (C4) for some $\bar{\rho}$. I first solve the best-fit program $\hat{Q} = \min_{(\boldsymbol{\phi},\boldsymbol{u})\in\mathcal{K}} \frac{1}{|\mathcal{M}|}\mathbf{1}^\top \boldsymbol{u}$, the smallest mean absolute error attainable under the coupling restriction. The approximately feasible set then allows the mean absolute error to exceed its minimum by at most $\epsilon_{\text{abs}} = 10^{-5}$, i.e. $\frac{1}{|\mathcal{M}|}\mathbf{1}^\top \boldsymbol{u} \le \hat{Q} + \epsilon_{\text{abs}}$, and the bounds are

$$\underline{\theta}_{Right} = \min\ \theta_{Right}(\boldsymbol{\phi}), \qquad \overline{\theta}_{Right} = \max\ \theta_{Right}(\boldsymbol{\phi}), \qquad \text{s.t.}\ \ (\boldsymbol{\phi}, \boldsymbol{u}) \in \mathcal{K},\ \ \frac{1}{|\mathcal{M}|}\mathbf{1}^\top \boldsymbol{u} \le \hat{Q} + \epsilon_{\text{abs}}.$$

These are linear-fractional programs, which I solve by bisection on the value of the target parameter: for a candidate value t, some feasible $\boldsymbol{\phi}$ attains $\theta_{Right}(\boldsymbol{\phi}) \le t$ if and only if the linear program with the additional constraint $W(\boldsymbol{\phi}; \mathcal{C}^{Right}) \le t\left[W(\boldsymbol{\phi}; \mathcal{C}^{Right}) + W(\boldsymbol{\phi}; \mathcal{C}^{AE})\right]$ is feasible, so $\underline{\theta}_{Right}$ is the smallest such t, found by bisection to the desired precision; $\overline{\theta}_{Right}$ is obtained symmetrically with the inequality reversed.

C.4 Model fit

Table C1 presents summary statistics on model fit at the municipality level. In every election year, the median residual is exactly 0 by construction. The standard deviation ranges from 0.016 in 2013 to just above 0.03 in 2017 and 2021, and 0.048 in 2025. This increase mirrors an increase in AfD vote shares. The correlation between predicted and observed vote shares is above 0.9 in the later years, and about 0.6 in 2013. Overall, these results suggest a reasonable model fit, especially in the later years. I also find that accounting for demographic heterogeneity improves out-of-sample prediction: holding out 10% of municipalities within each district, the 24-group model outperforms restricted specifications that merge groups by education, age, or gender. The loss from merging education or age groups is statistically significant ($p < 0.01$ in both cases), while the effect of merging gender groups is smaller but still significant at the 5% level ($p = 0.04$).

Table C1. Distribution of Residuals Across Municipalities

	2013	2017	2021	2025
Mean	0.001	0.001	0.001	-0.001
SD	0.016	0.033	0.032	0.048
P10	-0.017	-0.036	-0.034	-0.062
Median	0.000	0.000	0.000	0.000
P90	0.020	0.041	0.040	0.058

Note: The unit of observation is the municipality. Residual = observed AfD vote share − model-predicted share.

The correlations of predicted and observed municipal AfD shares with municipalities' demographic characteristics are almost identical (Figure C3). To sum up, both the model fit and the observed heterogeneity in the data suggest that the model performs well in explaining observed AfD voting patterns.

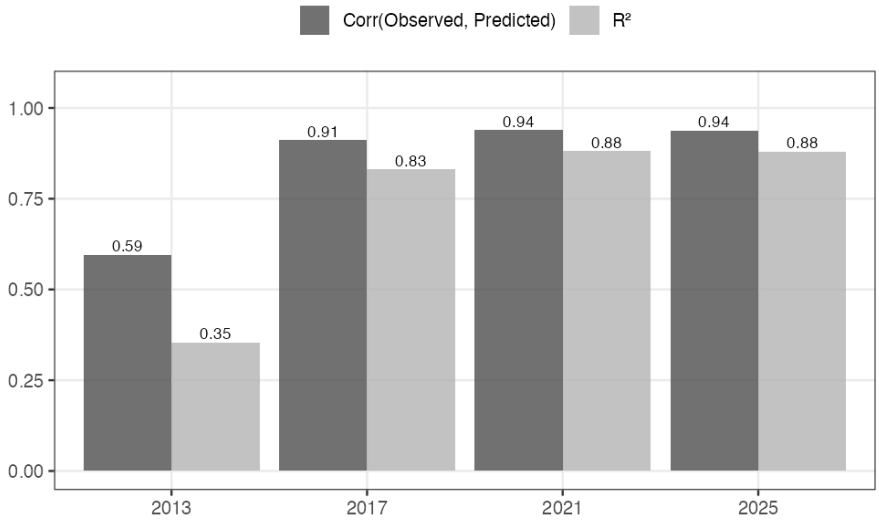


Figure C1. Correspondence between predicted and observed AfD vote shares across municipalities.

Notes. Bars plot the values for correlation between predicted and observed vote shares, and R^2 from a regression of one on the other.

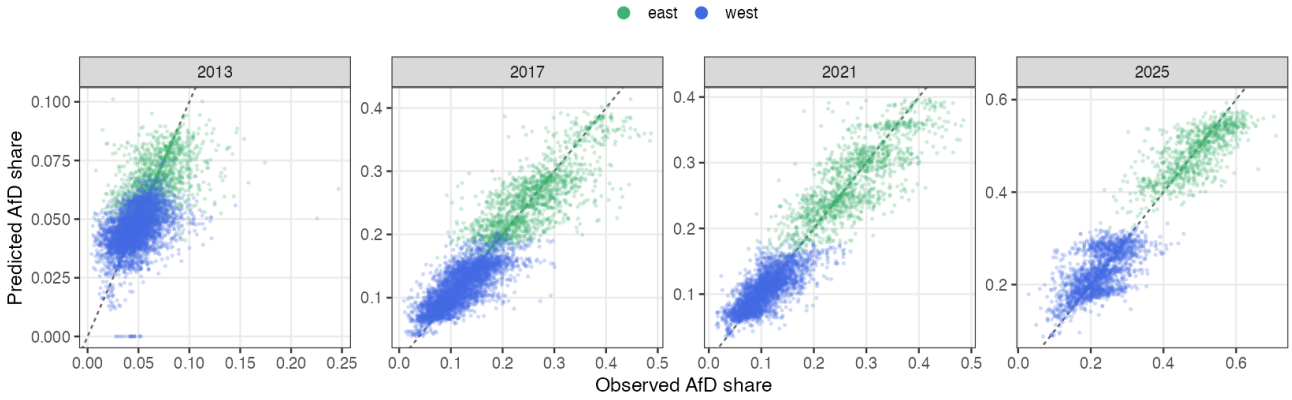


Figure C2. Correspondence between predicted and observed AfD vote shares in Eastern and Western Germany.

Notes. Each dot is a municipality, with the respective values for predicted and observed vote shares.

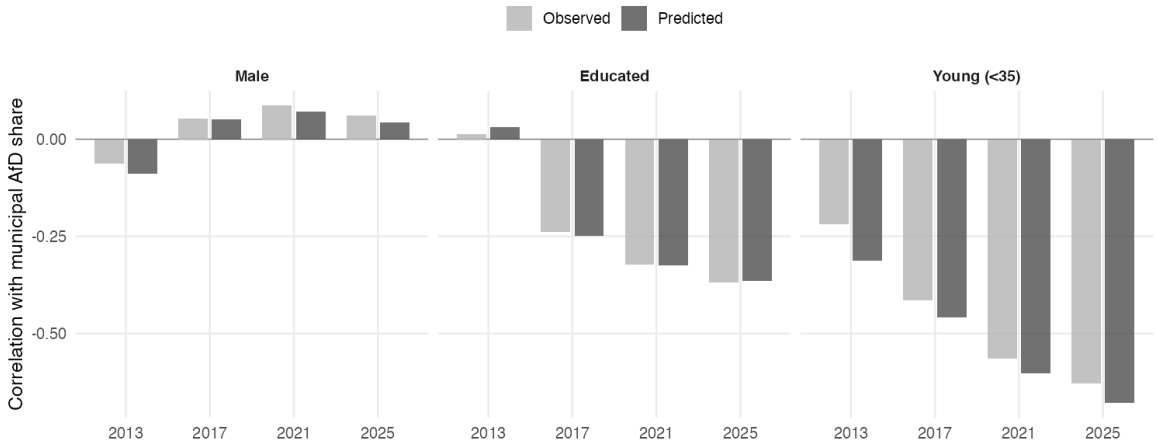


Figure C3. Correlation between predicted and observed AfD vote shares with demographic characteristics of municipalities.

Notes. Each bar is an estimate of a correlation between a predicted (observed) AfD vote share and a given municipality demographic characteristic.

D Weights estimation: Simulation

This section provides a simulation of weight estimation using the approach developed and adapted in this paper; and compares it to the one used in Merlo and De Paula (2017).

I simulate a population of 10,000 voters who vote in 24 elections.[46] Policy space is two-dimensional, and is restricted by a box of $[-2.5, 2.5]^2$. Each election has six candidates whose positions are drawn at random from a uniform distribution over the policy space; voters' ideal points are drawn from a mixture of two dominant clusters and a diffuse background.[47] Voters choose the candidate closest to their ideal point under the weighted Euclidean metric with true weight $w = 1$ as in (2); the econometrician observes only the resulting candidate vote shares in each election.

For each candidate value of w on a grid, the implied superimposed Voronoi tessellation determines the set of feasible cells (joint vote patterns). The question then is whether voter mass can be distributed across these cells so as to reproduce the observed shares exactly, subject to masses being non-negative and summing to one. The set of weights for which an exact fit exists is the identified set; in this simulation, the exact-fit set is itself an interval.

Figure D1 plots the fit against w. The loss is exactly zero on the interval $[0.80, 1.18]$, which contains the true value $w = 1$, and rises outside it. The cell geometry alone thus set-identifies the weight without any assumption on the distribution of ideal points.

The dotted lines mark the estimates obtained from the sieve estimator at polynomial orders $J_t = 2, 3$, which impose a flexible parametric form on the ideal-point density following the approach in (Merlo and De Paula, 2017). Because a low-order polynomial correction of a Gaussian base approximates the distribution of voters' ideal points not perfectly, the resulting weight estimates are biased, at 1.20 and 1.27 for $J_t = 2, 3$ respectively.

Number of elections. The identified set shrinks as the number of elections grows. Because each additional election contributes a further set of cells that the observed shares may or may not be able to populate, more elections impose more restrictions on which weights admit an exact fit. Figure D2 traces the identified set for w against the number of elections, holding the data-generating process fixed: the set narrows monotonically around the true value, from $[0.58, 1.18]$ with ten elections to $[0.93, 1.05]$ with one hundred. Note that while this simulation illustrates that the size of the identified set shrinks with the number of elections, its size at any given number of elections depends on the positions of the candidates and the distribution of preferences.

Noisy perceptions of candidate positions. In practice, however, the data are noisy. I next simulate the behavior of both estimators when each voter perceives every candidate's position with an idiosyncratic error, distributed $\mathcal{N}(\mathbf{0}, 0.1^2 I_2)$, while the estimation procedures remain unchanged.[48] Under noise an exact fit is no longer attainable at any weight. The question, however, is where the minimum is located and how large the basin delivering a near-identical fit is. The region of near-minimal losses becomes smaller: the loss now develops a strict interior minimum, and the near-feasible set (weights whose fit is within one percentage point of share mass of the best attainable) shrinks from the noise-free interval $[0.80, 1.18]$ to $[0.98, 1.06]$ (Figure D3). Intuitively, the smoothed shares are best approximated by the cell structure of the true weight, so mild

[46] As discussed in the main text, the weight is identified from variation in choice environments. Since I assume that the weight is common across socio-demographic groups, simulating the same voters across different choice environments is without loss of generality. The number of elections is close to the median number of electoral districts per area state in western Germany in my estimation sample.

[47] The two clusters intend to represent a more *left-leaning* and a more *right-leaning* group of voters. The choice of this distribution does not materially affect the results, as long as it is non-degenerate, that is, as long as it places mass across a broad region of the policy space.

[48] A related case is one in which voters observe positions perfectly but candidates' positions are measured with error.

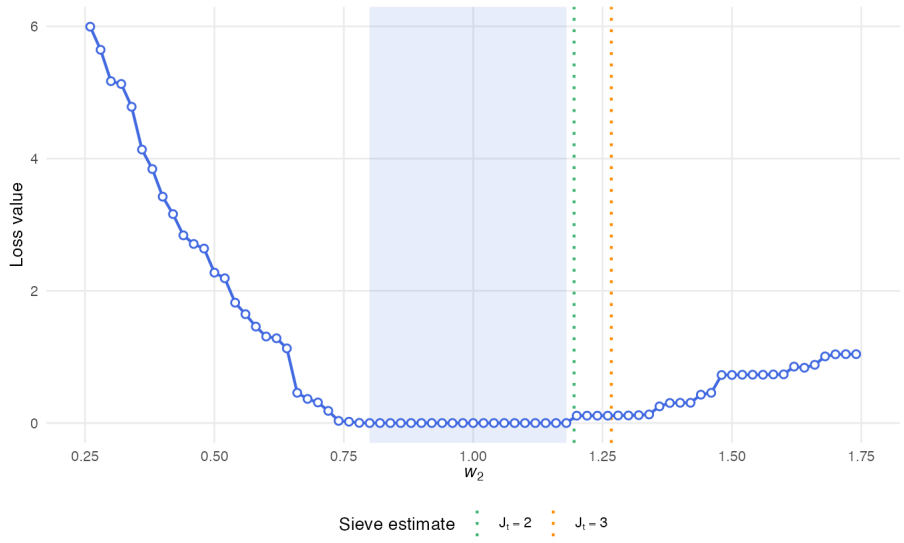


Figure D1. Set identification of the weight via the method used in this paper, and sieve estimates

Notes: Monte Carlo with 10,000 voters, true $w = 1$, twenty-four six-candidate elections; ideal points follow a two-cluster mixture with a diffuse background, and only vote shares are observed. The solid curve is the minimal L^1 distance between observed shares and shares generated by any distribution of voter mass over the cells of the superimposed tessellation at each candidate weight. The shaded region is the identified set $[0.80, 1.18]$, where the fit is exact; it contains the truth. The dotted lines mark the sieve estimates of Merlo and De Paula (2017) at $J_t = 2, 3$ ($\hat{w} = 1.20, 1.27$), both of which lie outside the identified set.

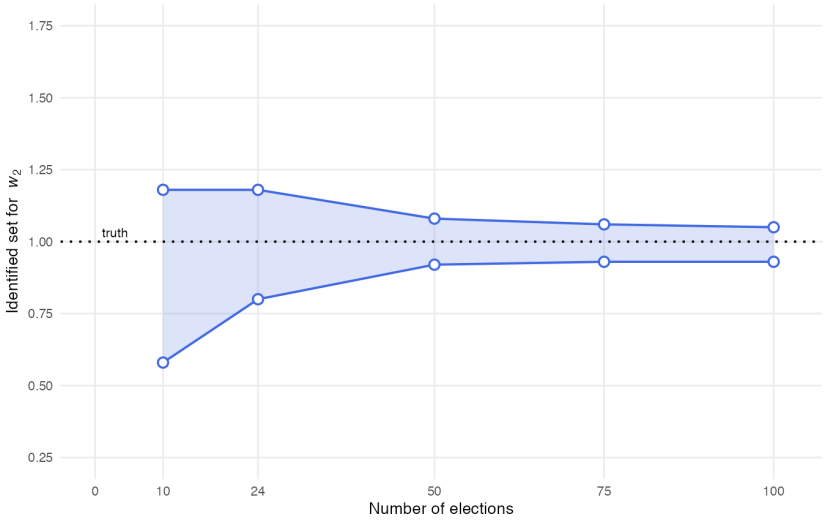


Figure D2. Identified set for the weight at different numbers of elections

Notes: Each point pair gives the lower and upper bounds of the identified set for w (the weights admitting an exact fit), computed as in Figure D1 but varying the number of elections. The dotted line marks the true value $w = 1$.

noise converts the flat exact-fit region into a well-defined basin centered near $w = 1$. The sieve estimator yields $\hat{w} = 1.09$ and 1.40 at $J_t = 2, 3$. Note that in the simulations the results obtained with $J_t = 2$ are closer to the truth; this suggests that additional flexibility in the estimator does not necessarily translate into more precise estimates.

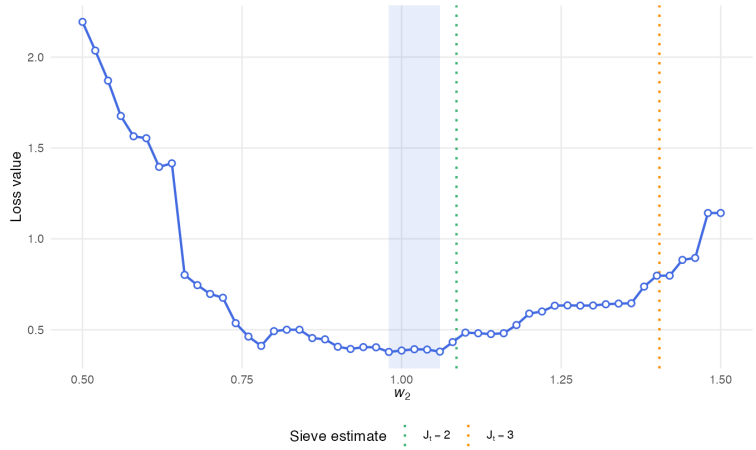


Figure D3. Set identification of the weight via the method used in this paper, and sieve estimates from Merlo and De Paula (2017) with voter perception noise

Notes: Same setup as Figure D1, except that each voter perceives every candidate's position with an idiosyncratic $\mathcal{N}(0, 0.1^2 I)$ mislocation error. The solid curve is the minimal L^1 distance between observed shares and shares generated by any distribution of voter mass over the cells of the superimposed tessellation at each candidate weight; with noise the loss is strictly positive everywhere. The shaded region is the near-feasible set $[0.98, 1.06]$ (fit within 0.01 of the minimum), which contains the true value $w = 1$. The dotted lines mark the sieve estimates of Merlo and De Paula (2017) at $J_t = 2, 3$ ($\hat{w} = 1.09$, 1.40), both outside the near-feasible set.

Removing the slivers. In the application, I remove the cells of negligible measure (see Section C for details). Figure D4 demonstrates that removing these cells does not compromise the recovered values. Using the noise-free setup and 24 elections, at each candidate weight I drop the smallest k% of the cells of the superimposed tessellation, ranked by area, and re-solve the exact-fit problem on the remaining cells. Dropping cells amounts to imposing that they carry no voter mass, so it can only remove weights from the identified set and not add them. In practice, the identified set is unaffected by removing up to 5% of the cells ($[0.80, 1.18]$) and contracts only slightly at 10–20% ($[0.82, 1.16]$). These cells are slivers covering less than 1% of the policy space. Even when half of the cells are removed, an exact fit remains attainable and the identified set $[0.86, 1.12]$ still contains the truth; at 60% it narrows to $[0.94, 1.06]$. Only beyond this point is no weight able to reproduce the observed shares exactly. The minimal loss is attained at $w = 0.98$, 0.92, and 0.94 when 70%, 80%, and 90% of the cells are removed, respectively. The loss profile remains relatively smooth even after removing large shares of the smallest cells, with the flat steps of observationally equivalent partitions most visible in the full-cell setup and fading as small cells are dropped.

Intuitively, these results follow because the mass of the removed cells can be reallocated to neighboring cells at little cost in fit, while the variation that identifies the weight is preserved. Moreover, when ideal points follow a continuous distribution, as in the simulation, smaller cells carry smaller mass than their larger neighbors, so dropping cells of negligible measure has a negligible effect. Dropping the largest cells instead matters more: the identified set changes shape once 20% of the cells are removed and collapses to the loss-minimizing weight at 40%, beyond which the loss function is no longer smooth too.

Number of cells and fit. The weight determines not only which voting profiles are available but also how many. In the application I fix the number of cells across weights (Section C). The simulations below show why this is warranted in the noisy setup.

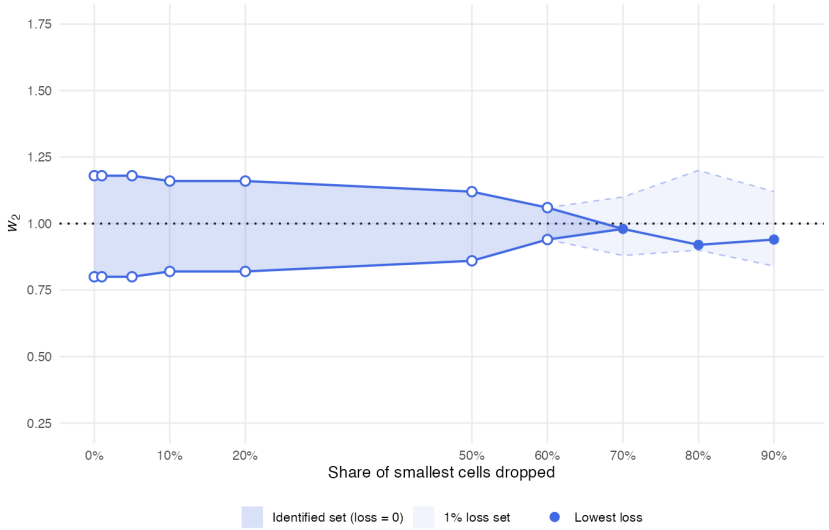


Figure D4. Identification of the weight removing the cells of negligible measure (in % of overall number of cells)

Notes: Same setup as Figure D1, except that at each candidate weight the smallest k% of the cells of the superimposed tessellation (by area) are removed before solving for the distribution of voter mass over cells. The removed cells cover 0.01%, 0.15%, 0.55%, 5.2%, and 9.1% of the policy space for $k = 1, 10, 20, 50, 60$, and 15%, 25%, and 43% for $k = 70, 80, 90$. The dark shaded region is the identified set (weights at which an exact fit is attainable), whose bounds are marked by hollow circles; for $k \geq 70$ no exact fit exists, and the filled circles mark the weight minimizing the L^1 loss. The light shaded region with dashed bounds is the set of weights whose fit is within one percentage point of share mass per election of the minimum. The dotted line marks the true value $w = 1$.

First, across weights the loss falls with the number of cells: a linear fit over all weights has $R^2 = 0.65$ at perception noise 0.2, and $R^2 = 0.48$ among the weights close to the true one, in $[0.80, 1.20]$. Second, the relationship across the whole weight grid reflects the weight, not the count. Keeping the same number of cells at every weight makes the R^2 almost unchanged.

Nevertheless, the count matters at the optimum. In 100 draws of candidate positions with noise 0.2, the lowest-loss weight is centred on the truth (mean 1.01, median 1.00). The lowest-loss weight induces about 4 more cells than the weights within 0.1 of it (95% confidence interval 1.4 to 5.6). A marginally richer mesh can thus tip which of two near-equivalent weights attains the minimum, but does not shape the loss profile. Keeping a fixed number of the largest cells removes this edge at no cost in accuracy: the lowest-loss weight lies within 0.10 of the truth in 62 of the 100 draws with a fixed number of cells (and in 51 with all cells). In a mesh with a fixed number of cells, differences in fit across weights therefore reflect the composition of the cells alone.

E Stated Preferences

E.1 Preference bins

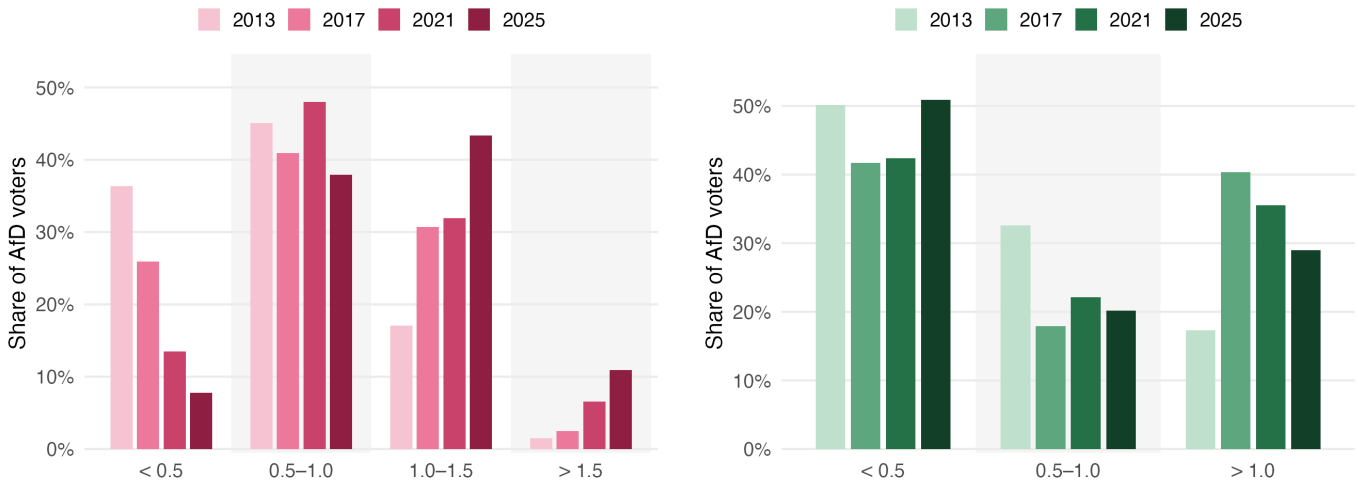
To see whether the evolution of the AfD electorate implied by the model is visible in stated attitudes, I locate the respondents of the GLES election surveys on the two dimensions. AfD voters are respondents who report an AfD party vote/party-vote intention. All shares use the GLES socio-demographic and regional weights.

Left–Right. The questions on the Left–Right dimension in the voter surveys are numerous and similar to those in the candidate surveys. The common items are the Left–Right self-placement, the immigration scale, the taxes-versus-social-services scale (2021 and 2025), and the issue statements on immigrants adapting to

German culture, the state staying out of the economy and reducing income differences, complemented in single years by same-sex marriage, immigration being good for the economy, defense spending and support in the euro crisis. Voter positions can therefore be reproduced well from the items common to the candidate and voter surveys.

Anti-Elite. Most anti-elite items of the candidate survey have no voter counterpart, so the Anti-Elite position of voters is built from the voter survey alone. It is the mean of three standardized items: (i) "the differences between the elite and the people are larger than the differences among the people" (2017–2025); (ii) "citizens should be able to call referendums at the federal level" (2013–2021; in 2025, when the item was not asked, it is replaced by "the people, not the politicians, should make the most important political decisions"); and (iii) "parties are only interested in votes, not in the opinions of voters" (all years). Voters' positions on this dimension are therefore measured with different items from candidates' positions, and the items also differ across years.

Figure E1 reports the share of AfD voters in fixed bins of each dimension. On the Left–Right dimension, the electorate shifts steadily to the right: the share below 0.5 falls from 36% in 2013 to 26%, 13%, and 8% in the following elections, while the share above 1.0 rises from 19% to 33%, 38%, and 54%. On the Anti-Elite dimension, 2013 stands apart, and it is also measured with a different set of items. Among the 2017–2025 elections, the most recent is clearly the least anti-elite for the AfD electorate: 51% of AfD voters lie below 0.5 in 2025, against 42% in 2017 and 2021, and the share above 1.0 falls from 40% in 2017 to 29% in 2025. These survey-based trends are therefore qualitatively consistent with the revealed-preference estimates in the main text.



(a) Left–Right (b) Anti-Elite

Figure E1. Distribution of AfD voters across fixed bins of the Left–Right and Anti-Elite dimensions, GLES voter survey 2013–2025.

Notes. Bars show the weighted share of AfD voters in each bin; shares sum to 100% within a year. Data: GLES Voters Surveys, 2013–2025. AfD voters report an AfD party vote after the election, or an AfD party-vote intention before the election. Panel (a): the distribution of the Left–Right positions among AfD voters. Panel (b): the distribution of the Anti-Elite positions among AfD voters.

E.2 Motivations exercise

This section runs an exercise similar to the one in Section 6.2.2 using the stated-preference positions of Appendix Section E.1.

Before turning to the exercise and its results, I want to stress that recovering the weights from these

data remains infeasible, so I take the estimated weights from the analysis in this paper. This is so for at least two reasons. First, the survey samples of AfD voters are small (fewer than 500 in every election before 2025, and 88 in 2013). Second, the problem of voters adjusting their preferences so as to be closer to the party is especially acute for weight estimation, since such closeness would be read as evidence that a dimension is important, which may be precisely the opposite of the truth if individuals are adjusting their stated positions on the issues that matter least to them.

Each AfD voter's ideal point is given by the Left–Right and Anti-Elite positions of Section E.1. The AfD position is given by the national AfD candidate medians used in the main exercise. To place voters on the candidate Anti-Elite scale, the voter index is rescaled to match the mean and standard deviation of a candidate-scale score built from the two anti-elite items common to both surveys; since voters are placed on the candidate Left–Right scale, the Left–Right coordinate applies directly.

Figure E2 plots the stated-preference shares over the results from Figure 9b. The survey reproduces the shift of the AfD electorate towards the Right region by 2025, but not the earlier change between 2013 and 2017; the latter might reflect the small 2013 survey sample and the different anti-elite items available that year, as well as voters' preference misreporting (in particular, not reporting radical-right views).

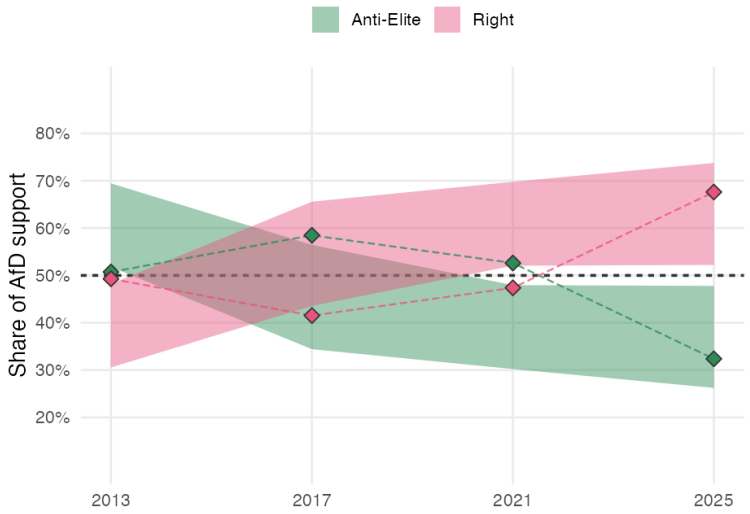


Figure E2. Comparison to the survey results

Notes: Survey benchmark constructed from the GLES surveys of Figure E1 (2013, 2017, 2021 and 2025). Each AfD voter's ideal point is given by the Left–Right position (candidate scale) and the Anti-Elite index of Section E.1. The AfD position is the national AfD candidate median; the two counterfactual positions lie one unit away from it, as in Section 6.2.2. AfD voters are assigned to the closer counterfactual position under the region-year weights w_{AE} estimated in the main analysis. The resulting shares of AfD voters are point estimates, not partial-identification bounds. The estimates are weighted throughout using the GLES socio-demographic and regional weights.

F Comparison with the Merlo–De Paula Point Estimates

The bounds of Section 6.2.2 are derived under no assumptions on the functional form of the distributions of ideal points. Merlo and De Paula (2017) instead obtain point estimates of the same spatial model by restricting the distribution to the semi-nonparametric (SNP) class of Gallant and Nychka (1987) and estimating it by sieve minimum distance. This section applies their estimator to the same data and target parameters. The ideal-point density of each of the 24 socio-demographic groups is an SNP density, $f_g(t) \propto P_g(z)^2\,\phi(z)$ with $z = R^{-1}(t - b_g)$, where P_g is a polynomial of degree at most two, b_g is the group mean and R a common scale matrix; the municipal vote shares are the mixture of the group densities integrated over the parties' areas of attraction, with the turnout-adjusted demographic composition as mixture weights, exactly as in

the linear program. The weight matrix W is estimated jointly. Following Merlo and De Paula (2017), group means are anchored at the party centroids with a bounded deviation. The AfD composition shares are then computed from the estimated densities under the same counterfactual masks as in the main text.

Figure F1 overlays the resulting point estimates on the bounds. Their point estimates lie inside or close to the bounds in 2017–2025, but not in 2013, and display no systematic trend: the dominant motive switches between Anti-Elite and Right from one election to the next. They should be read as an illustration of what the parametric restriction adds rather than as an alternative estimate: the SNP class fixes the shape of each group's density, so the point estimates depend on the anchoring and on the polynomial order, whereas the bounds do not. Figure F2 shows the net gains of the CDU and the AfD in the repositioning game under the point estimates of the ideal-point densities. Whatever the CDU's move, the AfD's best response is to move right, which captures the most voters on the right. This is because the fitted density places a large mass of voters at far-right, pro-elite positions.

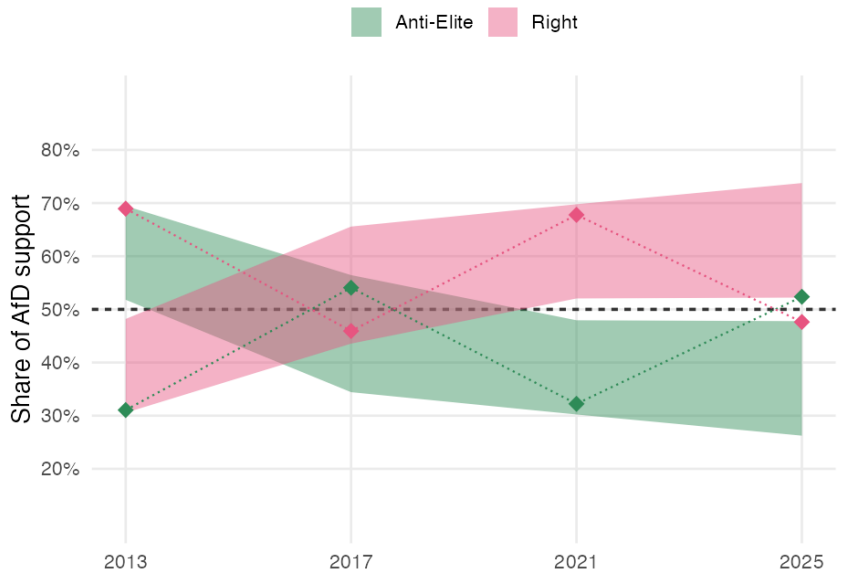


Figure F1. Merlo–de Paula point estimates vs. estimated bounds using the methodology of this paper

Notes. Bands are the bounds of Figure 9b. Diamonds are the sieve minimum-distance point estimates of the same shares under the SNP restriction of Merlo and De Paula (2017) (anchored group means, polynomial order two), aggregated over East and West with AfD-vote weights. Estimation uses the same municipality sample, socio-demographic groups and counterfactual masks as the bounds; the weight matrix is estimated jointly and agrees with the grid-searched w_{AE} of the main text.

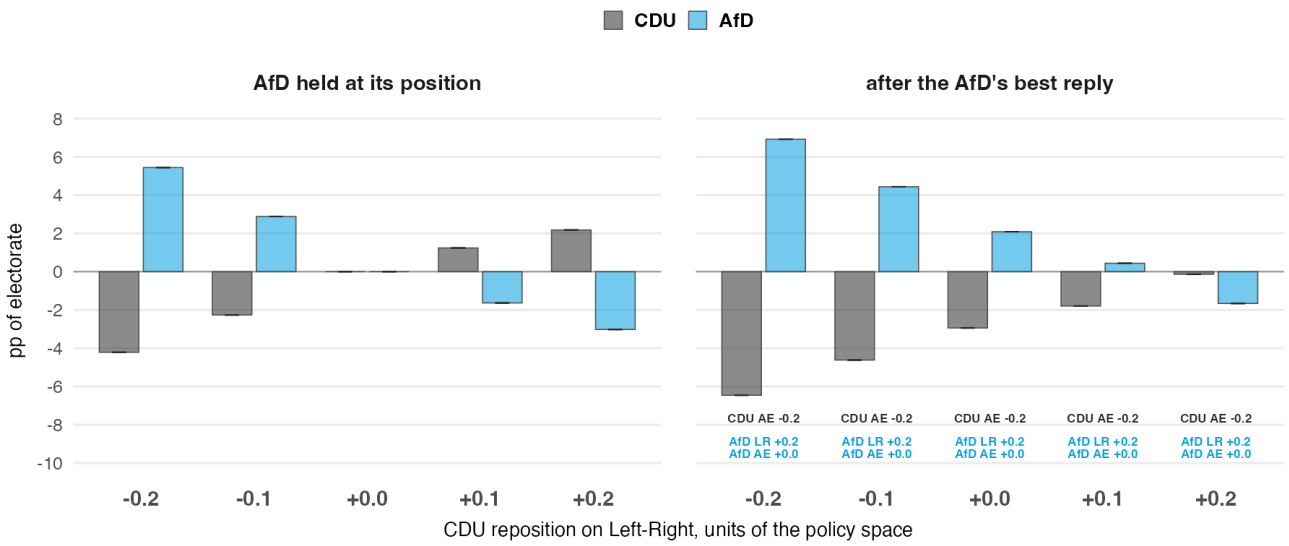


Figure F2. The repositioning game under the Merlo–de Paula point estimates

Notes. Counterpart of Figure 12 with the ideal-point distribution replaced by the sieve minimum-distance point estimates of Merlo and De Paula (2017). Left: the CDU moves on the Left–Right axis, the AfD stands still. Right: the CDU pairs each Left–Right move with its best anti-elite move (restricted to steps away from anti-elite rhetoric), and the AfD replies on both dimensions (anti-elite steps upward only); the profile under each pair of bars is the AfD's reply.

G The Role of the Supply Side: Shapley Values

Let $\theta_{Right}(Pref_t, Comp_t, Pos_t)$ denote the share of AfD voters whose ideal points fall in the *Right* region (Figure 9a), given voter preferences $Pref_t$ across socio-demographic groups, the demographic composition of AfD voters $Comp_t$, and national median party positions Pos_t. I decompose the change in this share between t and t' as follows:

$$\begin{aligned}\theta_{Right}(Pref_{t'},\, Comp_{t'},\, Pos_{t'}) - \theta_{Right}(Pref_t,\, Comp_t,\, Pos_t) \\ = \phi(Pref) + \phi(Comp) + \phi(Pos)\end{aligned} \tag{G1}$$

where each Shapley value, e.g., $\phi(Pos)$, is the weighted average marginal contribution of this factor across all 2^2 coalitions of the other two factors:

$$\begin{aligned}\phi(Pos) \;=\; & \tfrac{1}{3}\Big[\theta_{Right}(Pref_{t'}, Comp_{t'}, Pos_{t'}) - \theta_{Right}(Pref_{t'}, Comp_{t'}, Pos_t)\Big] \\ & + \tfrac{1}{6}\Big[\theta_{Right}(Pref_t, Comp_{t'}, Pos_{t'}) - \theta_{Right}(Pref_t, Comp_{t'}, Pos_t)\Big] \\ & + \tfrac{1}{6}\Big[\theta_{Right}(Pref_{t'}, Comp_t, Pos_{t'}) - \theta_{Right}(Pref_{t'}, Comp_t, Pos_t)\Big] \\ & + \tfrac{1}{3}\Big[\theta_{Right}(Pref_t, Comp_t, Pos_{t'}) - \theta_{Right}(Pref_t, Comp_t, Pos_t)\Big],\end{aligned}$$

and $\phi(Comp)$, $\phi(Pref)$ are defined analogously. The weights $\frac{1}{3}, \frac{1}{6}$ follow from $\frac{|S|!\,(n-|S|-1)!}{n!}$ with $n = 3$ ($|S| = 0, 2$ give $\frac{1}{3}$; $|S| = 1$ gives $\frac{1}{6}$). By construction $\phi(Pref) + \phi(Comp) + \phi(Pos)$ equals the left-hand side of (G1) exactly.

To evaluate each counterfactual share, I combine three ingredients. Party positions (Pos_t) are the estimated national median positions in year t; voter preferences ($Pref_t$) enter as group-specific ideal-point

distributions, estimated as a primitive of the model, $f(x \mid s, d)$.

Composition ($Comp_t$) comes from the Representative Electoral Statistics. These are unique data on voting choices for 12 age–gender groups per federal state.[49] These data do not permit disaggregation by education, which I therefore omit.

Table G1 presents the results for $\phi(Pos)$ across every pair of consecutive years and regions of the country. The last row provides conservative bounds aggregated over the two parts of the country. Table G2 also presents the Shapley values for the other two channels (composition and preferences), but bounds for those channels are wide and uninformative.

Table G1. Shapley Decomposition of the Change in the Right Share of AfD Support

	2013→2017	2017→2021	2021→2025
West			
Positions	[−0.202, −0.026]	[−0.047, +0.000]	[+0.023, +0.117]
East			
Positions	[−0.103, −0.018]	[+0.000, +0.018]	[−0.041, −0.005]
National (AfD-electorate weighted)			
Positions	**[−0.188, −0.025]**	**[−0.038, +0.004]**	**[+0.008, +0.092]**

Notes: Bounds on the Shapley contribution of party positions, $\phi(Pos)$ in equation (G1). Each bound is the minimum or maximum of $\phi(Pos)$ over the identified set, with one distribution of ideal points per year entering all counterfactual shares, the AfD regions under both years' positions imposed in the same linear program, and each group's AfD total held at its fitted value. Positive values raise the Right share θ_{Right} (Section 6.2.2). The national row weights the East and West bounds by each region's share of the AfD electorate and, since this share differs between t and t', reports the envelope over the two years' weights.

Table G2. Shapley Decomposition of the Change in the Right Share of AfD Support (full)

	2013→2017	2017→2021	2021→2025
West			
Composition	[−0.102, +0.141]	[−0.138, +0.119]	[−0.122, +0.106]
Positions	[−0.202, −0.026]	[−0.047, +0.000]	[+0.023, +0.117]
Preferences	[−0.287, +0.506]	[−0.432, +0.466]	[−0.345, +0.295]
East			
Composition	[−0.090, +0.048]	[−0.066, +0.042]	[−0.031, +0.045]
Positions	[−0.103, −0.018]	[+0.000, +0.018]	[−0.041, −0.005]
Preferences	[−0.384, +0.234]	[−0.233, +0.312]	[−0.302, +0.157]

Notes: Bounds on the Shapley contributions $\phi(Comp)$, $\phi(Pos)$ and $\phi(Pref)$ of equation (G1). Each bound is the minimum or maximum of the Shapley value over the identified set, with one distribution of ideal points per year entering all counterfactual shares, the AfD regions under both years' positions imposed in the same linear program, and each group's AfD total held at its fitted value, so that all shares are linear in the cell masses. Positive values raise the Right share θ_{Right} (Section 6.2.2).

The bounds on the Shapley values for the Positions channel suggest that the supply-side channel contributes positively to the increase in the share of Right-motivated voters over 2021–2025, while its contribution for the other pairs of years leans negative. What these results make clear is that the evolution over time in the share of Right-motivated voters documented in this paper is not primarily driven by parties changing their positions.

[49] They are based on a sample of polling districts in which marked ballots allow votes and turnout to be tabulated by gender and age group. Hence, these data are based on actual ballots, not surveys.

H Coalition-Building and the Firewall Against Cooperation with the Far Right

The sections in the main text asked how mainstream parties could foreclose the rise of populism by repositioning in a two-dimensional policy space. This section studies the welfare consequences of a different measure the establishment parties could use in their electoral competition with populists: the firewall against cooperation with them. Since the AfD entered the Bundestag in 2017 there has been an unwritten rule of no cooperation with it;[50] comparable rules were enacted in France, Sweden and Belgium, and in most cases abandoned once the radical-right party became pivotal for some governing majority.

Until 2025 the German rule was largely non-binding at the national level: the AfD's vote share was modest and its parliamentary group small. In the 2025 elections it outpolled the SPD (historically one of the two largest establishment parties) in the party-list vote, and, absent the firewall, would have been a plausible coalition partner for the CDU (i.e. CDU/CSU).

Among other things, critics of the firewall policy argue that it leaves a large share of the electorate unrepresented, and this section tests that argument via coalition formation. Given the vote shares, only two coalitions were feasible: the centrist CDU–SPD and the right-wing CDU–AfD. Taking the estimated 2025 voter preferences as given, I ask (i) which coalition maximizes social welfare of the whole electorate and (ii) which coalition is preferred by a majority of voters.

Figure H1a illustrates the logic behind the coalition analysis. The policy implemented by the coalition is assumed to be the midpoint of its members' positions – blue and red diamonds for the CDU+AfD and CDU+SPD coalitions, respectively. As before, the analysis uses cells as the unit of analysis; some cells are unambiguously closer to one coalition policy than to the other, while some cells are divided (as illustrated in the figure). I evaluate the coalitions using two different metrics – social welfare and majoritarian choice. I focus in this exercise on voters of all parties.

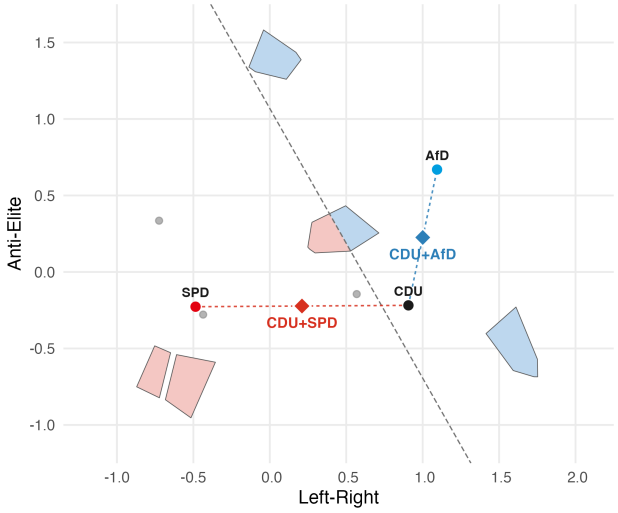
Social welfare from coalitions. Let $Y_{\text{CDU+AfD}}$ and $Y_{\text{CDU+SPD}}$ denote the two coalitions' implemented policies; and as before let $f(x \mid s, d)$ be the density of ideal points among all voters who turned out in state s and demographic group d, and $F_{S,D}$ their joint distribution. The target parameter is the difference of weighted Euclidean distances $d_w(\cdot,\cdot)$ from voters' ideal points to the two coalitions' policies:[51]

$$\theta_{\text{Welfare}} = \iint_{\mathcal{R}} \Big[d_w\big(x, Y_{\text{CDU+AfD}}\big) - d_w\big(x, Y_{\text{CDU+SPD}}\big) \Big] \, f(x \mid s, d)\, dx\, dF_{S,D}(s, d).$$

The bounds for θ_{Welfare} are in the upper table of the right panel of Figure H1, where negative values mean that the CDU+AfD policy is closer to voters on average than the CDU+SPD policy. The estimated bounds at the national level straddle zero, so under this welfare rule neither coalition can be ranked unambiguously, although the bounds lean negative. This negative lean is driven by Eastern Germany, where the difference is bounded strictly below zero; in the West the bounds are roughly centered around zero, favoring neither coalition.

[50] The rule is informal at the federal level, but the CDU codified it in its 2018 incompatibility resolution (*Unvereinbarkeitsbeschluss*). At the municipal level compliance is high but not universal: Schroeder, Ziblatt and Bochert (2025) find no cooperation in about 81% of district-council votes in which cooperation was feasible between 2019 and 2024.

[51] As before, I report sharp bounds: for each cell c of the partition I compute the range $[D_{\min}(c),\ D_{\max}(c)]$ of the difference $D(x) = d_w(x, Y_{\text{CDU+AfD}}) - d_w(x, Y_{\text{CDU+SPD}})$ over $x \in c$, and the lower (upper) bound on θ_{Welfare} is the minimum (maximum) of the cell-mass-weighted sum of $D_{\min}(c)$ ($D_{\max}(c)$) over the identified set of the main linear program. These bounds are sharp: they impose a common preference distribution for both coalitions and, within each cell, a common ideal-point location for both distances.



(a) Illustration: Coalition policy positions and random cells of voters' locations

Welfare difference, CDU+AfD − CDU+SPD	
National	$[-0.18,\ 0.03]$
East	$[-0.29,\ -0.10]$
West	$[-0.13,\ 0.08]$
Share of the electorate closer to CDU+AfD	
National	[61.9%, 72.8%]
East	[69.9%, 83.8%]
West	[58.5%, 68.1%]

(b) Bounds on the target parameter values, θ_{Welfare} and θ_{Majority}

Figure H1. Welfare difference and majority preference between the two 2025 coalitions

Notes: Illustration of the coalition-welfare exercise (Panel (a)). Dots mark the 2025 national median positions of the parties. Each coalition is assumed to implement the midpoint of its members' positions (diamonds): CDU+AfD in blue, CDU+SPD in red. The dashed gray line is the set of voters indifferent between the two coalition policies. The polygons are illustrative cells of the preference partition: a cell is shaded blue (red) when every point in it is closer to the CDU+AfD (CDU+SPD) policy; the cell crossed by the indifference line is split, with each part shaded by the closer policy; within-cell locations are not identified, so such cells can only be bounded. Panel (b) reports bounds over the identified set of preference distributions. Distances are in the region-specific w_{AE}-weighted metric; no assumption is made on the location of mass within a cell. The welfare-difference rows bound $\mathbb{E}[d_w(x, Y_{\text{CDU+AfD}}) - d_w(x, Y_{\text{CDU+SPD}})]$ directly in one program per bound (sharp); negative values mean the CDU+AfD policy is closer on average. The majority rows bound the mass of cells lying entirely closer to the CDU+AfD policy (lower) and with any point closer (upper); they are invariant to any increasing transformation of the distance. National bounds combine East and West; the regional programs are independent, so the combination is sharp.

Note that this welfare estimation is not invariant to the distance metric; Table H1 provides the same welfare analysis for the squared Euclidean distance.

Table H1. Welfare difference between the two 2025 coalitions under quadratic loss: equal-weighted policy

Welfare difference (quadratic), CDU+AfD − CDU+SPD	
National	$[-1.43,\ 0.07]$
East	$[-3.33,\ -0.53]$
West	$[-0.62,\ 0.33]$

Notes: Sharp bounds over the identified set of preference distributions on $\mathbb{E}\left[d_w^2(x, Y_{\text{CDU+AfD}}) - d_w^2(x, Y_{\text{CDU+SPD}})\right]$; negative values mean the CDU+AfD policy is closer to voters on average. Each coalition implements the simple average of its members' 2025 national median positions. National bounds combine East and West.

Majoritarian choice of the coalition. Another way to estimate which coalition is preferred by society is via the share of the electorate strictly closer to one coalition policy than to the other. To estimate this, I

formulate the target parameter as the share of voters preferring the CDU+AfD policy,

$$\theta_{\text{Majority}} = \iint_{\mathcal{R}} \mathbb{1}\Big\{ d_w\big(x, Y_{\text{CDU+AfD}}\big) < d_w\big(x, Y_{\text{CDU+SPD}}\big) \Big\}\, f(x \mid s, d)\, dx\, dF_{S,D}(s, d).$$

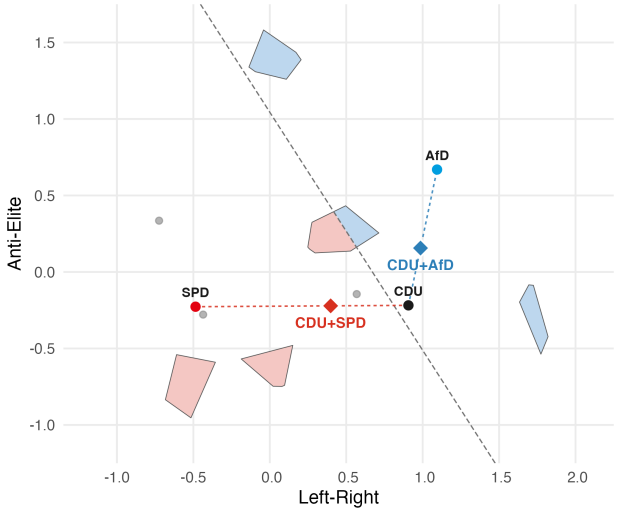
For θ_{Majority}, the lower bound counts the mass of voters lying in the cells c with $D_{\max}(c) < 0$ (entirely closer to CDU+AfD, as in the color coding in Figure H1a) and the upper bound the mass of cells with $D_{\min}(c) \leq 0$.

The results are in the lower part of Figure H1(b). For every preference distribution in the identified set, a majority of voters in both parts of the country is closer to the CDU+AfD policy: at least 69.9% in the East and at least 58.5% in the West. Hence the CDU+AfD coalition would win a pairwise vote against the CDU+SPD coalition. The underrepresentation argument is therefore nuanced and depends on the decision rule one has in mind when making such claims. This analysis nevertheless highlights once again the regional differences: in the worst scenario only 16.2% of voters in Eastern Germany actually prefer the coalition with the SPD over the one with the AfD; this highlights that coalition building under the firewall policy could create an additional feeling of exclusion and underrepresentation among Eastern Germans in particular.

The contrast between the two measures of coalition preferences highlights an essential fact about the present-day political landscape in Germany: many voters on the left of the political spectrum hold very strong preferences against a coalition of the center with the AfD, strong enough to keep the welfare comparison ambiguous even though a clear majority sides with CDU+AfD.[52]

Figure H2 provides the same analyses as above, but under the assumption that the implemented coalition policy is the vote-share-weighted average of the coalition parties' own positions.

[52] The results broadly align with survey data, both in terms of regional heterogeneity and in terms of the strong dislike of a CDU–AfD coalition among some groups of voters. Figure H3 plots the survivor functions of the coalition desirability ratings in the West and the East. In both regions many respondents place the CDU–AfD coalition at the very bottom of the scale (47% in the West, 36% in the East, against 21% and 26% for CDU–SPD). But in the West the CDU–SPD curve dominates over most of the scale, whereas in the East the CDU–AfD curve lies above it at every positive threshold.



(a) Illustration: Coalition policy positions and voters' locations

Welfare difference, CDU+AfD − CDU+SPD	
National	[−0.13, 0.05]
East	[−0.24, −0.08]
West	[−0.08, 0.10]
Share of the electorate closer to CDU+AfD	
National	[54.8%, 70.8%]
East	[69.9%, 78.6%]
West	[48.3%, 67.5%]

(b) Bounds under the vote-share-weighted policy

Figure H2. Welfare difference and majority preference between the two 2025 coalitions: vote-share-weighted policy

Notes: Same objects as Figure H1, but each coalition implements the vote-share-weighted average of its members' positions (diamonds in Panel (a)), with weights from the 2025 second-vote shares. Both policy points shift toward the CDU. Panel (b) reports bounds over the identified set of preference distributions, constructed exactly as in Figure H1.

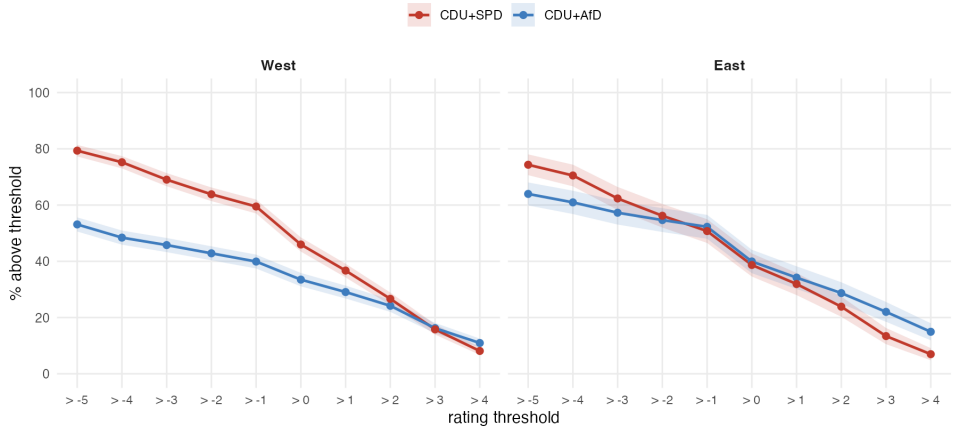


Figure H3. Desirability of a CDU–SPD and a CDU–AfD coalition, West and East Germany

Notes. The figure plots the share of respondents rating each coalition strictly above a given threshold, for every threshold on the eleven-point desirability scale. Respondents are asked how desirable each coalition would be, on a scale from −5 ("not at all desirable") to +5 ("extremely desirable"). The same respondents rate both coalitions, so the two curves within a panel are directly comparable. Data are from the GLES Tracking, fielded 12–20 February 2025, just before the federal election. Shaded bands are 95% confidence intervals.

I Tables & Figures

Table I1. Marginal (observed) shares and Fréchet bounds on the joint cell shares

	Y_{1e_2}	Y_{2e_2}	Total
Y_{1e_1}	[0, 0.1]	[**0.6**, **0.7**]	0.7
Y_{2e_1}	[0, 0.1]	[0.2, 0.3]	0.3
Total	0.1	0.9	1

Notes: The Total row and column report the observed marginals (vote shares) in two hypothetical elections, e_1 (rows) and e_2 (columns). Each interior cell reports the Fréchet bounds on the joint share of voters choosing the corresponding pair of candidates, $\max\{0, \sigma_r + \sigma_c - 1\} \leq \sigma_{rc} \leq \min\{\sigma_r, \sigma_c\}$, where σ_r and σ_c denote the row and column marginals. The bold cell shows that at least 60% of voters vote for Y_{1e_1} in e_1 and for Y_{2e_2} in e_2.

Table I2. Correlation of LR Positions between Candidates (estimated) and Parties (reported), GLES 2021, estimation sample

	Perceived Party Position, L-R			*Own L-R Position*
	All Parties	Chosen Parties	Chosen AfD	Chosen Parties
	(1)	(2)	(3)	(4)
Panel A: no individual controls				
Candidate, L-R	−0.007	0.276**	1.207***	0.106
	(0.080)	(0.135)	(0.213)	(0.166)
Demographic controls	No	No	No	No
Panel B: with demographic controls				
Candidate, L-R	−0.015	0.359**	1.243**	0.174
	(0.078)	(0.140)	(0.408)	(0.169)
Demographic controls	Yes	Yes	Yes	Yes
State F.E.	Yes	Yes	Yes	Yes
Party F.E.	Yes	Yes	No	Yes
Observations	19,992	2,269	210	2,219

Note: Coefficient on the estimated L–R position of the respondent's district candidate of the party; GLES 2021 pre- and post-election cross-sections, survey weights. Columns (1)–(3): dependent variable is the respondent's placement of the party on the 1–11 left–right scale, for all parties, for the party chosen with the second vote (party-list), and for AfD voters only; column (4): the respondent's own left–right self-placement. Panel B adds gender, age, age^2 and education level controls. Sample: districts of the estimation sample (districts with at most three imputed candidates). Standard errors clustered by party × state in parentheses. *p<0.1; **p<0.05; ***p<0.01.



Figure I1. Permutation test on fit: observed versus randomly reassigned district positions

Notes. The histogram shows the distribution of the pooled fit error over 100 permutations in which candidate positions are randomly reassigned across districts within each state, separately for each year; the vertical line is the fit error under the observed positions. The fit error of a region-year is the mean absolute error of the fitted municipal AfD vote shares at the best-fitting preference distribution, in percentage points; the pooled error is the geometric mean of the eight region-year errors, with permutation draws matched across region-years. The observed assignment ranks second among the 101 values ($p = 0.02$).

Table I3. Candidates' LR positions are uncorrelated with voters' self-reported views

	Candidate LR Position		
	All Parties	Chosen Parties	Chosen AfD
	(1)	(2)	(3)
Panel A: no individual controls			
Self, LR	0.0008	0.0023	0.0005
	(0.0011)	(0.0037)	(0.0171)
Demographic controls	No	No	No
Panel B: with demographic controls			
Self, LR	0.0017	0.0039	−0.0000
	(0.0012)	(0.0040)	(0.0169)
Demographic controls	Yes	Yes	Yes
State F.E.	Yes	Yes	Yes
Party F.E.	Yes	Yes	No
Observations	19,450	2,219	205

Note: Estimation is based on the GLES 2021 (pre- and post-election cross-sections), which records each individual's electoral district; survey weights. Dependent variable: estimated L–R position of the party's district candidate. Self-reported views take values from 1 (very left) to 11 (very right). Panel B adds gender, age, age^2 and school education (four groups). Sample: districts of the estimation sample (districts with at most three imputed candidates). Standard errors clustered by party × state in parentheses. *p<0.1; **p<0.05; ***p<0.01.

Table I4. Small-minus-large difference in second-vote shares, pooled over 2013–2025, at three levels of comparison

	CDU/CSU	SPD	Greens	FDP	Left	AfD
	(1)	(2)	(3)	(4)	(5)	(6)
Raw difference	4.37	−3.16	−2.38	−0.60	−1.00	2.09
	(0.24)	(0.19)	(0.16)	(0.06)	(0.17)	(0.18)
Within state	4.44	−1.77	−2.23	−0.34	−1.47	1.11
	(0.18)	(0.13)	(0.12)	(0.05)	(0.06)	(0.11)
Within district	3.55	−1.59	−1.53	−0.07	−1.33	0.72
	(0.15)	(0.11)	(0.09)	(0.04)	(0.05)	(0.09)

Each cell is the coefficient on $\mathbf{1}\{\text{inhabitants} < 10{,}000\}$, in percentage points, from a separate regression of the municipality's second-vote share on that indicator, pooling the 2013, 2017, 2021 and 2025 elections. *Raw difference* holds only the election year fixed; *within state* adds a state × year effect; *within district* adds a district × year effect. Standard errors clustered by municipality in parentheses.

Table I5. Dispersion of estimated candidate positions by year and region

		Total SD		Between-party		Within-party		AfD only	
Year	Region	L-R	anti-el.	L-R	anti-el.	L-R	anti-el.	L-R	anti-el.
2013	East	0.63	0.60	0.66	0.58	0.21	0.29	0.22	0.11
2013	West	0.64	0.64	0.65	0.61	0.24	0.31	0.25	0.38
2017	East	0.66	0.58	0.67	0.56	0.24	0.28	0.29	0.26
2017	West	0.72	0.64	0.75	0.62	0.22	0.30	0.26	0.28
2021	East	0.72	0.63	0.71	0.56	0.32	0.36	0.48	0.43
2021	West	0.71	0.61	0.74	0.58	0.24	0.30	0.29	0.29
2025	East	0.69	0.42	0.72	0.33	0.22	0.29	0.23	0.32
2025	West	0.75	0.48	0.80	0.41	0.19	0.30	0.19	0.35

Notes: Standard deviations of district-level (Wahlkreis) candidate positions; between-party = SD of party means, within-party = SD of positions around the own-party mean.

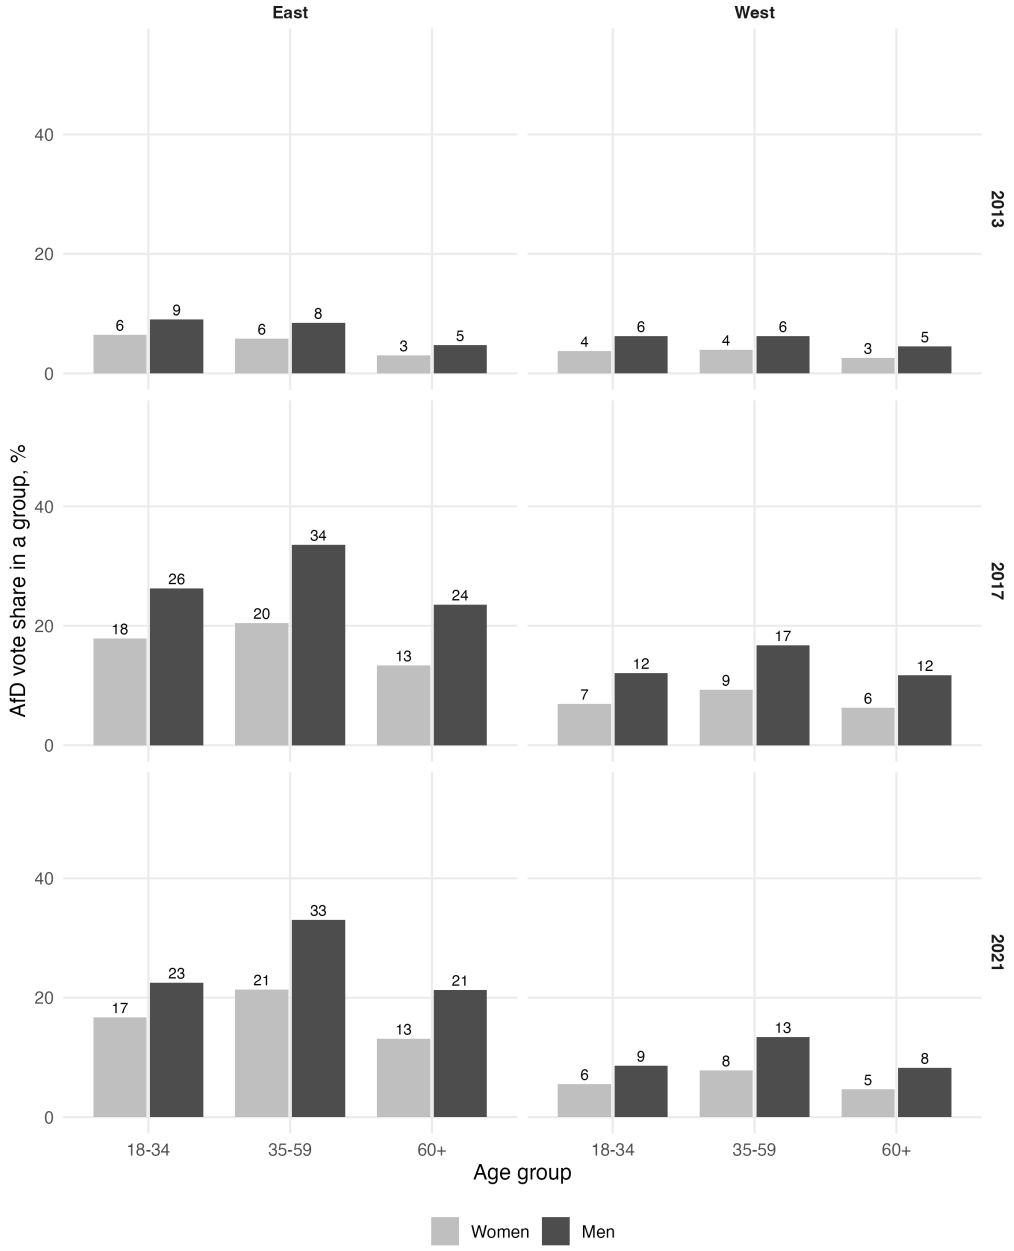


Figure I2. AfD party vote share by region, age group and gender, 2013, 2017 and 2021.

Notes. Bars show the AfD share of valid second votes cast within each group. Source: Repräsentative Wahlstatistik of the Federal Returning Officer, 2013–2021, an official statistic in which ballots in a random sample of polling districts are marked by the voter's gender, birth-year group and state. The six published age groups (18–24, 25–34, 35–44, 45–59, 60–69, 70+) are pooled into three by summing votes for visualization. East comprises Brandenburg, Mecklenburg-Vorpommern, Saxony, Saxony-Anhalt and Thuringia; West comprises the remaining states, including Berlin.

Table I6. Item-level potency (2017): shift in the dimension score from a one-point move on each survey item

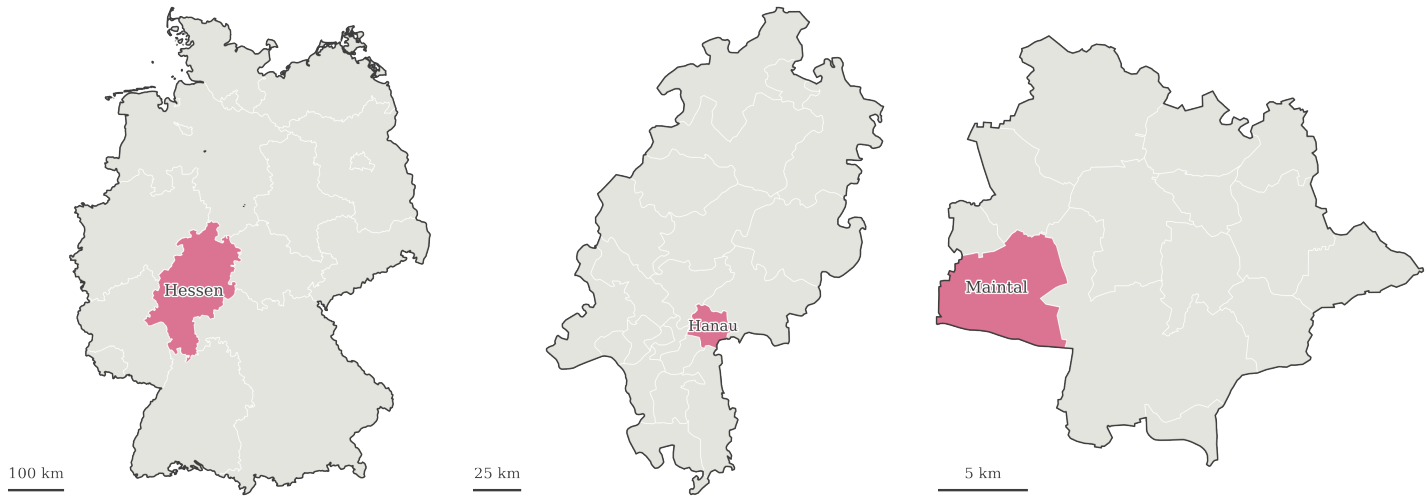
Survey item (question)	Theme	GESIS item	Scale	σ (1 pt)
Anti-elite dimension				
Satisfaction with how democracy works in Germany	Current instit.	c20	1–4	0.175
Parties are citizens' central intermediaries	Current instit.	d6c	1–5	0.144
Legislation reflects citizens' interests	Current instit.	d6b	1–5	0.137
Citizens have lost trust in politics	Current instit.	d6h	1–5	0.137
Referendums lead to badly designed laws	Direct democracy	d6g	1–5	0.138
There are enough opportunities to participate	Direct democracy	d6a	1–5	0.117
There should be federal-level referendums	Direct democracy	d6f	1–5	0.098
Left–right dimension				
Left–right self-placement	X	c3	1–11	0.038
Own position on immigration (self-placement)	Immigration	c6	1–11	0.032
Immigrants are good for the German economy	Immigration	c2i	1–5	0.093
Immigrants should adapt to German culture	Immigration	c2a	1–5	0.073
Same-sex marriage should be banned	Culture	c2d	1–5	0.090
Women should decide themselves about abortion	Culture	c2j	1–5	0.085
Providing social security is a core government task	Econ. policy	c2g	1–5	0.081
Government spending on development aid	Econ. policy	c18h	1–5	0.080
The government should stay out of the economy	Econ. policy	c2b	1–5	0.071
The government should reduce income differences	Econ. policy	c2h	1–5	0.070
Government spending on defense	Econ. policy	c18c	1–5	0.070

Notes. σ is the change in the standardized dimension score (Left–Right or Anti-Elite) induced by moving a single item one scale point, for a candidate answering all items; it is identical across parties because the score is linear in the answers. Moving a dimension score by a 0.2/0.5 step with a single item requires about $0.2/\sigma$ and $0.5/\sigma$ scale points. Units and item set are mostly fixed across years; the 2017 survey items are provided as an example.

Table I7. Item-level potency (2025): shift in the dimension score from a one-point move on each survey item

Survey item (question)	Theme	GESIS	Scale	σ (1 pt)
Anti-elite dimension				
Satisfaction with how democracy works in Germany	Current instit.	c25	1–4	0.245
Parties are citizens' central intermediaries	Current instit.	d4c	1–5	0.201
Legislation reflects citizens' interests	Current instit.	d4b	1–5	0.192
Citizens have lost trust in politics	Current instit.	d4g	1–5	0.192
There should be federal-level referendums	Direct democracy	d4f	1–5	0.138
Left–right dimension				
Left–right self-placement	X	c3	1–11	0.041
Own position on immigration (self-placement)	Immigration	c6	1–11	0.035
Immigrants are good for the German economy	Immigration	c2i	1–5	0.102
Immigrants should adapt to German culture	Immigration	c2a	1–5	0.080
Marriage should be only for heterosexual couples	Culture	c2d	1–5	0.099
Women should decide themselves about abortion	Culture	c2j	1–5	0.093
Providing social security is a core government task	Econ. policy	c2g	1–5	0.089
The government should stay out of the economy	Econ. policy	c2b	1–5	0.078
The government should reduce income differences	Econ. policy	c2h	1–5	0.076
Own position on taxes vs. social spending	Econ. policy	c9	1–11	0.030

Notes. σ is the change in the standardized dimension score (Left–Right or Anti-Elite) induced by moving a single item one scale point, for a candidate answering all items; it is identical across parties because the score is linear in the answers. Moving a dimension score by a 0.2/0.5 step with a single item requires about $0.2/\sigma$ and $0.5/\sigma$ scale points. Items are standardized with their pooled cross-year standard deviation; σ is larger in 2025 than in earlier years because fewer items are asked (five instead of seven on the Anti-Elite dimension). The 2025 items are shown as an example.



(a) **Federal states**. Hessen highlighted in pink.

(b) **Electoral districts** within Hesse. Hanau highlighted in pink.

(c) **Municipalities** within Hanau. Maintal highlighted in pink.

Figure I3. Administrative and electoral hierarchy in Germany.

Municipalities are nested within electoral districts, which are in turn nested within federal states. The figure illustrates these levels for one example state and one example electoral district.

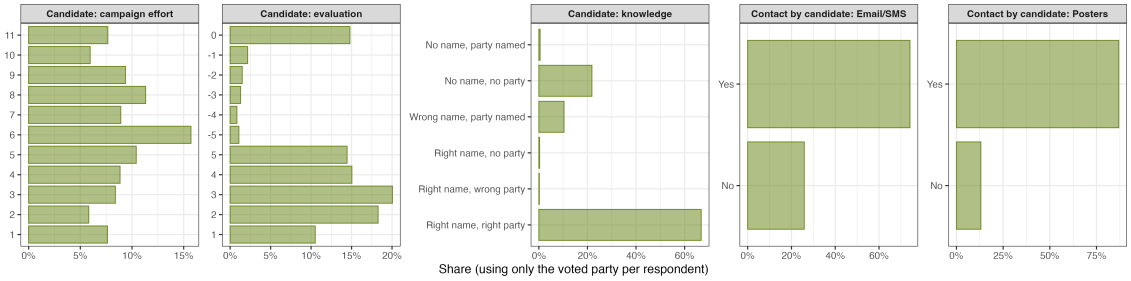


Figure I4. Voters' knowledge of and interaction with SMD candidates

Notes. **Campaign effort (1–11):** 1 = mostly about the candidate; 11 = mostly about the party; **Candidate knowledge:** respondents' ability to name the constituency candidate and their party; **Contact:** whether the respondent was contacted via the listed means. Data are for the 2021 Bundestag election.

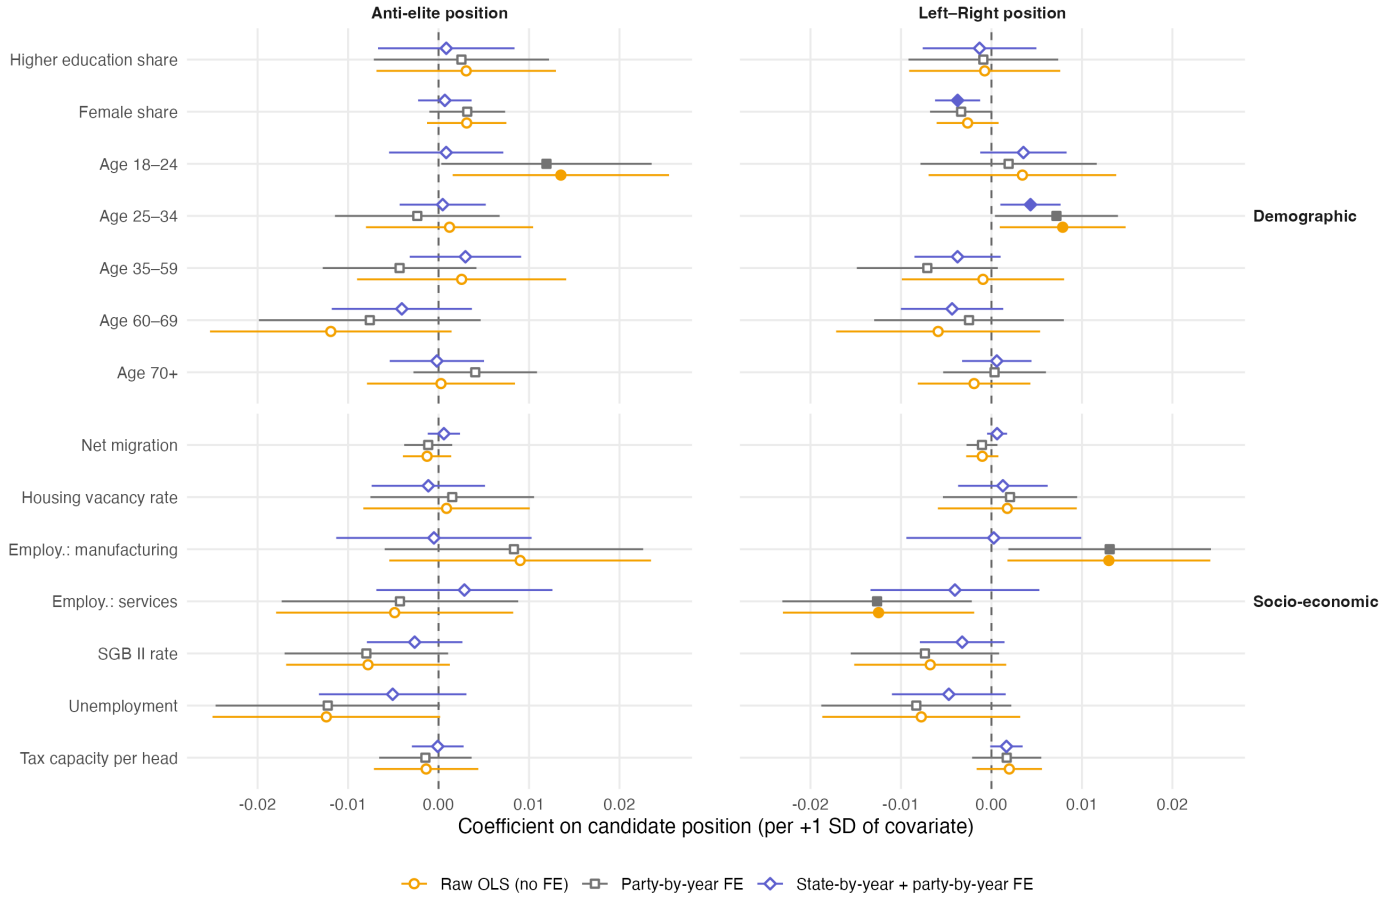


Figure I5. Association between candidates' positions on each dimension and municipalities' demographic and socio-economic characteristics.

Notes. Each marker is the coefficient from a separate regression of a candidate position on one characteristic. The unit of observation is a municipality-by-party cell in the 2013–2025 elections. The sample consists of municipalities used in the estimation of the model (see Sections 4.2 and 5 for details). All covariates are standardized to mean zero and unit standard deviation. The specifications are (i) Raw OLS with no fixed effects, (ii) party-by-year fixed effects, and (iii) state-by-year plus party-by-year fixed effects. Standard errors are clustered by district-by-year. Horizontal bars are 95% confidence intervals. Filled markers denote coefficients significant at the 5% level.

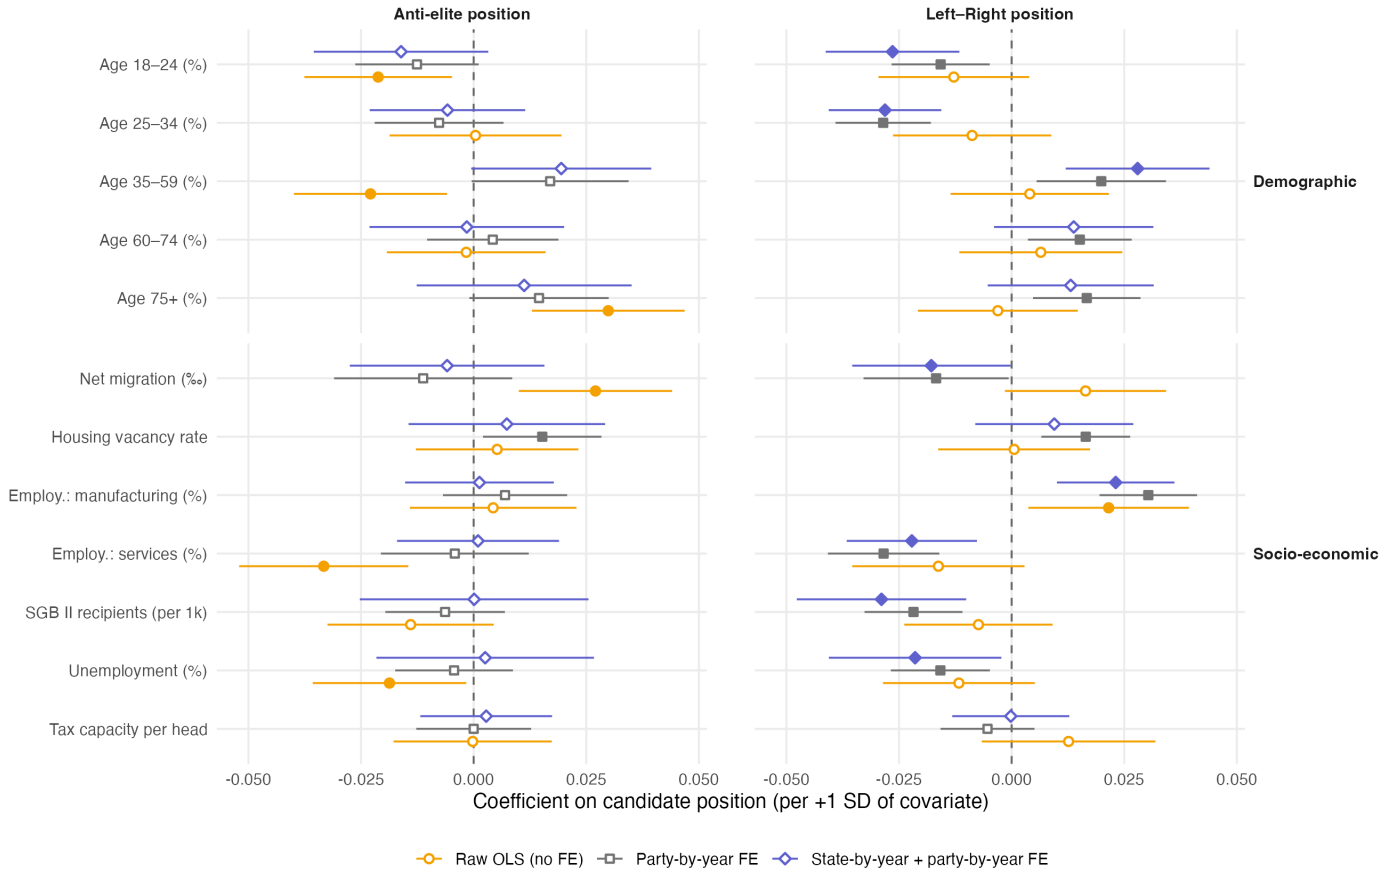


Figure I6. Association between candidates' positions on each dimension and their electoral districts' demographic and socio-economic characteristics.

Notes. Each marker is the coefficient from a separate regression of a candidate position on one characteristic. The unit of observation is an electoral district-by-party cell in the 2013–2025 elections. The sample consists of electoral districts used in the estimation of the model (see Sections 4.2 and 5 for details). All covariates are standardized to mean zero and unit standard deviation. The specifications are (i) Raw OLS with no fixed effects, (ii) party-by-year fixed effects, and (iii) state-by-year plus party-by-year fixed effects. Standard errors are clustered by district-by-year. Horizontal bars are 95% confidence intervals. Filled markers denote coefficients significant at the 5% level.

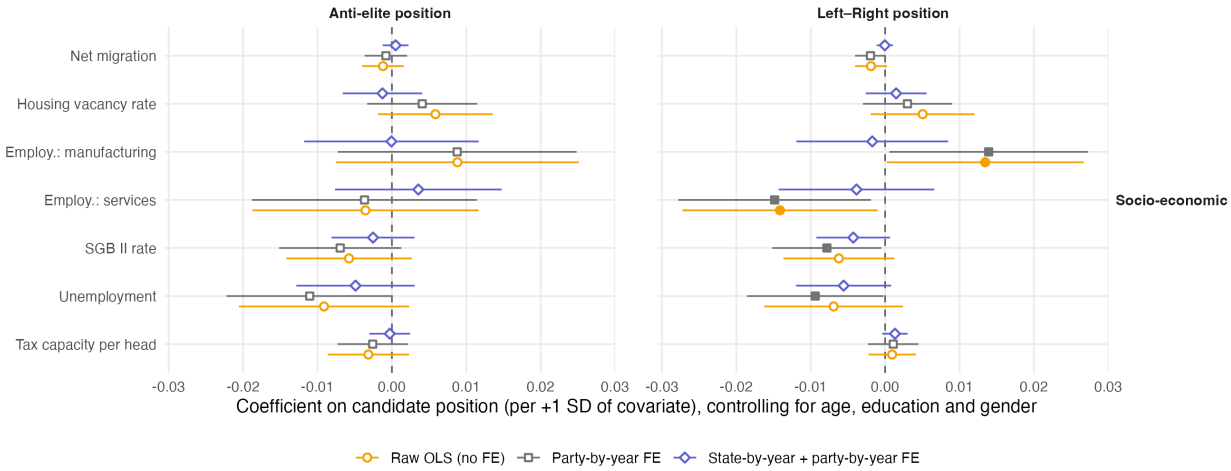


Figure I7. Association between candidates' positions on each dimension and municipalities' socio-economic characteristics, conditional on socio-demographic composition.

Notes. Each marker is the coefficient from a separate regression of a candidate position on one characteristic, controlling for the municipality's socio-demographic composition: the share of eligible voters with higher education, the female share, and the age-band shares. The unit of observation is a municipality-by-party cell in the 2013–2025 elections. The sample consists of municipalities used in the estimation of the model (see Sections 4.2 and 5 for details). All covariates are standardized to mean zero and unit standard deviation. The specifications are (i) OLS with no fixed effects, (ii) party-by-year fixed effects, and (iii) state-by-year plus party-by-year fixed effects. Standard errors are clustered by district-by-year. Horizontal bars are 95% confidence intervals. Filled markers denote coefficients significant at the 5% level.

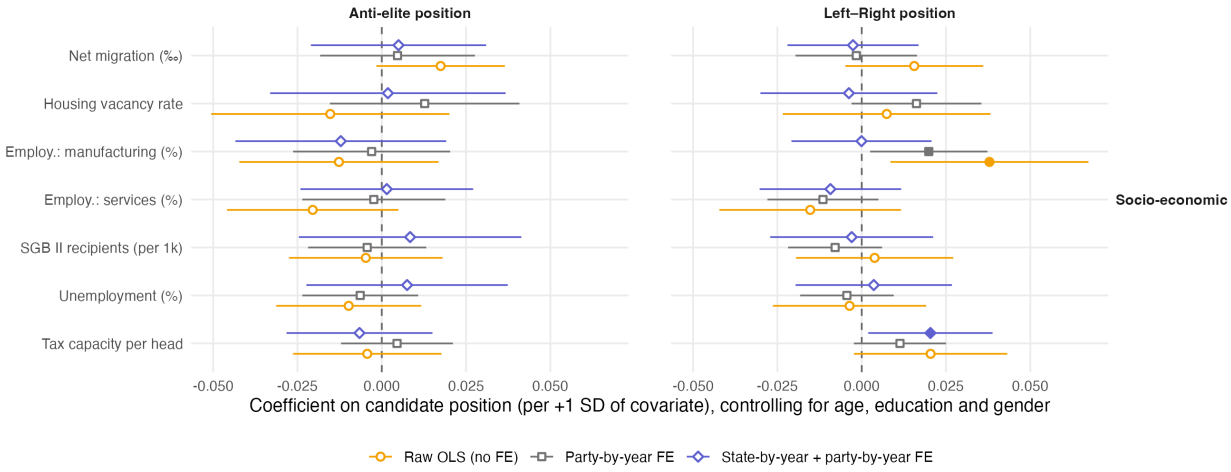
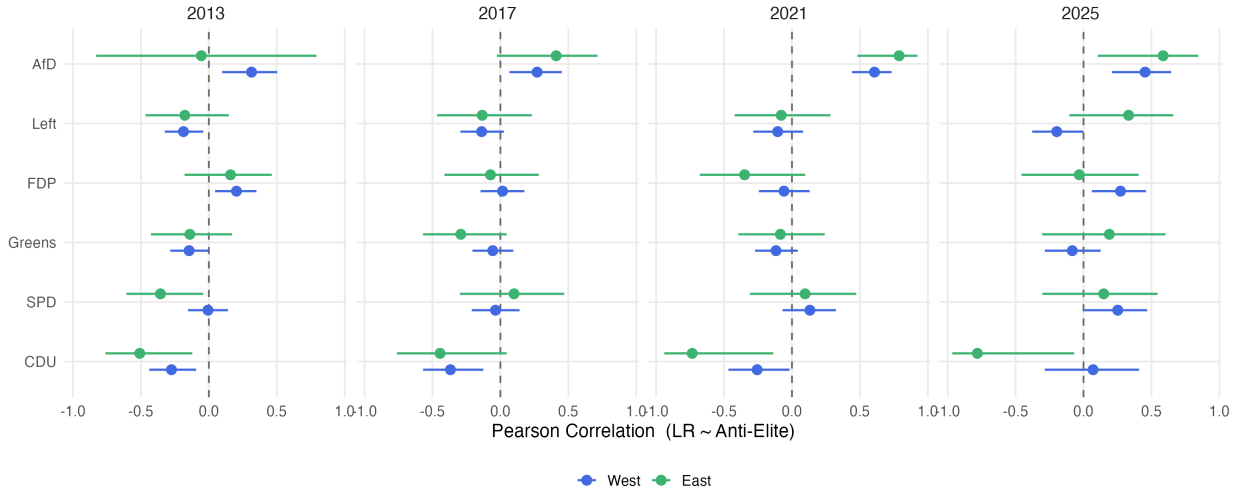
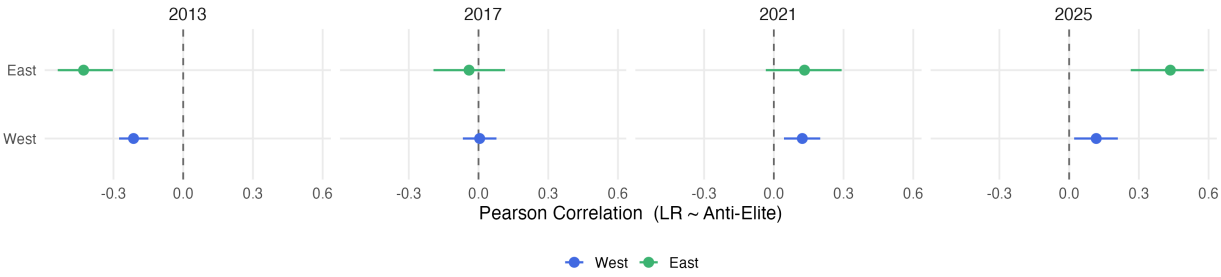


Figure I8. Association between candidates' positions on each dimension and their electoral districts' socio-economic characteristics, conditional on socio-demographic composition.

Notes. Each marker is the coefficient from a separate regression of a candidate position on one characteristic, controlling for the district's socio-demographic composition: the age-band shares of the population, and the share of eligible voters with higher education and the female share. The unit of observation is an electoral district-by-party cell in the 2013–2025 elections. The sample consists of electoral districts used in the estimation of the model (see Sections 4.2 and 5 for details). All covariates are standardized to mean zero and unit standard deviation. The specifications are (i) OLS with no fixed effects, (ii) party-by-year fixed effects, and (iii) state-by-year plus party-by-year fixed effects. Standard errors are clustered by district-by-year. Horizontal bars are 95% confidence intervals. Filled markers denote coefficients significant at the 5% level.



(a) Within parties



(b) Across parties

Figure I9. Correlation between candidates' Left–Right and Anti-Elite positions

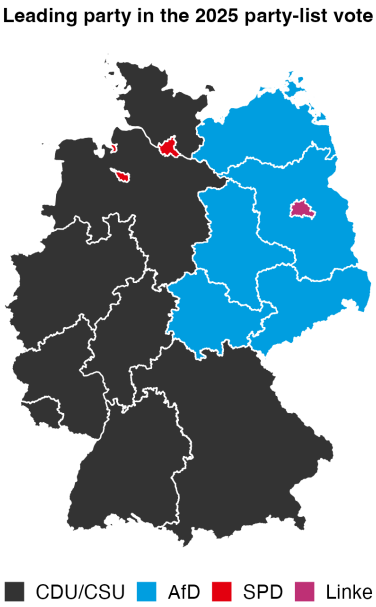


Figure I10. Leading party in the 2025 federal election, by federal state

Notes: Each federal state is colored by the party that received the largest share of party-list (second) votes in the 2025 Bundestag election. Vote counts come from the Official election returns. The AfD leads in all five East German states (Brandenburg, Mecklenburg-Vorpommern, Saxony, Saxony-Anhalt and Thuringia), with 32–39% of the vote and the CDU/CSU in second place at 18–20%. The CDU/CSU leads in every West German state except Berlin, Bremen and Hamburg.

J Additional Results

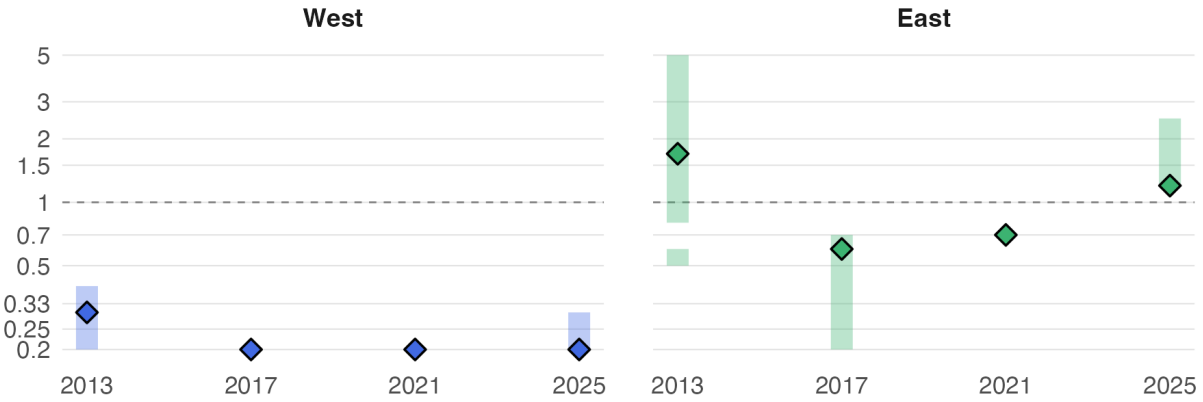


Figure J1. Sets of best-fitting weights on the Anti-Elite dimension, w_{AE}, relative to the Left–Right dimension ($w_{LR} = 1$): weights whose fit error is within 1% of the minimum, by region and year

Notes. $w_{LR} = 1$ is fixed throughout; w_{AE} should be interpreted as the relative weight placed on the Anti-Elite dimension. Points denote the point estimates corresponding to the best-fitting weight, shaded intervals denote weights within 1% of the minimum objective-function value. Separate exercises are run for East and West Germany. I allow ideal-point distributions to vary across federal states – that is, each socio-demographic group d in each state s has its own ideal-point distribution $f(x \mid s, d)$, while imposing the same weight w_{AE} for all groups within a given year and region. To run this exercise, the grid of possible weights is constructed. The grid is symmetric around 1 and implies relative weights across the two dimensions of $(0.2, 0.3, \ldots, 0.9, 1)$ and the corresponding inverse ratios. See Figure J3 for full results over the grid.

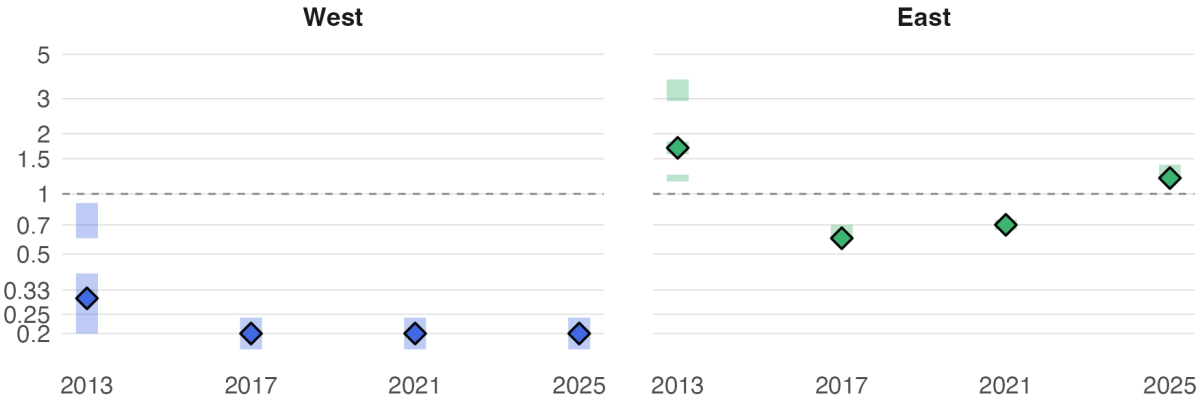
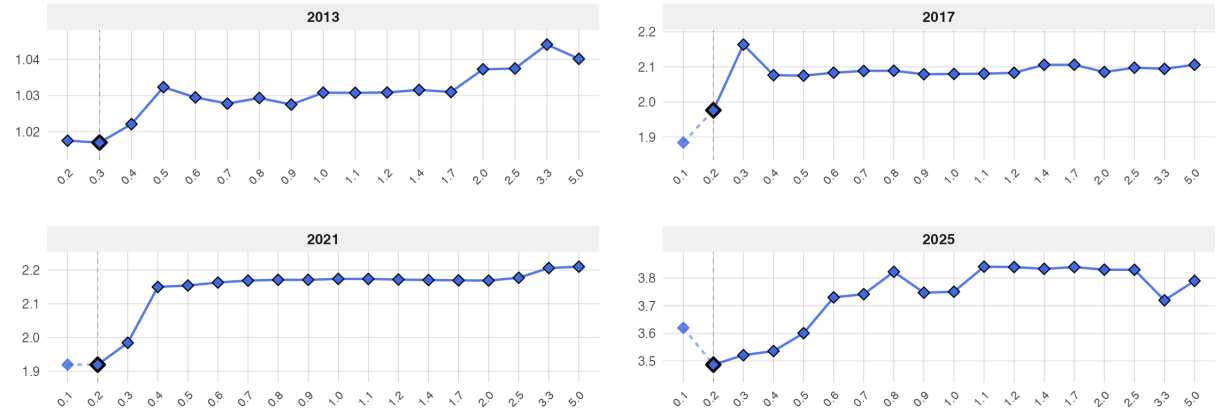
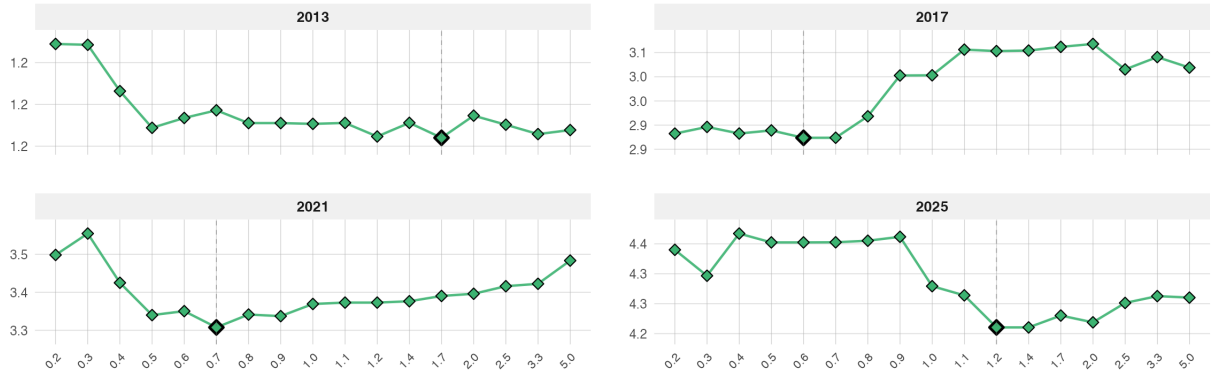


Figure J2. Confidence sets for the Anti-Elite weight w_{AE}, calibrated from the subsample re-runs

Notes. $w_{LR} = 1$ is fixed throughout; w_{AE} should be interpreted as the relative weight placed on the Anti-Elite dimension. Diamonds denote the point estimates corresponding to the best-fitting weight, identical to those in Figure J1. Shaded bars denote the confidence set CS, calibrated from the $B = 10$ subsample re-runs. Each draw's MAE curve is recentered on its own minimum, $D_{bj} = E_{bj} - \min_j E_{bj}$, which removes the level shift between draws and retains only the curvature; the cutoff is the $(1 - \alpha)$ empirical quantile (inverse-ECDF, $\alpha = 0.05$) of D_{bj^*}. Separate exercises are run for East and West Germany.



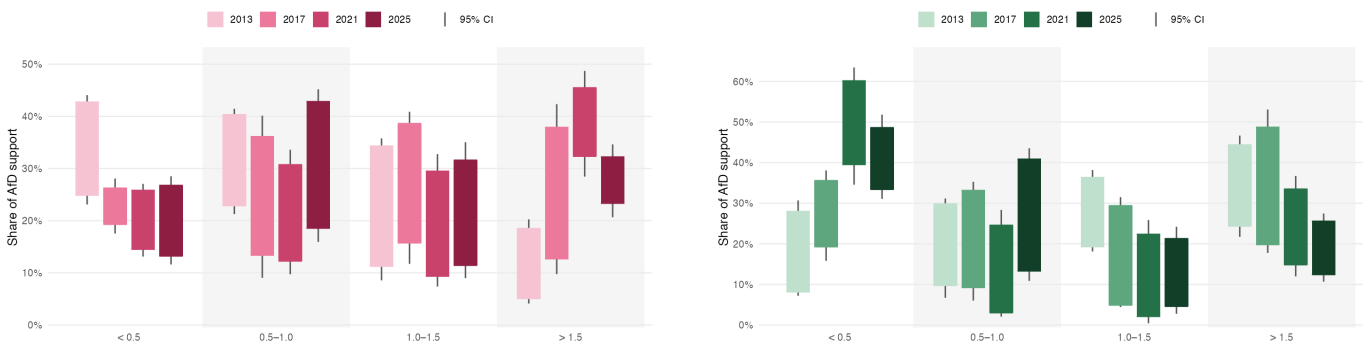
(a) West



(b) East

Figure J3. Model Error Value by different relative weight values, years and parts of the country

Notes. $w_{LR} = 1$ is fixed throughout; w_{AE} should be interpreted as the relative weight placed on the Anti-Elite dimension. Separate exercises are run for East and West Germany. I allow ideal-point distributions to vary across federal states – that is, each socio-demographic group d in each state s has its own ideal-point distribution $f(x \mid s, d)$, while imposing the same weight w_{AE} for all groups within a given year and region. Note that the weights are symmetric around 1 and imply relative weights across the two dimensions of $(0.2, 0.3, \ldots, 0.9, 1)$ and the corresponding inverse ratios. The grid is extended with a new point $w_{AE} = 0.1$ (where the corner $w_{AE} = 0.2$ is the best-fitting weight).



(a) Left–Right dimension (b) Anti-Elite dimension

Figure J4. Sampling uncertainty behind the dimension bins

Notes: Bars reproduce the bounds of Figure 8; thin gray lines are 95% Imbens and Manski (2004) confidence intervals. Endpoint standard errors come from a municipality bootstrap with $B = 50$ replications: municipalities are resampled with replacement within electoral districts, holding candidate positions and the partition fixed.

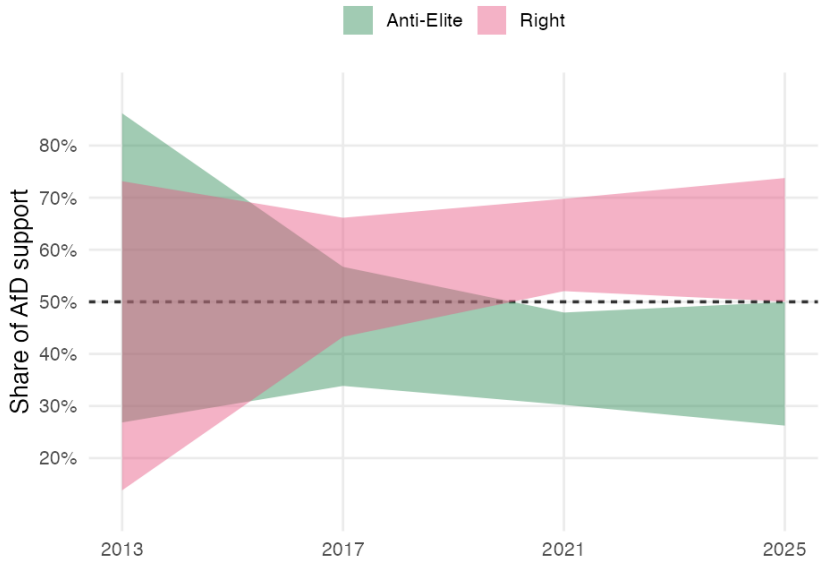


Figure J5. National bounds on the composition of AfD support with set-identified weight.

Notes. For each region and year, the anti-elite weight w_{AE} is treated as set-identified: every grid value whose fit is within 1% of the best fit is admissible (Figure J1). The bounds of the main specification are re-estimated at the smallest, middle and largest admissible weight, and the figure reports the smallest lower and the largest upper bound across admissible weights. Region-years whose admissible set is a single weight (East 2021, West 2017, West 2021) keep the main-specification bounds.



Figure J6. Bootstrapped 95% confidence intervals (Imbens–Manski), national

Notes: Bars show the identified set for θ_{Right} (Section 6.2.2); whiskers are 95% Imbens and Manski (2004) confidence intervals. Endpoint standard errors come from the municipality bootstrap with $B = 100$ replications.

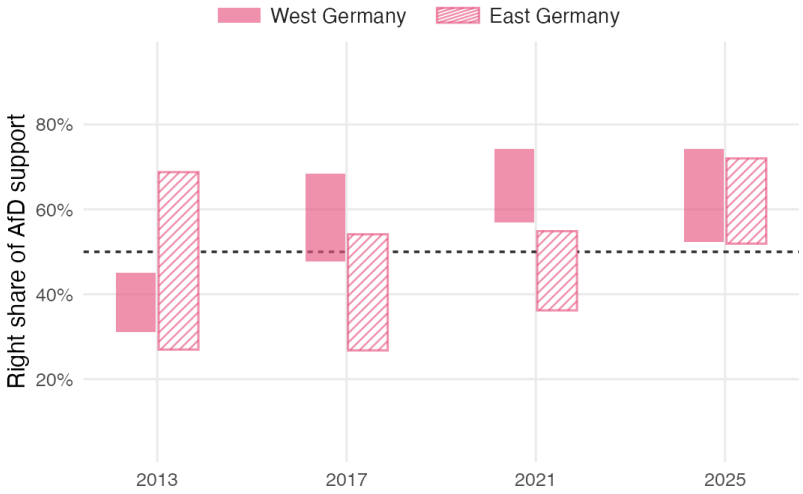
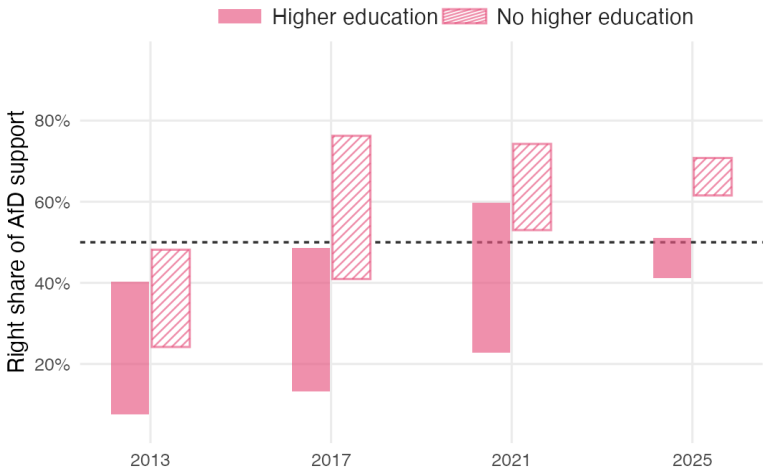
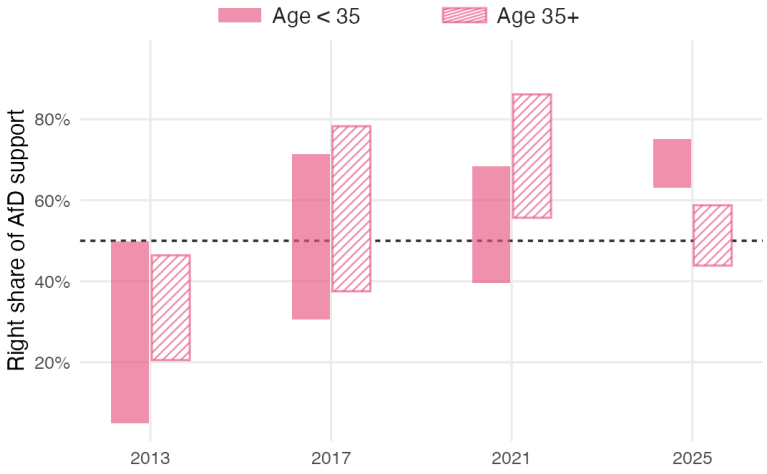


Figure J7. Regional Heterogeneity: Results separately for West and East Germany.

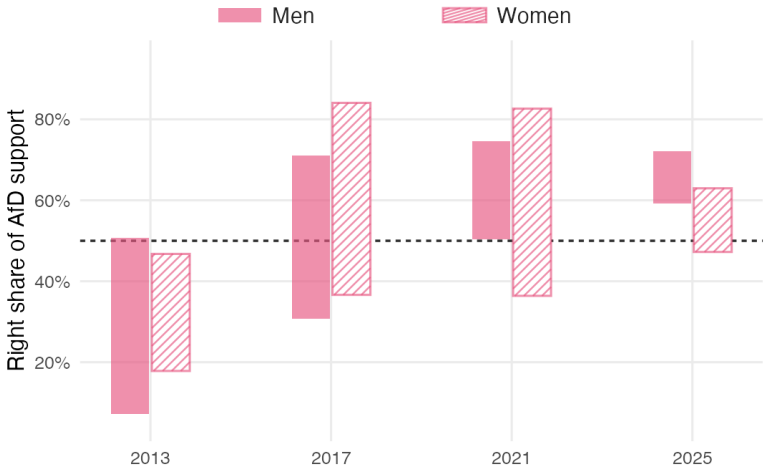
Notes: Bars plot bounds on the share of AfD voters whose ideal points lie in the attraction area of a hypothetical *Right* party that retains the AfD's median national position on the economic dimension while adopting a more pro-elite position; the remaining AfD support lies in the attraction area of the *Anti-Elite* party (see Figure 9 for an illustration). Each pair of bars compares West (solid) and East Germany (hatched) in a given year.



(a) By education



(b) By age



(c) By gender

Figure J8. Right share of AfD support by socio-demographic groups, Western Germany

Notes. Bounds on the share of each group's AfD voters whose ideal points lie in the pink region of Figure 9a across years are plotted; the anti-elite share is the complement. Each pair of bars compares the two groups in a given year. Results are for Western Germany only.

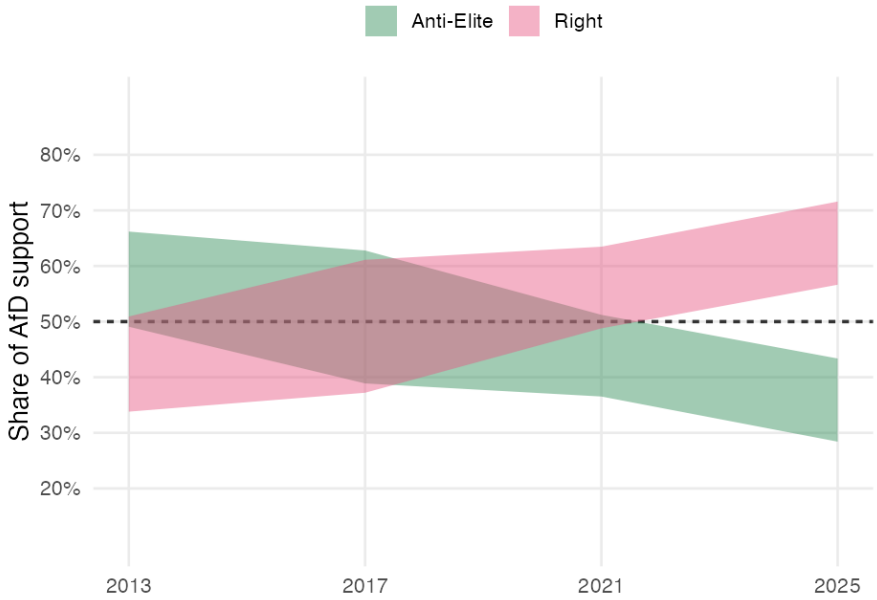


Figure J9. Estimated shares of AfD voters in both subspaces on the sample of municipalities with fewer than 5,000 inhabitants.

Notes. Bounds on the shares of voters whose ideal points lie in the green and pink regions in Figure 9a across years are plotted.



Figure J10. Estimated shares of AfD voters in both subspaces on the sample of non-contender electoral districts.

Notes: The figure plots bounds, across years, on the shares of voters whose ideal points lie in the green and pink regions of Figure 9a. Only districts in which the AfD candidate's vote share falls short of the winner's by more than 5 (10) p.p. in the East (West) are retained, for the years 2017 and 2021 (2025). These restrictions drop 12 districts in the East in 2017 and 28 in 2021, and 22 districts in the West in 2025. The East in 2013 and the West in 2013–2021 are omitted because there were no serious contenders, so the results from the main text stand; the East in 2025 is omitted because the AfD won in almost every district, making the analysis infeasible.

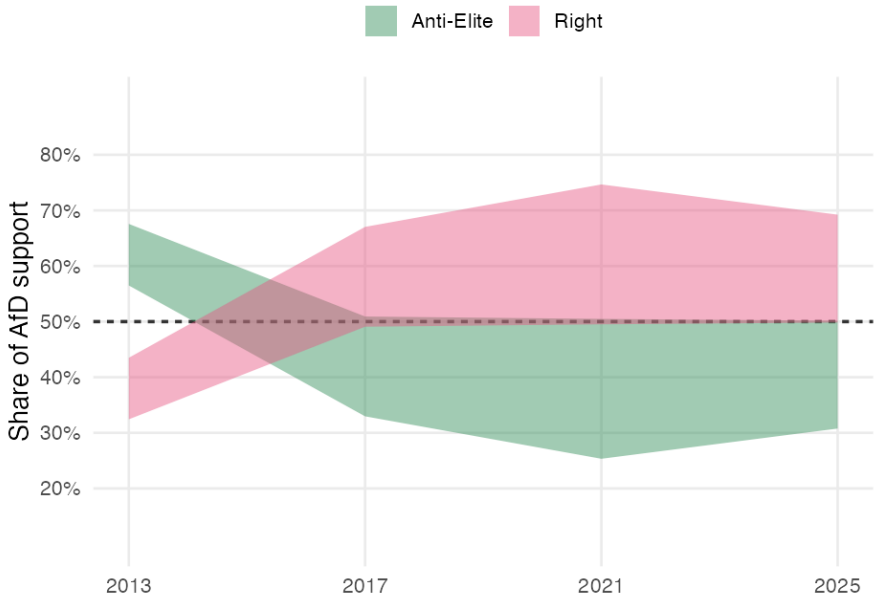


Figure J11. Estimated shares of AfD voters in both subspaces allowing for the differences in ideal points distributions between *conservative* and *moderate* electoral districts.

Notes: Bounds on the shares of voters whose ideal points lie in the green and pink regions in Figure 9a across years are plotted. The analysis allows the estimated ideal-point mass of each socio-demographic group, in each partition cell, to differ by up to a factor of three between conservative and moderate electoral districts of the same state, i.e. the conservative-district mass may range from one-third to three times the moderate-district mass in every cell. A district is classified as conservative if its AfD candidate is positioned above the within-state median AfD position on either the Left–Right or the Anti-Elite dimension, and moderate otherwise. This analysis is similar to the imperfect-instrument approach in Tebaldi, Torgovitsky and Yang (2023).

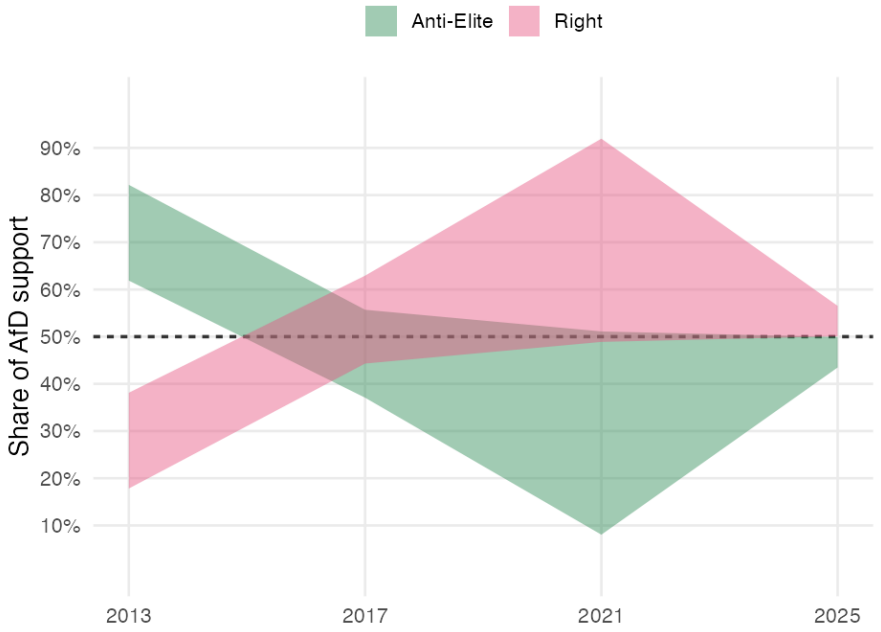


Figure J12. Estimated shares of AfD voters in both subspaces, estimated on the sample of border municipalities.

Notes. Bounds on the shares of voters whose ideal points lie in the green and pink regions in Figure 9a across years are plotted. The sample is restricted to municipalities adjacent to a border between two electoral districts within the same federal state. The preference distribution of each socio-demographic group is imposed to be identical across the municipalities of the same border pair, and is allowed to differ across border pairs within a state by up to 20%; across states the coupling is the same as in the baseline specification.

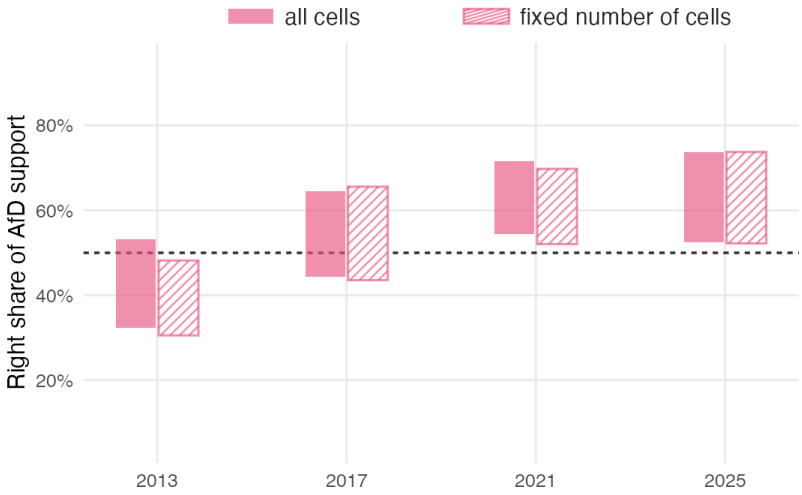
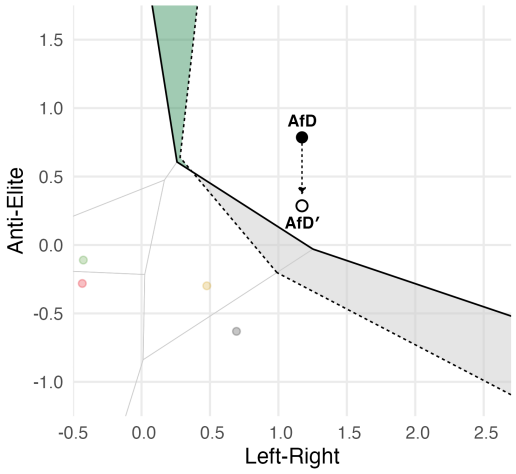
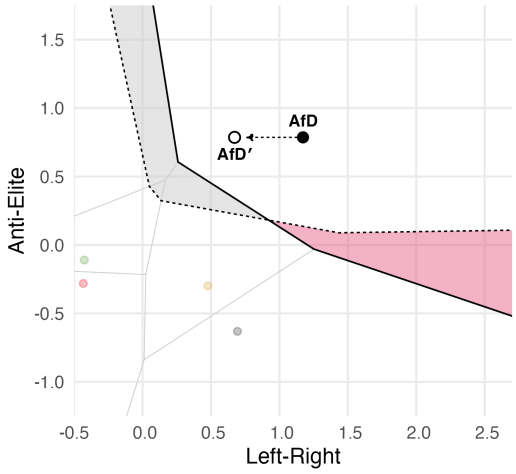


Figure J13. Bounds on the Right share of AfD support: full and merged partitions

Notes: Bars plot bounds on the share of AfD voters whose ideal points lie in the attraction area of a hypothetical *Right* party that retains the AfD's median national Left–Right position while adopting a more pro-elite position; the remaining AfD support lies in the attraction area of the *Anti-Elite* party (see Figure 9 for an illustration). Each pair of bars compares, within a year, the full Minimal Relevant Partition (solid; sharp bounds) with the merged partition of 1,000 cells used in the main text (hatched). See Appendix C.1 for details.



(a) AfD shift of 0.5 towards a more pro-elite position



(b) AfD shift of 0.5 towards a more left-wing position

Figure J14. Illustration of counterfactual moderating moves by the AfD by 0.5 units separately on two dimensions

Notes: The panels present an illustration of the unweighted cropped policy space, with partitions determined by median party positions across years. Shaded dots mark the median positions of parties other than the AfD. **AfD** denotes the AfD's median position across years, and **AfD'** denotes the shifted AfD position, either towards the pro-elite position (Panel (a)) or towards a more left-wing position (Panel (b)). Solid lines mark the existing partition, dashed lines the counterfactual partition. The illustrated shift is 0.5 units along each dimension. Since positions are standardized by construction, a 0.5-unit shift is comparable across dimensions. The exercise estimates the share of voters the AfD would lose under the shift, i.e. those located in the green or pink areas in each panel.

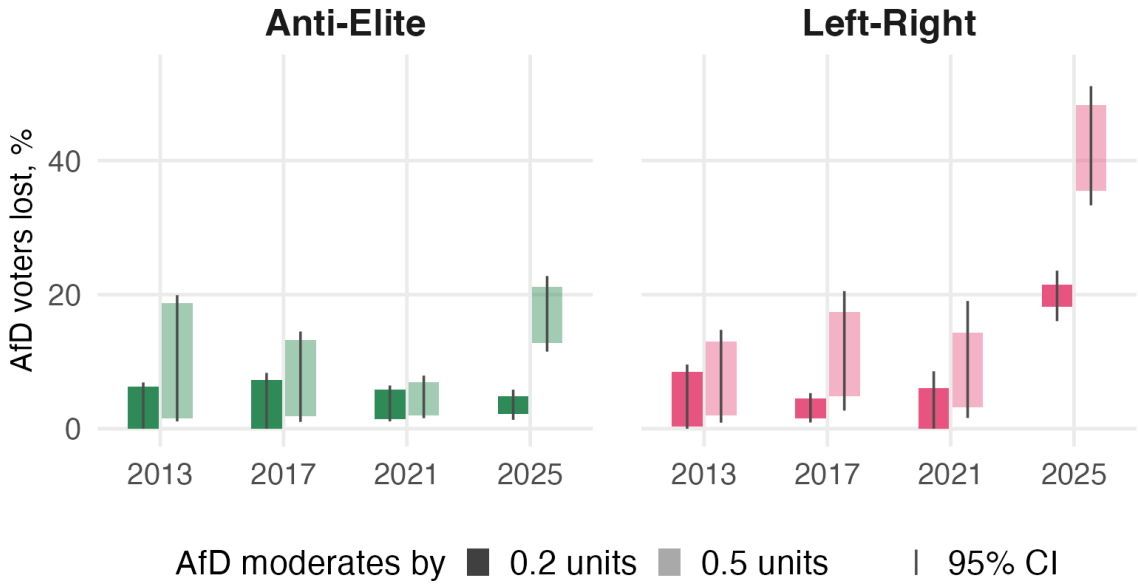
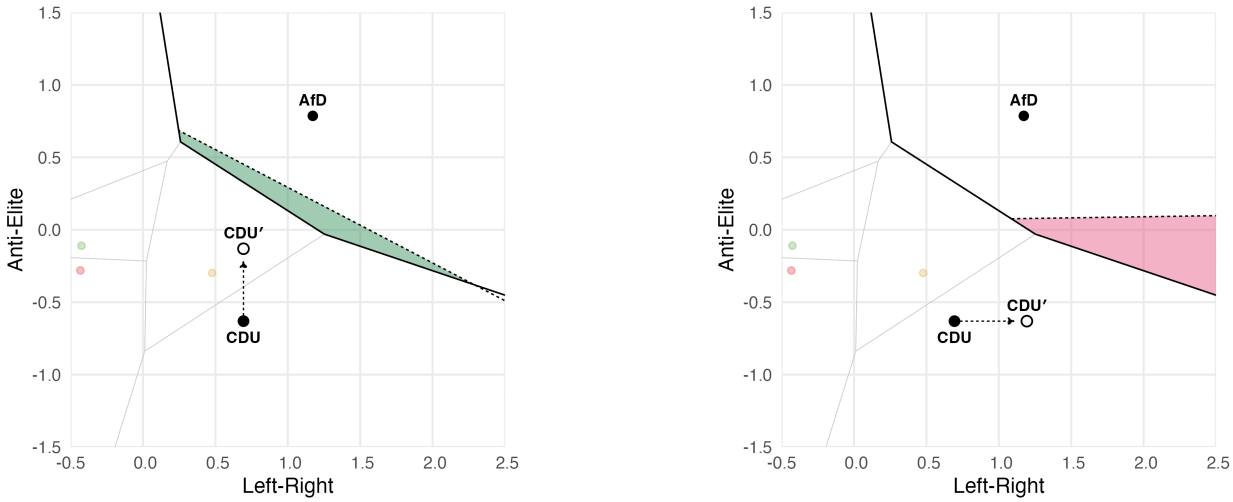


Figure J15. AfD moderates: share of AfD voters lost, identified sets and bootstrap confidence intervals

Notes. Bars are identified sets of the share of AfD voters lost when the AfD moderates by 0.2 or 0.5 units along the Anti-Elite (left) or Left–Right (right) dimension. Vertical lines are 95% Imbens and Manski (2004) confidence intervals from the municipality bootstrap with $B = 50$ replications, truncated at zero.



(a) CDU shift of 0.5 towards a more anti-elite position (b) CDU shift of 0.5 towards a more right-wing position

Figure J16. Illustration of counterfactual moves by CDU by 0.5 units separately on two dimensions

Notes: The panels present an illustration of the unweighted cropped policy space, with partitions determined by median party positions across years. Shaded dots mark the median positions of parties other than the AfD and CDU. **CDU** denotes the median CDU position across years, and **CDU'** denotes the shifted CDU position—either towards the anti-elite position (Panel (a)) or towards a more right-wing position (Panel (b)). Solid lines mark the existing partition, dashed lines the counterfactual partition. The illustrated shift is 0.5 units along each dimension. Since positions are standardized by construction, a 0.5-unit shift is comparable across dimensions. The exercise estimates the share of voters the AfD would lose under the shift, i.e. those located in the green or pink areas in each panel.

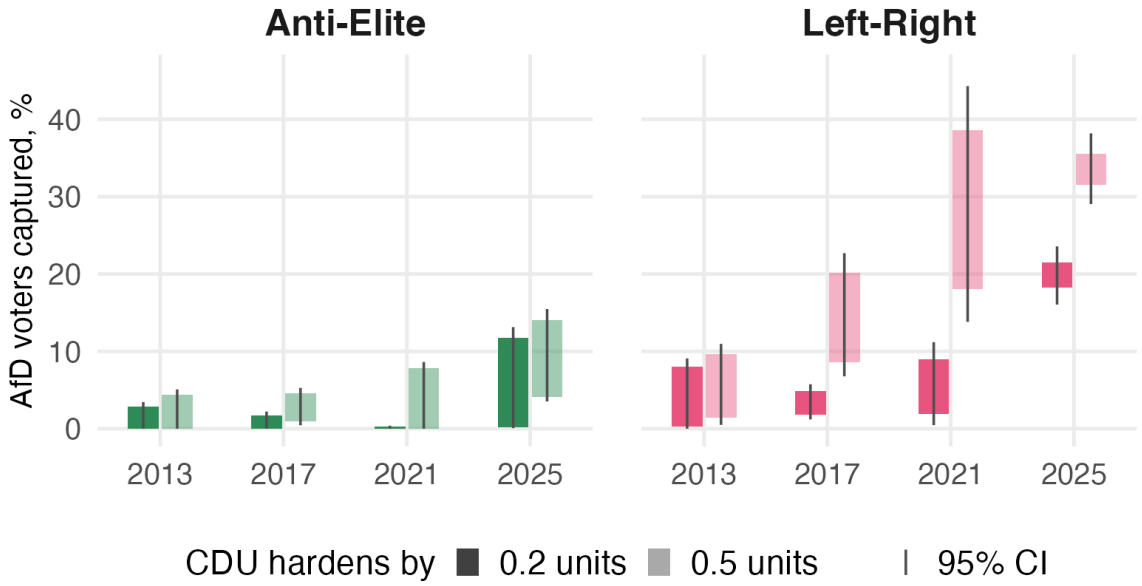


Figure J17. CDU hardens: share of AfD voters captured, identified sets and bootstrap confidence intervals

Notes. Bars are identified sets of the share of AfD voters captured when the CDU hardens its position by 0.2 or 0.5 units along the Anti-Elite (left) or Left–Right (right) dimension. Vertical lines are 95% Imbens and Manski (2004) confidence intervals from the municipality bootstrap with $B = 50$ replications, truncated at zero.

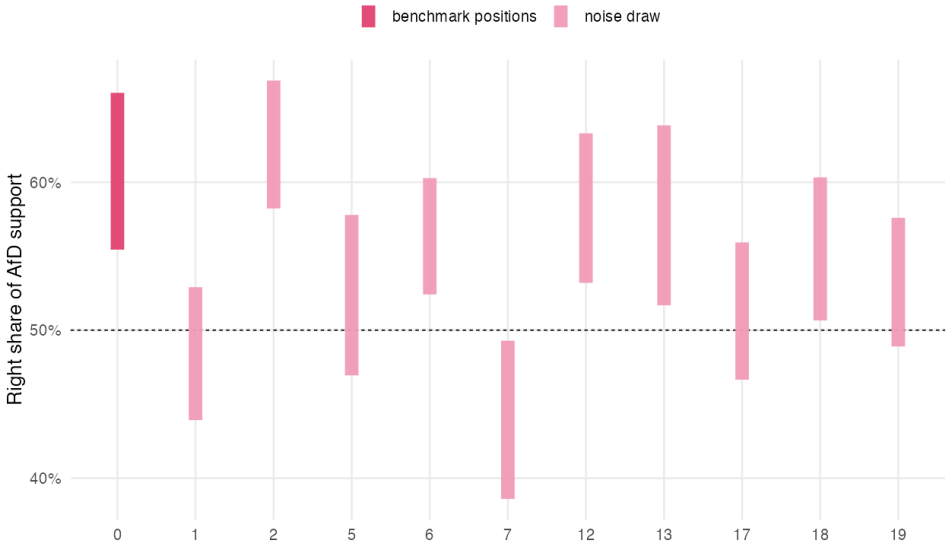


Figure J18. Robustness of the target parameter to measurement error in candidate positions

Notes. The figure shows the behavior of the target parameter from Section 6.2.2, estimated on the West 2025 data, when candidate positions are perturbed by measurement error. In each of 10 replications, an i.i.d. $N(0, \delta^2)$ shock with $\delta = 0.01$ is added to every estimated candidate position on both dimensions and to the anchor of the hypothetical parties, and the parameter is re-estimated; the resulting values are plotted against the baseline obtained without perturbation. Relative to the spread of the West 2025 district positions, δ equals $0.013 \times$ s.d.(Y_{LR}) (s.d. 0.76) and $0.021 \times$ s.d.(Y_{AE}) (s.d. 0.48); relative to the AfD's own candidates, $0.05 \times$ s.d. and $0.03 \times$ s.d. respectively.

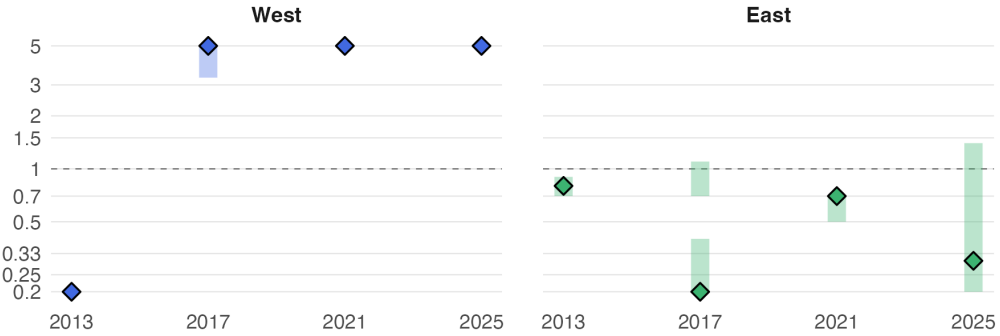
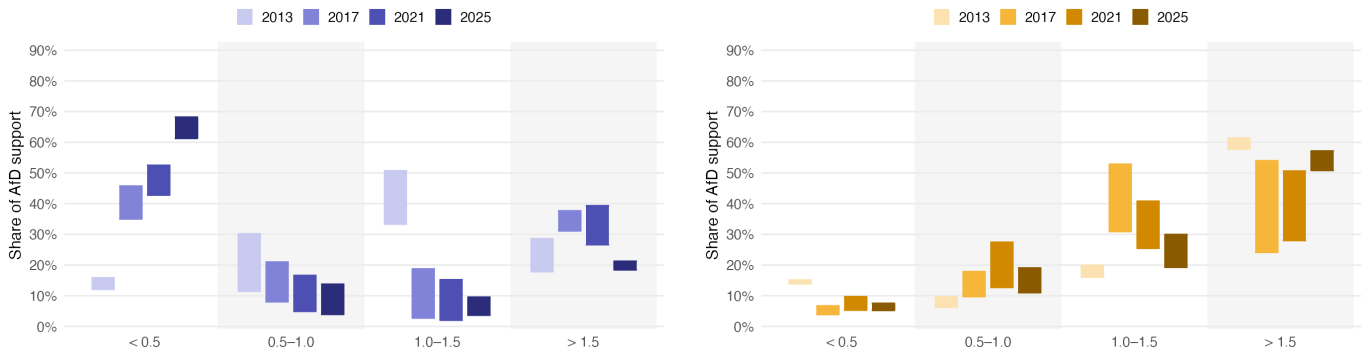


Figure J19. Weights for the Cultural dimension w_{cult} relative to Economic $w_{econ} = 1$ across years and parts of the country

Notes. $w_{econ} = 1$ is fixed throughout; w_{cult} should be interpreted as the relative weight placed on the cultural dimension. Separate exercises are run for East and West Germany. I allow ideal-point distributions to vary across federal states – that is, each socio-demographic group d in each state s has its own ideal-point distribution $f(x \mid s, d)$, while imposing the same weight w_{cult} for all groups within a given year and region. To run this exercise, the grid of possible weights is constructed. The grid is symmetric around 1 and implies relative weights across the two dimensions of $(0.2, 0.3, \ldots, 0.9, 1)$ and the corresponding inverse ratios.

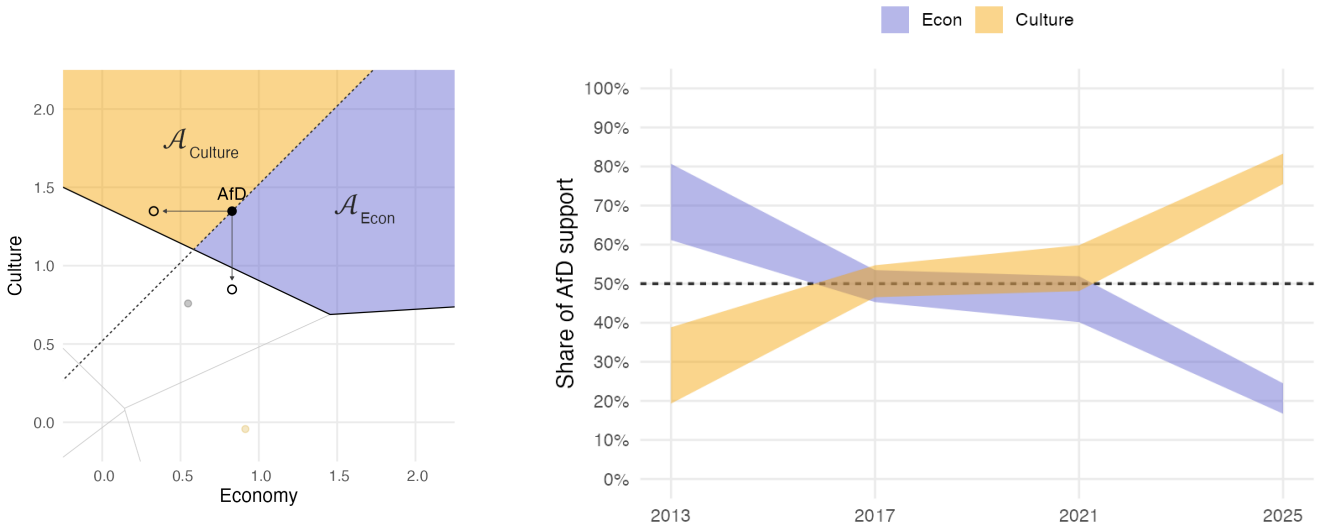


(a) Economics dimension

(b) Culture dimension

Figure J20. Distribution of AfD supporters' ideal points along the Economics and Culture dimensions, 2013–2025

Notes: The analogue of Figure 8 for the Economics × Culture split of the Left–Right dimension. Each panel plots bounds on the share of AfD voters whose ideal points lie in each 0.5-wide interval of one dimension, the other dimension left unrestricted; within each interval, bars run from 2013 (lightest) to 2025 (darkest). The strip below each panel places the intervals on the dimension, with the AfD's national median position by year (open circles) and the pooled medians of the other parties (faded dots).



(a) Illustration

(b) Estimated shares of voters preferring the hypothetical Economically Conservative vs. Culturally Conservative party among AfD voters

Figure J21. Illustration of a target parameter (a) and results (b)

Notes: Panel (a) shows party positions at their national medians; hollow circles mark the counterfactual AfD shifts. The shaded areas contain AfD voters motivated more by the party's economic stance (purple) or by its cultural stance (yellow). The illustration uses $w_{cult} = 1$; the estimation uses the weights estimated in Appendix Figure J19. Panel (b) plots bounds on the shares of AfD voters whose ideal points lie in the purple and yellow regions, by year.

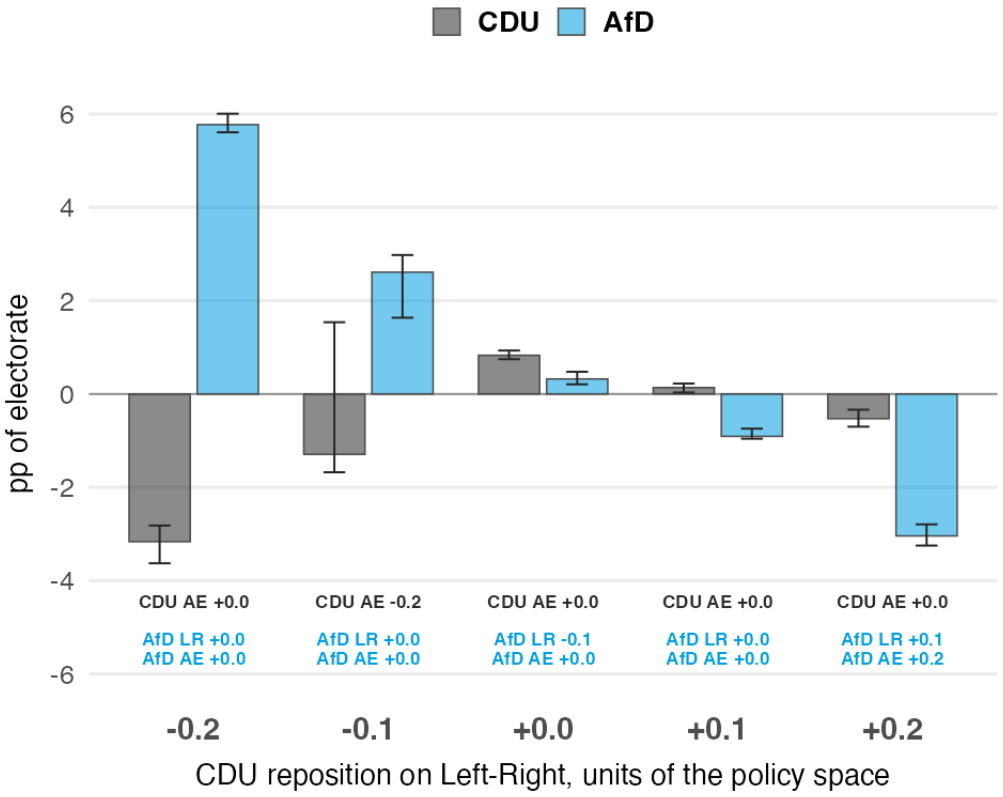


Figure J22. Estimated net CDU and AfD gain relative to the status quo by CDU Left–Right moves

Notes. Bars show the median net change in each party's vote share across 100 draws from the identified set, in percentage points of the vote share, after the AfD's best reply to each Left–Right move of the CDU (paired with the CDU's best Anti-Elite move in that draw); whiskers show interquartile ranges. National payoffs are population-weighted sums of the East and West changes. Both parties move on both Left–Right and Anti-Elite, in steps of 0.1 up to ±0.2. The labels under the bars give the CDU's Anti-Elite move and the AfD's modal reply across draws, as (Left–Right, Anti-Elite): the AfD replies $(-0.1, 0)$ to the CDU's status quo and $(+0.1, +0.2)$ to the CDU's +0.2 step, and stands still otherwise.

K Candidate Data

K.1 Candidate nominations across state lists and single-member districts

This section provides descriptive information on nomination patterns across state lists and single-member districts (SMDs). The data on all nominated politicians, with their names and party affiliations, were obtained by scraping the nonpartisan platform *Wen Wählen?* (*Whom to Vote For?*). I focus on candidates from the main parties considered in the paper (AfD, CDU/CSU, FDP, Left, Greens, SPD).

Under the mixed system, party nominations proceed through two channels: single-member district nomination and state-list nomination. The nomination of *single-member district* candidates is formally decentralized to constituency-level members' or delegates' assemblies, and state party actors do not nominate district candidates themselves, though they retain gatekeeping power. This is because the composition of state lists is fully controlled by state party actors, and the two channels are closely linked: about 64% of candidates who appear on state lists also run in an SMD, and about 74% of those who run in an SMD also appear on a state list (Table K1 reports the exact shares across parties and years). The state lists, chosen by the state leadership, therefore constrain the district-level choices made by local party units.

Candidates nominated in single-member districts also occupy higher list positions: across all years and lists their median position is 10, against 22 for list-only candidates. District strength predicts both access and rank. Candidates placed first or second in their district hold a list place in 84% of cases, against 68% of those placed third or lower; conditional on being listed, they stand at the 28th percentile of their list at the median, against the 44th for the remainder. The pattern also holds within lists: among candidates on the same state–party list, the top tercile of district vote share sits at a median position of 7, against 12 for the bottom tercile, and an additional percentage point of district vote share is worth 0.44 list places (s.e. 0.09).

I use these data to explore the patterns of candidate nomination and renomination to single-member districts. Across 2013, 2017, 2021, and 2025 there were 7,001 unique nominations at the candidate–election level.[53] Of these, 1,546 were repeat nominations from one electoral cycle to the next: the same candidate could be nominated in 2013 and 2017, in 2017 and 2021, or in 2021 and 2025. Repeat nominations account for 29.6% of all nominations in 2017, 2021, and 2025. Among them, only 91 (5.9%) are in a different electoral district, and only 9 (0.6%) in a different federal state.

Renomination patterns differ by electoral success. Among candidates who win their district, 56.3% are renominated, against 24.1% of those who lose; within each of the largest parties the rate rises with the first-vote share: from 17.1% in the bottom quartile of a party's candidates to 35.4% in the top, or 1.0 percentage points per percentage point of first-vote share (s.e. 0.08). Winning candidates who are renominated return to the same district in 99.6% of cases.[54] Among losing candidates who are renominated, 91.5% return to the same district.

The strong attachment of candidates to their electoral districts (and federal states) is due in part to where they reside at the time of the elections. To examine this, I use the question in the GLES Candidate Survey that asks Bundestag candidates how long they have lived in their electoral district. Participation in the survey is voluntary, so not all candidates respond, and the information is self-reported. The share of candidates born in the district in which they run was 28%, 16%, and 31% in 2013, 2017, and 2021,

[53] The maximum possible number is 7,176. It is slightly lower because some parties, chiefly the AfD and the Greens, did not field candidates in all districts in some years.

[54] There are two cases in which a candidate was nominated to a *different* district following a change in district boundaries. Albert Helmut Weiler won the district of Gera–Jena–Saale-Holzland-Kreis in 2013, which was dissolved before the 2017 elections; he was nominated in, and won, the district of Saalfeld-Rudolstadt–Saale-Holzland-Kreis–Saale-Orla-Kreis, which was partly composed of the district in which he had run in 2013. Kay-Uwe Ziegler won the district of Anhalt in 2021, which was merged into Anhalt–Dessau–Wittenberg before the 2025 elections; he was nominated in, and won, the neighboring district of Mansfeld, whose boundaries were redrawn at the same time.

respectively, and the mean number of years of residence in the district was 25.0, 25.3, and 25.9 across the same years. About 20% of candidates had lived in their district for less than ten years in every year. While most candidates thus have long-standing ties to their district, a non-negligible share moved not long before the elections, possibly for reasons related to their political careers.

Table K1. Overlap between state-list and single-member-district nominations

	2013	2017	2021	2025
SPD	69%	62%	67%	65%
Left	76%	71%	71%	77%
Greens	76%	77%	64%	78%
FDP	80%	80%	77%	82%
CSU	41%	41%	25%	28%
CDU	52%	47%	46%	46%
AfD	28%	57%	76%	71%
All	62%	63%	63%	66%

(a) Of state-list candidates, share also nominated in a district

	2013	2017	2021	2025
SPD	96%	99%	96%	97%
Left	48%	46%	46%	45%
Greens	72%	71%	85%	85%
FDP	86%	90%	88%	95%
CSU	64%	64%	50%	55%
CDU	81%	79%	81%	82%
AfD	49%	47%	45%	56%
All	73%	72%	73%	76%

(b) Of district candidates, share also nominated on a state list

Notes. The data are drawn from the nonpartisan candidate platform Wen Wählen?, which provides information on the candidate lists and state lists for federal elections. These data provide the universe of all party nominations. Panel (a) is computed over state-list candidacies, panel (b) over district candidacies, matched by candidate name and year. The Piraten are excluded from both panels, including the *All* rows.

K.2 Estimation of candidates' positions

1. Estimation of candidates' positions using GLES Candidate Survey Data Candidate positions on two dimensions, left–right and anti-elite, are estimated from the GLES Candidate Studies for 2013, 2017, 2021 and 2025. For each dimension, a set of survey items is first harmonized (see Table K2 for the core items, which recur across waves; several year-specific items are added in individual waves). Each item is then z-scored using a pooled mean and standard deviation computed across all waves in which it appears. This strategy places all items on a common metric and preserves comparability across years.

Each candidate's score on a dimension is the unweighted mean of all available z-scored items within that dimension. The anti-elite dimension comprises up to seven questions (five in 2025) on satisfaction with democracy and current institutions; the Left–Right dimension uses questions covering redistribution, state intervention, immigration, and social and cultural issues, such as abortion and same-sex marriage. It is also supplemented by a left–right self-placement scale.

Because of the low response rate to this survey (around 35% across all years and parties), the survey was complemented with a nonpartisan candidate survey, which also asks candidates about their positions on different issues in each electoral year.

2. Complementing GLES Candidate Survey Data with the nonpartisan candidate survey *Wen Wählen?* For each election year, candidates from the GLES Candidate Survey are merged with candidate data from the nonpartisan platform *Wen Wählen?*. Because the two surveys cannot be linked at the level of the individual candidate, the merge is at the district–party level: each *Wen Wählen?* candidate is matched to the mean GLES-derived score of their party's candidates in their district. The GLES scores themselves come from the pooled cross-year normalization, so the ideological scale is comparable across waves. The elastic net procedure is then used to impute ideological scores for those candidates missing in the GLES Survey but present in *Wen Wählen?*. Since the *Wen Wählen?* questionnaire changes from election to election, the

models are estimated separately by wave (2013, 2017, 2021, 2025).

Two separate elastic net models are estimated per wave, one for each ideological dimension. The outcome y_i is the GLES-derived district–party score attached to candidate i. Only candidates whose district–party cell contains at least one GLES-surveyed candidate enter the training set.

The feature matrix $\mathbf{X}$ contains, for candidate i, the full battery of that wave's policy questions (65 in 2013, 76 in 2017, 69 in 2021, 50 in 2025; the same set enters both dimension models). Each question is encoded as one-hot indicators for the five response categories: *Yes*, *Rather Yes*, *Undecided*, *Rather No*, *No*, plus an explicit indicator for a missing answer. No year indicators are needed since each wave is fit separately.

The elastic net minimizes a penalized least-squares criterion that interpolates between Ridge (ℓ_2) and Lasso (ℓ_1) regularization:

$$\hat{\boldsymbol{\beta}}(\lambda, \alpha) \;=\; \arg\min_{\boldsymbol{\beta}} \frac{1}{n} \sum_{i=1}^{n} \big(y_i - \mathbf{x}_i^\top \boldsymbol{\beta}\big)^2 \;+\; \lambda \left[\alpha \sum_j |\beta_j| + \frac{1-\alpha}{2} \sum_j \beta_j^2 \right]$$

where n is the number of candidates in the training set, $\lambda \geq 0$ controls the overall amount of regularization, and $\alpha \in [0, 1]$ determines the mix between the two penalty types. Setting $\alpha = 1$ recovers the Lasso, which drives many coefficients to exactly zero (implicit variable selection); $\alpha = 0$ recovers Ridge, which shrinks all coefficients towards zero while retaining all predictors. Intermediate values of α combine both effects.

Both λ and α are selected jointly by 10-fold cross-validation on the labeled candidates (those with a GLES-based score). The procedure is:

1. The labeled training set is randomly partitioned into 10 roughly equal folds.
2. A full Cartesian grid of 550 hyperparameter combinations is constructed:
   - $\lambda \in \{10^{-4}, \ldots, 10^{2}\}$ (50 values, log-uniform spacing);
   - $\alpha \in \{0.0, 0.1, 0.2, \ldots, 1.0\}$ (11 values).
3. For each of the $550 \times 10 = 5\,500$ fold–hyperparameter combinations the fit is run exclusively on the nine training folds and then applied to the held-out validation fold.
4. The root mean squared error (RMSE) between predicted and observed scores is recorded for each held-out fold.
5. The pair $(\hat{\lambda}, \hat{\alpha})$ that minimizes the mean RMSE across the 10 folds is selected.

The winning model is then refit on the wave's full labeled training set using $(\hat{\lambda}, \hat{\alpha})$, and this final fit is used to predict scores for the wave's unlabeled candidates. The selected mixtures are Ridge-leaning ($\hat{\alpha}$ between 0.1 and 0.5 across waves and dimensions), retaining 20–119 non-zero coefficients per model. The most potent questions for each year and dimension are presented in Table K3. The most potent questions and directions are as expected given the underlying dimensions.

The model performs well, explaining a large share of variation in observed GLES-based Left–Right position scores (cross-validated $R^2 = 0.76$ pooled across waves), with a root mean squared error of 0.336 against a standard deviation of the target score of 0.689 (RMSE/s.d. = 0.49). This indicates that the elastic net predictions are able to capture meaningful within-party variation in candidate positions across districts.

Performance is somewhat lower for the anti-elite score: an RMSE of 0.450 against a standard deviation of 0.648 (RMSE/s.d. = 0.69) and a cross-validated $R^2 = 0.53$ pooled. The anti-elite score is harder to map from the policy-stance questions available in *Wen Wählen?*. Nonetheless, the model explains about half of the variance in observed anti-elite scores, a clear improvement over mean imputation.

To mitigate overfitting, all performance metrics are computed from out-of-fold cross-validation predictions. Across the four waves, 1,147 *Wen Wählen?* candidates carry a GLES-based training score and 1,348 receive an elastic-net imputed score. Combining the two candidate surveys yields a locally estimated position for about 60% of all district×party cells. Missing positions are imputed with the within-state-wave median; if that median is also unavailable, the national within-wave median serves as a fallback.[55]

Table K2. Survey questions used to estimate candidate positions (GLES Candidate Study)

Dimension	**Question**	**Values**
Anti-Elite	Satisfaction with democracy in Germany	1–4
Anti-Elite	Democracy offers enough opportunities for citizens to participate	1–5
Anti-Elite	Laws reflect citizens' interests	1–5
Anti-Elite	Political parties are the central intermediary between citizens and government	1–5
Anti-Elite	There should be nationwide referendums (direct votes) at the federal level	1–5
Anti-Elite	Referendums tend to lead to bad laws	1–5
Anti-Elite	Citizens have lost trust in politicians	1–5
	Economy:	
Left–Right	In a crisis, the government should spend to stimulate economic growth	1–5
Left–Right	Government should stay out of the economy	1–5
Left–Right	Providing social security should be a key goal of government	1–5
Left–Right	Government should reduce income differences	1–5
	Culture:	
Left–Right	Immigrants should adapt to German culture	1–5
Left–Right	Same-sex marriage should be banned	1–5
Left–Right	Immigration is good for the German economy	1–5
Left–Right	Women should be free to decide about abortion	1–5
Left–Right	Immigration policy: make immigration easier vs. restrict it	1–11
	Overall:	
Left–Right	Self-placement on the left–right scale	1–11

Notes. The table lists the core items used to construct the Left–Right and Anti-Elite scores. They recur across the waves of the GLES Candidate Study (2013–2025) with the same wording and scales, although not every item is fielded in every wave (e.g., in 2025 the Anti-Elite score is built from five of the seven items). In addition, each wave includes several year-specific items that enter the Left–Right score (e.g., government spending on development aid and on defense in 2017, and the taxes vs. social spending scale in 2025). All items are standardized with their pooled mean and standard deviation across the waves in which they appear.

[55] The national median fallback is applied in less than 2% of the electoral districts used in the estimation, and only in Eastern Germany in 2013 and 2017.

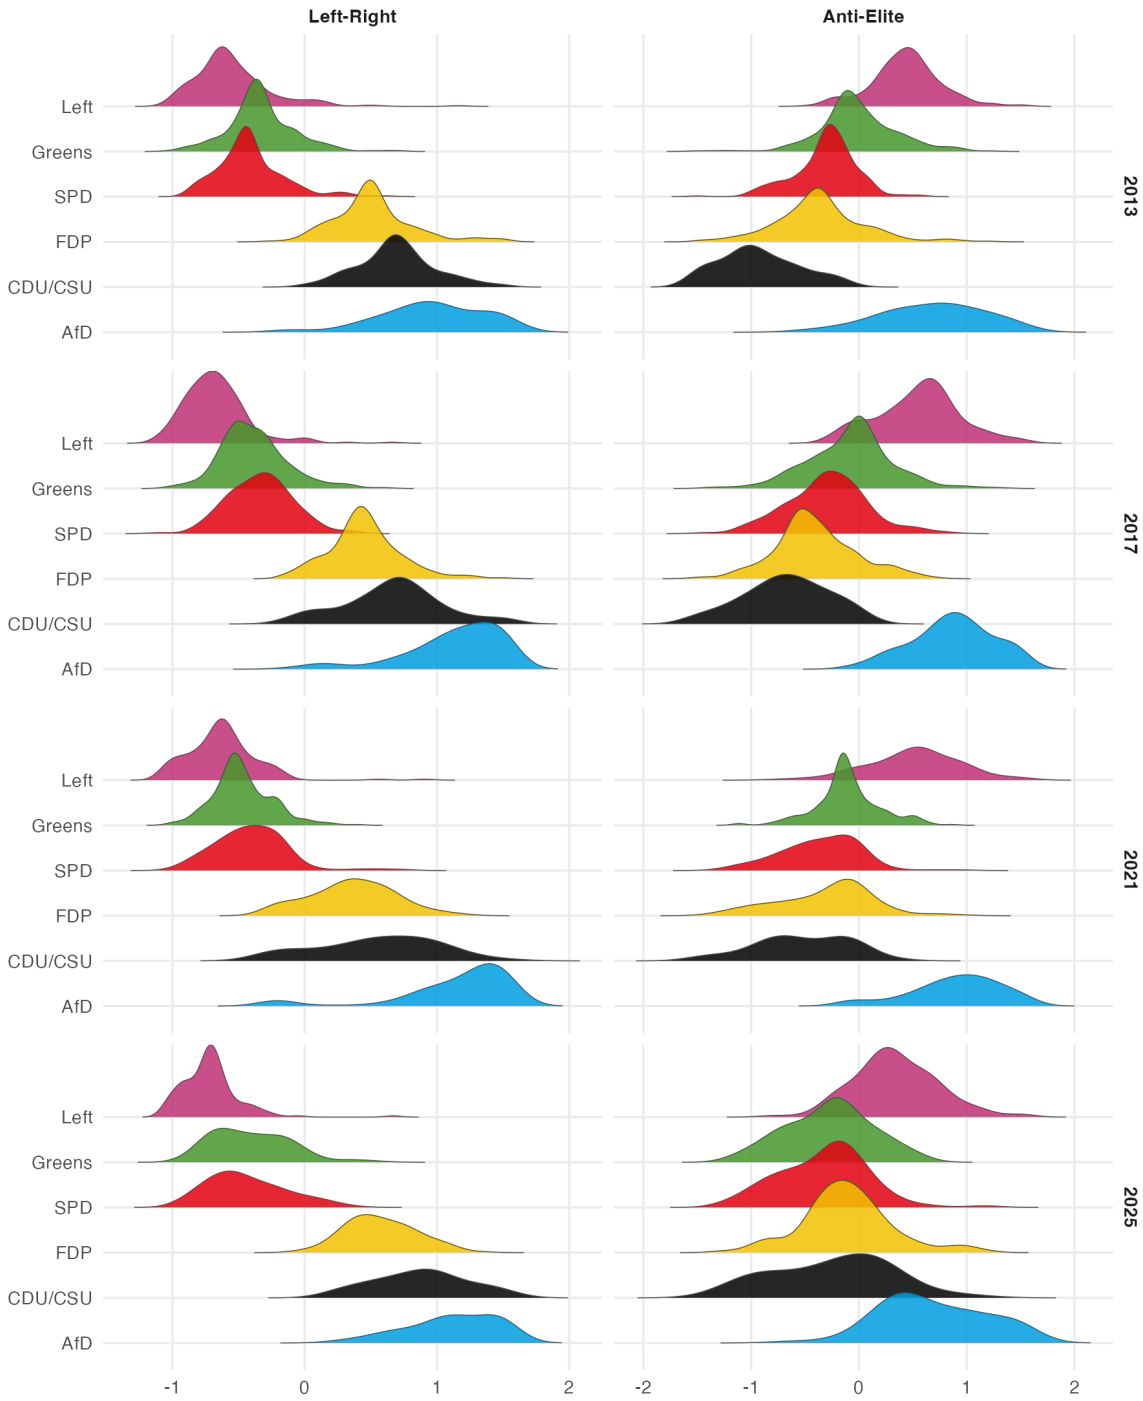


Figure K1. Distribution of candidate positions by party and election year

Notes. Kernel density estimates of the positions of single-member district candidates by party, separately for each election (rows) and dimension. Positions are estimated from the GLES Candidate Survey and the nonpartisan platform *Wen Wählen?*. Each density uses a Gaussian kernel with the bandwidth selected by Scott's rule computed separately for each party, year and dimension.

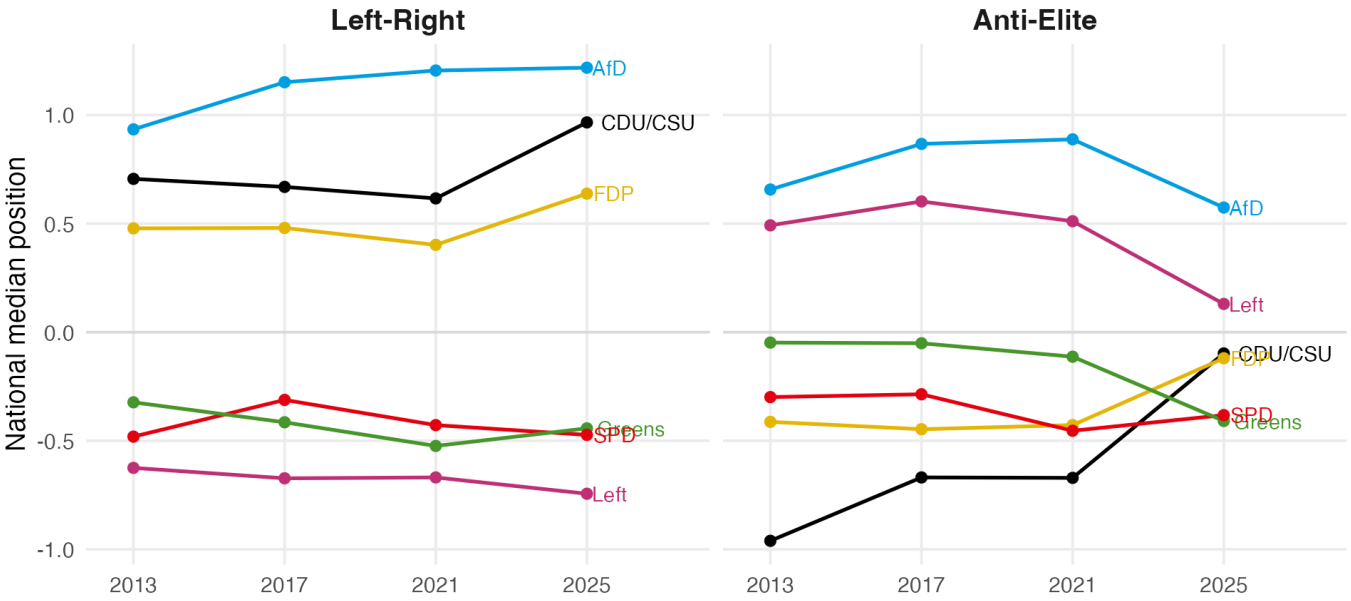


Figure K2. Medians of left–right and anti-elite positions across parties over years

Notes. The data are based on the estimated single-member district candidate positions from the GLES Candidate Survey and the nonpartisan platform *Wen Wählen?*. Within-party dynamics broadly follow the Chapel Hill Expert Survey (Bakker et al., 2015): for instance, the rightward and anti-elite shift of the CDU and the moderation of the AfD's anti-elite positioning between 2021 and 2025 are mirrored in the change between the 2019 and 2024 CHES waves.

Table K3. Most Predictive *Wen Wählen?* Items in the Elastic-Net Imputation

Year	Item (translated)	Answer	Coef.
Left–Right (positive = further right)			
2013	The welfare state is more important than reducing debt!	Yes!	−0.08
2017	Repeal "marriage for all"!	Yes!	+0.12
2021	Ban gender-inclusive language!	Yes!	+0.20
2025	Raise the minimum wage substantially!	Yes!	−0.39
Anti-Elite (positive = more anti-elite)			
2013	Referendums at the federal level!	Yes!	+0.13
2017	Referendums at the federal level!	Yes!	+0.25
2021	Build a common European army!	No!	+0.14
2025	Invest more in defense!	Rather yes	−0.17

Notes. For each election year and dimension, the *Wen Wählen?* item with the largest contribution to the variance of the imputed score, $\sum_a |\beta_a| \,\mathrm{sd}(D_a)$ over the item's answer dummies D_a; the answer with the largest single contribution is shown. Coef. is the elastic-net coefficient β_a on that answer dummy: the change in the predicted GLES score for a candidate giving this answer rather than any other, holding all other answers fixed. One elastic net per wave and dimension, λ and α chosen by 10-fold cross-validation; "No answer", "Undecided", missing and pooled-rare levels are excluded from the ranking.

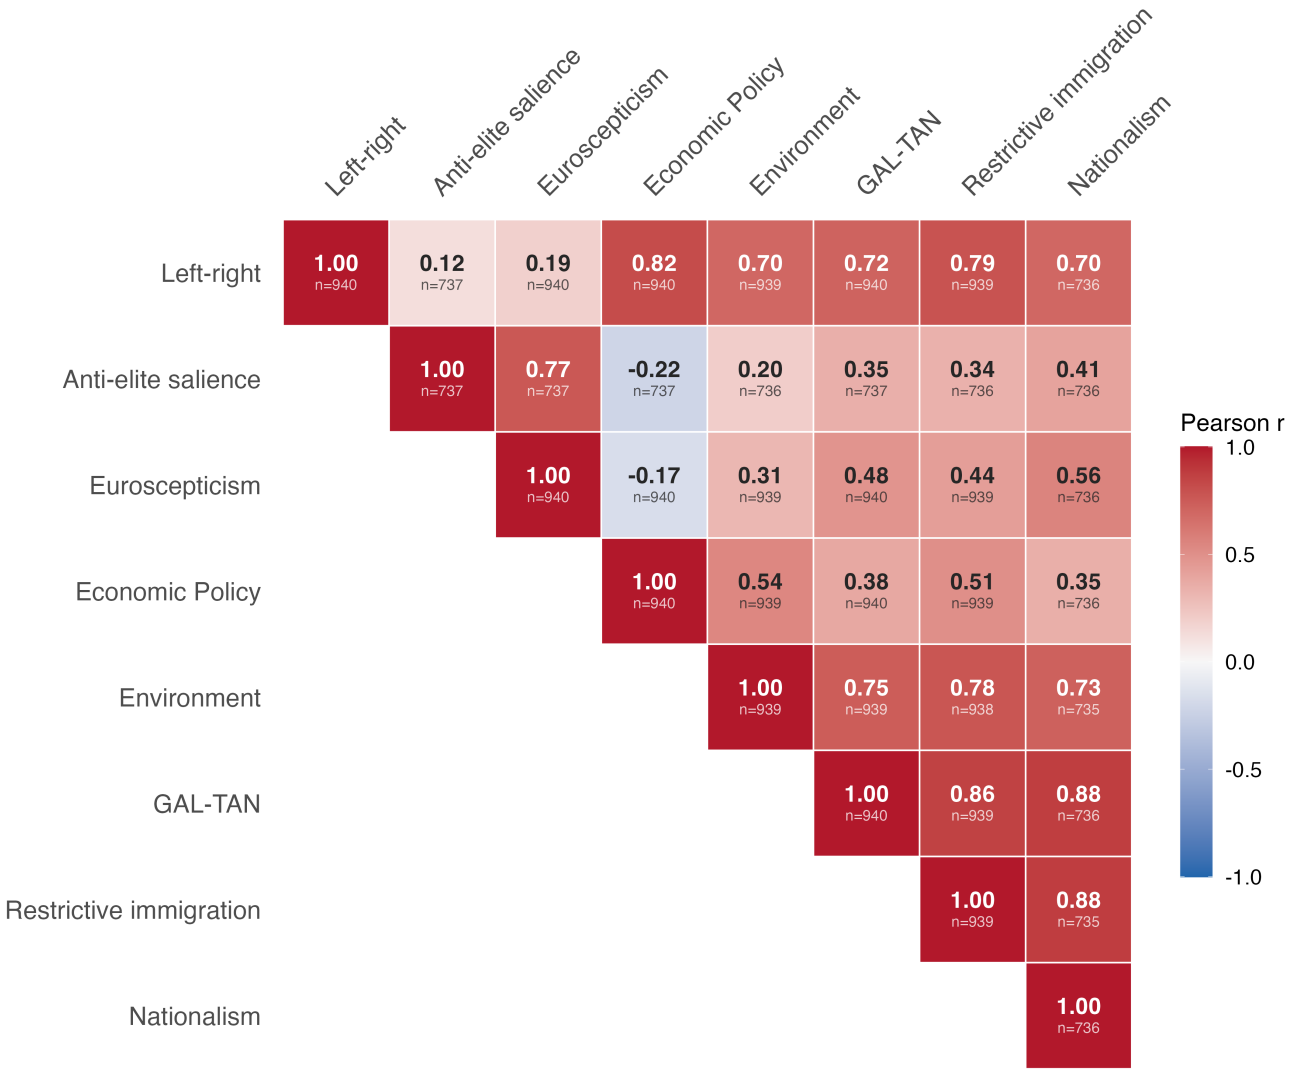


Figure K3. Correlations between expert-estimated party positions across policy dimensions

Notes. Each cell reports the pairwise-complete Pearson correlation between two expert-estimated party-position dimensions, computed across party–year observations. Data are the Chapel Hill Expert Survey 1999–2024 (party-level expert means), restricted to European countries and to the waves from 2010 onwards ($N = 940$ party–year observations, 408 parties, 28 countries). All CHES policy scales run 0–10 and are coded so that higher values are more right, conservative or restrictive; *Economic Policy* is the solely economic left–right scale and *Euroscepticism* reverses the EU-position item. *Anti-elite salience* and *Nationalism* were not fielded in the 2010 wave. Selected correlations (95% Fisher-z CIs in brackets): Left–Right $\sim$ Anti-elite salience 0.12 $[0.05, 0.19]$; Left–Right $\sim$ Euroscepticism 0.19 $[0.13, 0.25]$; Anti-elite salience $\sim$ Euroscepticism 0.77 $[0.74, 0.80]$.