Residential Sorting and the Reproduction of Political Geography

Lara Andermann Mathias Bühler*

August 5, 2026

Abstract

Historical persistence is commonly attributed to durable institutions or cultural transmission. We identify residential sorting as a distinct mechanism, linking household heads across Munich's directories, 1893–1961. We show that neighborhood occupational composition persists although wartime destruction and turnover left under one percent of 1912 residents at the same address by 1953. In each period, the composition is regenerated by newcomers choosing neighborhoods matching their own type. Sorting remains strong even among tenants pushed out by ownership changes, but strong sorting does not rule out that places convert whoever they receive. To separate the two, we use post-war administrative housing assignment, which replaced residential choice for displaced households, and find that composition persistence falls by roughly a quarter where more residents were assigned. The regenerated composition, in turn, predicts precinct voting across regime changes and population replacement. Political persistence is stronger where residents chose their neighborhoods than where the housing office assigned them.

JEL Codes: D72, R23, N93, N94
Keywords: Political persistence, residential sorting, neighborhood effects, voting behavior

*Ludwig-Maximilians-Universität Munich. Email: mathias.buehler@econ.lmu.de and lara.andermann@econ.lmu.de. We thank Davide Cantoni, Andrew Dickens, Joris Müller, and Jonathan Norris for insightful comments on early versions, and seminar participants at ASREC, the Big Data Seminar at Ifo, the CES Workshop on Microdata in Economic History, CRC Berlin Schwanenwerder, the University of Bayreuth, and internal seminars for helpful comments. Bühler acknowledges financial support from the LMUexcellent Investment Fund and from the European Research Council (ERC) under the European Union's Horizon Europe research and innovation programme (grant agreement no. 101044546, "CityRising"). Andermann acknowledges financial support from the Deutsche Forschungsgemeinschaft through CRC TRR 190 (project no. 280092119). Stefano von Igel provided excellent research assistance.

1 Introduction

Historical persistence is among the most robust findings in political economy (Siegfried, 1913; Voigtländer and Voth, 2012; Acharya et al., 2016). Existing explanations run through institutions that outlive the regimes that built them (Acemoglu et al., 2001) or cultural transmission across generations (Bisin and Verdier, 2001; Becker et al., 2020; Fouka and Voth, 2023). We show that newcomers sort into neighborhoods matching their own type, so part of what the literature credits to surviving institutions or transmitted culture is instead selection.

Isolating sorting empirically is difficult because persistence is observed in aggregate outcomes, whereas sorting is an individual decision (Tiebout, 1956; Bayer and Timmins, 2007). Aggregate data cannot distinguish whether a place's political character persists because institutions and local culture shape residents or because residents choose places matching their own characteristics. While the idea that sorting contributes to persistence is not new (Cantoni and Yuchtman, 2021), existing data rarely observe residential choice directly. Redding and Sturm (2024), for example, use wartime destruction to identify neighborhood effects in London from building-level and census records, showing how places shape residents but not whether residents chose those places. Ochsner and Roesel (2020) show the converse, that migrants carry their politics with them, but measure the migration only in the aggregate.

Munich's directories, by contrast, name every household head, their occupation, and their address. Linking the same individuals across yearly editions, we observe each residential choice in the year it is made, and with it the neighborhood as it stood when the mover chose it. By 1953, fewer than 1% of the city's 1912 residents still lived at the same address, and wartime bombing had destroyed roughly a third of its housing (Haus der Bayerischen Geschichte, 2024). Yet a precinct's vote in 1912 predicts its vote in 1961, three regime changes later, with a rank correlation near 0.8 (Figure 1). Its occupational composition, the average occupational standing of its household heads, persists just as strongly. With the people gone, newcomers must have rebuilt the pattern by sorting onto durable features of

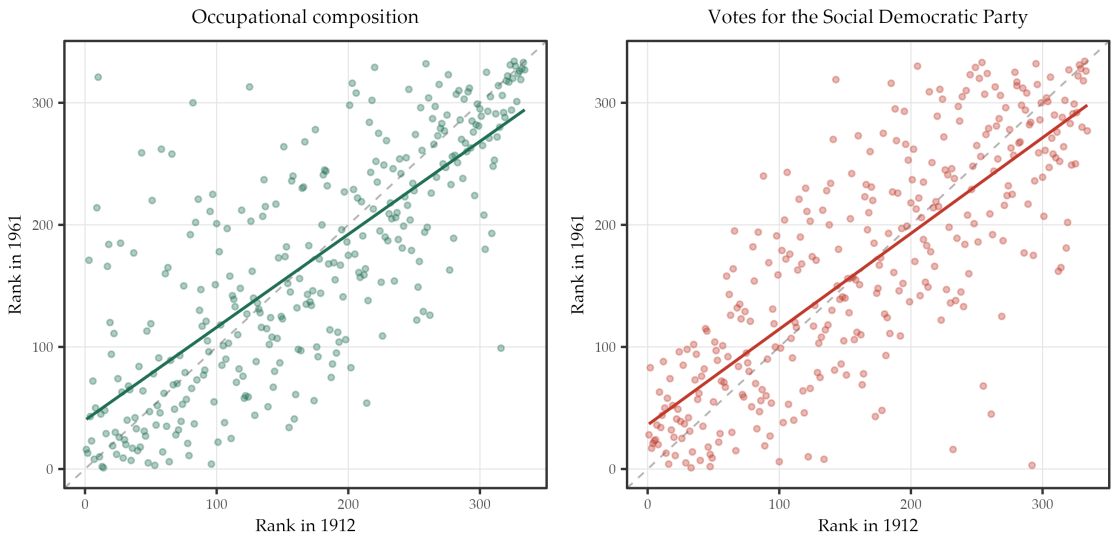

Figure 1: Occupational composition and the SPD vote share, 1912–1961

Notes: Each panel ranks Munich's 334 precincts within year and plots the 1912 rank against the 1961 rank. The left panel shows the occupational composition measured using the mean occupational standing of a precinct's household heads, lowest to highest. The right panel shows the Social Democratic Party's (SPD) vote share, lowest to highest. Dashed lines mark perfect persistence and solid lines are linear fits. Across the half century, the rank correlation is 0.76 for composition and 0.78 for the SPD vote share. Within each year, higher-status precincts vote less SPD.

place.

The Catholic milieu has been shown to be such a durable feature (Spenkuch and Tillmann, 2018; Becker and Voth, 2023; Braun and Öztürk, 2026): it does not need the people to stay, and it survived both the Nazi regime and the war. Its parishes stayed in place, and its bloc re-formed after the war as the Christian Social Union (CSU). The socialist milieu had no such continuity, as the 1933 seizure of power destroyed its party, its unions, and its associations. Yet which precincts vote strongly for SPD survived the regime and the war at least as strongly as which vote Catholic. Institutional continuity is therefore not necessary to explain persistence.

Transmission by residents is also implausible, because passing on local norms needs incumbents who stay, but few residents stayed long enough at the same address. Across the Weimar Republic's fifteen yearly directories, a precinct shed between a third and half of its residents from one year to the next. Of the renters listed in 1933, only 3% had lived in

their precinct in 1912, the anchor year of the voting map. Yet each arriving cohort chose its precinct by matching the type of the residents already there, and so came to resemble the stock it replaced without ever being taught by it or absorbing their norms. Transmission requires residents who stay long enough to teach; sorting requires only that residents of a similar type already be present when a newcomer chooses where to live.

The durable feature is then a precinct's occupational composition, which outlasts its residents. We show that newcomers sorted into neighborhoods according to the occupations of existing residents, and the regenerated composition predicts how a precinct votes. A precinct one standard deviation higher in 1910 composition voted 5.7 percentage points less for the Social Democrats and 3.8 points more for the Catholic bloc in the Weimar elections, and 4.1 and 3.1 points respectively after the war. But this pattern is equally consistent with places shaping whoever arrives and with newcomers arriving already sorted. We show that the latter is the case: where people could not sort, both composition and voting persist less across decades.

Sorting is observed only when a household moves, and every move combines a pull and a push: something draws a household to a new address, and something drives it out of its old one. In our setting, the pull is the draw of a neighborhood matching the household's own occupational type, and the push comes from ownership changes. Following households from one address book to the next, we estimate a standard discrete choice model in which each mover chooses among all precincts, described by their average occupational composition (Schelling, 1971; Bayer et al., 2007). A one-standard-deviation closer match between a mover's own occupational type and a destination's composition raised the probability of choosing that precinct by 30%, net of distance, amenities, and the built stock itself. Summed over thousands of such choices, the people refilling each precinct carried the composition of those they replaced.

The push is our first source of variation in residential choice and tests whether the matching just reflects who chooses to move. Tenant departures rose in the year a building

changed owners. In the years before the change, buildings that would later change hands and those that would not show similar turnover rates. Yet tenants pushed out when their building owner changed sorted to the occupational type as strongly as voluntary movers did. The matching is driven by where movers go, not by what made them leave.

Sorting alone can recreate persistence, but do newcomers bring their preferences with them, or do places convert whoever they receive? A worker moving into a working-class precinct might have been a Social Democrat who chose the precinct because of its neighbors, or might have become one by living there. If newcomers bring their preferences, persistence is consistent with selection and needs the sorting. But if a place's durable features convert newcomers, it is treatment and needs only the place. Conversion by the neighbors themselves needs the sorting too: neighbors can convert newcomers only where sorting has already assembled them. Our second source of variation is who fell under the rule that decided where households lived; the places themselves stay fixed. After the war, Munich's housing office (*Wohnungsamt*) assigned displaced households to the vacancies, so the households it placed could not act on the pull. By statute, the office could not weigh where anyone wished to live. Placed residents therefore met the same occupational composition, the same churches and schools, as those who chose to be there.

If a place's durable features converted whoever they received, how residents arrived should not matter for what a precinct became. Comparing precincts within the same district and holding pre-war amenities fixed, a precinct one standard deviation higher in 1910 composition sits 0.87 standard deviations higher in post-war composition where the office assigned few households, but only 0.64 where it assigned many. The same precinct votes 0.65 standard deviations less SPD where residents chose but only 0.37 where the office placed them. This pattern is robust to holding war destruction fixed or flexibly estimating the gradient within each quantile of destruction. Assignment also did not follow the pre-war map: within a district, neither a precinct's composition nor its vote before the war predicts how many households the office placed there. Newcomer sorting weakens with placement only after the

office began operating.

Munich's political map, then, persisted because each cohort of newcomers rebuilt it. Where placement removed their choice, much of the persistence went with it; exposure to the place alone does not sustain the pattern. What households select on is the design's limit: the match may reward similar neighbors, the amenities such neighborhoods share, housing a mover can afford, or proximity to work, and the levers cannot separate these. The argument does not need them separated, because on any of the four accounts newcomers arrive already sorted.

Sorting does not depend on anything specific to Munich, so it can sustain persistence elsewhere. As long as residents are free to sort onto a feature that also predicts the vote, newcomers can rebuild a city's political geography, in confessional cities or industrial ones. Even across cities, sorting may explain persistent political divisions such as East-West, Catholic-Protestant, or urban-rural splits in Germany or elsewhere.

This paper connects to three literatures. The first establishes that historical events generate persistent spatial differences in political and economic outcomes. Boundary designs at the mita border (Dell, 2010), medieval pogrom sites (Voigtländer and Voth, 2012), and colonial institutions (Acemoglu et al., 2001) show long-term persistence but cannot isolate the channel: at every boundary, institutions, infrastructure, and population composition differ simultaneously. Acharya et al. (2016) and Spenkuch and Tillmann (2018) extend this approach to political preferences and religious institutions. A related literature in spatial economics shows that places recover their economic size after temporary shocks (Davis and Weinstein, 2002; Bleakley and Lin, 2012), but the mechanism there is agglomeration, not political sorting. Heblich et al. (2021) show that historical pollution patterns drove persistent neighborhood sorting in English cities long after the pollution source disappeared, establishing that stable physical features can anchor residential sorting across generations. This paper extends their logic to political outcomes and adds a test that shuts down sorting directly, taking persistence as given and asking what sustains it.

The second literature exploits population shocks to narrow the set of explanations. Alesina and Fuchs-Schündeln (2007) document East-West German preference gaps surviving reunification, but displacement is partial, so transmission from the residents who stayed cannot be ruled out as the source (Bisin and Verdier, 2001). Newcomers can also rebuild the pattern themselves: Ochsner and Roesel (2020) trace seventy years of far-right voting in the destinations of migrating extremists to the local party branches those migrants founded. Bazzi et al. (2020) decompose persistence of frontier individualism into selective migration and causal exposure, finding both channels at work. This paper tests the sorting channel separately. When administrative placement removes residential choice, the composition–vote gradient falls by more than 40%, so exposure alone does not sustain the pattern.

Third, the neighborhood effects literature estimates place effects on individuals from movers (Chetty and Hendren, 2018) or quasi-random assignment (Damm and Dustmann, 2014; Chyn and Haggag, 2023). Cantoni and Pons (2022) apply an analogous decomposition directly to political outcomes, showing that place effects explain a substantial share of geographic variation in U.S. voting behavior. Their design identifies that place effects exist but cannot distinguish whether they operate through sorting, treatment, or institutions. Observing the sorting channel directly and shutting it down via placement tests sorting separately from place-based treatment, where partisan registration data alone can only bound their joint magnitude (Martin and Webster, 2020). Yet these mover designs answer a different question from this paper: whether places change the people who live in them. A place could have zero treatment effect on any individual and still sustain aggregate persistence through sorting.

The remainder of the paper proceeds as follows. Section 2 gives the historical background, Section 3 the data, and Section 4 the framework. Section 5 establishes that a precinct's prewar occupational composition predicts its vote through 1961. Section 6 turns to individual records and shows that turnover leaves no carriers for continuity or transmission, and movers sort to occupational type. Section 7 tests whether places convert whoever they receive: where

administrative placement removed the choice of destination, both composition and the vote persist less. Section 8 concludes.

2 Historical Background

Munich grew from 350,000 residents in 1890 to over 800,000 by 1933 as the royal capital industrialized. Because German law before 1951 did not recognize ownership of individual flats, nearly all of Munich's households rented. Working-class districts expanded east and south along rail corridors and factory sites, bourgeois quarters concentrated near the river Isar and the city center, and Catholic parishes clustered in the old town and the western neighborhoods. Each of these areas maintained its own dense institutional infrastructure of churches, schools, and associational life. By the turn of the century, Munich's neighborhoods were socially distinct, each with a characteristic mix of occupations, confessions, and institutions.

The SPD (Social Democrats) anchors the analysis because its voter base maps most directly onto the city's occupational geography. It drew its support from trade unions, workers' associations, and the industrial working class. As the alternative we study the Catholic bloc: Zentrum before 1918, the Bavarian People's Party (BVP) after 1918, and the CSU after 1945. Unlike the SPD, the bloc rested not on class but on confession, organized through the parish. Catholic associational networks (clubs, charities, newspapers) mobilized voters through institutional channels that operated independently of the labor market.[1]

Between 1893 and 1961, Munich voted under four political regimes. World War I ended the Empire and produced the Weimar Republic, which brought universal suffrage and a new party system. The social geography we document below survived the transition. The Nazi seizure of power in 1933 then suppressed free elections for sixteen years. The period between

[1] We set aside the right (DNVP, NSDAP) and the communists (KPD). Each was banned at some point and did not contest elections under all four regimes, so cross-regime persistence is undefined for them. The liberals contested throughout but their vote share collapsed below 2% by the early 1930s. Their spatial ordering persists nonetheless, nearly as strongly as the SPD's, which we show in Appendix A.

1933 and 1949 brought a major break in the social composition of Munich's neighborhoods. Allied bombing destroyed between 27 and 31% of Munich's housing stock (Haus der Bayerischen Geschichte, 2024), and the regime murdered or expelled virtually the entire Jewish population. Wartime evacuation and post-war refugee inflows reshuffled the city's residents. When federal elections resumed in 1949, the people voting in Munich's precincts were largely different from those who had voted in the same precincts before the war.

3 Data

Directories Munich's municipal address books (*Adressbücher*) list every household head by name, occupation, and street address, covering Imperial Germany (1893–1914), the First World War (1914–1917), the Weimar Republic (1919–1933), the Nazi period (1934, 1935 and 1938), and the early Federal Republic (1953 and 1961). The 1910 edition lists 175,947 household heads, the 1933 edition 288,958, and the 1953 edition 366,918.[2] Because each record pairs an occupation with a precise address, the same source shows who lives in each precinct and, by linking across editions, where residents move within the city. Each edition also lists the owner of every building, so the same source records who owns the housing stock and when a building changes hands.

Occupational Composition We classify each occupation using HISCAM, a social-stratification scale that measures occupational standing (not income) from 1 to 99, where a higher score marks a more prestigious occupation (Lambert et al., 2013).[3] Aggregated to the precinct, the

[2]The address books list household heads, overwhelmingly men in this period, and do not record women's occupations consistently. The occupational measure is therefore a precinct's occupational *type*, read from its male occupational structure, and not a statement about how women voted. Where the genders diverged at the ballot, the gap is known and bounded: Munich's 1924 returns, counted separately by sex, put women's vote for the confessional center about ten points above men's (Bühler and Sabet, 2026). That gap tracks religious identity rather than material circumstance, so it shifts the level of a precinct's vote and not the cross-precinct ordering on which the persistence result rests.

[3]The classification pipeline predicts standardized HISCO codes (Van Leeuwen et al., 2002) from raw German occupation strings using the OccCANINE machine-learning model (Dahl et al., 2024), then maps HISCO codes to HISCAM scores using the universal (u2) release of Lambert et al. (2013). HISCAM covers 93.4% of occupation strings by frequency.

mean HISCAM of its household heads gives the precinct's occupational composition, which anchors our analysis throughout. The pre-war anchor is the 1910 directory, the last edition Munich's address-level census can be matched to. We read it against the 1912 election, the last before the war. Figure G.1 shows the distribution of this measure. We standardize each precinct's composition by the 1910 mean (63.0) and standard deviation (6.15), so the regressions that follow measure effects per standard deviation of 1910 composition.

Voting Data The election data span 18 contests across the same four regimes: five Imperial Reichstag elections (1893–1912), eight Weimar Reichstag elections (1919–1932), the March 1933 election, and four Federal elections (1949–1961). We focus on the two blocs that contest all four regimes from a stable political position: the SPD and the Catholic bloc (ZENT/BVP/CSU). Vote shares equal party votes divided by valid votes cast and do not sum to one.[4]

Munich's precincts organized the count but elected no one, because seats were decided in constituencies and on party lists far larger than any precinct. No seat, candidate, or majority ever turned on where a precinct boundary ran, so boundary changes carried no political motive. As precincts are the finest geography at which the vote is reported, we use them to track the vote over time. Because precinct boundaries were redrawn for each election, precincts are administrative slices of the city's geography, not units into which residents could sort ex ante. All precinct-level data are harmonized to 1933 boundaries, which provide the finest consistent geography across the panel. The crosswalk maps each year's addresses to their 1933 precinct using the 1933 address book as the reference, yielding a panel of 380 precincts observed across 18 elections.[5] The 380 precincts nest within 30 older city districts (*Stadtbezirke*). A typical district in 1910 held about 500 buildings, 5,000 households, and 20,000 residents. A precinct held only about 36 buildings and 420 households.

[4]Vote distributions are shown in Figure G.1. The remaining vote went to parties we set aside for the reasons given in Section 2: the liberals, the right (DNVP, NSDAP), and the communists (KPD).

[5]The persistence and composition regressions require a precinct's 1910 occupational anchor (at least ten listed occupations) together with a post-war vote, a condition the full panel does not meet everywhere. The working sample is about 335 precincts, falling to 297 once the full pre-war amenity set is required.

Amenities We construct 39 precinct-level pre-WWI amenity variables from six sources: address books (establishment densities, housing quality), the 1910 municipal census (population, demographics, religion), parish records (church locations), school records, a 1912 factory directory (211 firms, 86% geocoded), and a 1900 tramline shapefile. The set combines stable physical features of a precinct (establishment densities, distances, housing form) with population measures from the 1910 census (total population and shares by religion, age, sex, and place of birth). Distance variables measure the mean Euclidean distance from all addresses in a precinct to the nearest instance of each amenity (churches, schools, pubs, factories, tram stops, post offices, and others), with summary statistics in Appendix G. Section 5 uses the full set of 39 to characterize composition exploratorily. From Section 6 onward, wherever amenities serve as controls or as LASSO candidates, we restrict them to the 32 that are physical features of place (densities, distances, housing form, population total), excluding the religion, age, sex, and birthplace shares, since those measures are themselves compositional and could absorb the sorting we aim to isolate.

Destruction We measure wartime destruction as the percentage change in listed households between the 1933 address book and the first post-war directory of 1953. We take precincts that lost a larger share of their listed households as having suffered more physical damage, a reading that captures destruction at precinct resolution without relying on retrospective surveys or aggregate bomb-damage assessments. The measure conflates bombing damage with deportation, evacuation, voluntary wartime migration, and the post-war inflows that preceded the 1953 directory. It is therefore a composite shock to the local population rather than a measure of physical destruction alone.[6]

[6]The mean change is -17.4% (sd = 29.5%), with 65.5% of precincts showing net population loss. The hardest-hit precincts lost all listed households; some peripheral precincts gained population through refugee settlement.

4 Framework

We read persistence off votes, but votes are counted only in aggregate. Lacking finer data, most work uses municipalities, counties, or constituencies (Cantoni et al., 2020; Ochsner and Roesel, 2020). In contrast, our panel resolves the vote into 380 precincts within a single city. A Protestant rather than a Catholic majority or a different institutional inheritance (Spenkuch and Tillmann, 2018) separates regions but cannot distinguish precincts that all sit in Catholic Munich under the same franchise and party system. Yet even a precinct containing only a few hundred voters reports a single tally to protect the secrecy of the ballot.[7] As a consequence, aggregate results can only show that a place's composition and its vote move together. Separating places treating people from people choosing places requires tracking how a place's composition regenerates as residents sort in and out.

We argue that residential sorting builds a precinct's composition, and composition in turn predicts the vote (Figure 2). We provide evidence for this chain at both the aggregate and the individual level. We first ask at the aggregate level whether composition predicts the vote beyond stable precinct characteristics and past voting results.[8] Then, at the individual level, we ask whether newcomers rebuild that composition once a precinct's prior residents are gone, by sorting into precincts that match their own occupational status.

Conceptually, however, observing a move does not reveal why a household made it. A push forces a household out of where it lived, independent of where it lands, while a pull draws it toward a matching destination, independent of where it started. Because we see the move but not its cause, we exploit an exogenous lever on each margin in turn. The owner-change design asks whether sorting survives when the decision to leave comes from outside the household: a change in a building's owner pushed sitting tenants out for reasons unrelated to where they would go. The placement design asks whether persistence weakens

[7] Some studies observe whether a given person turned out to vote (Cantoni and Pons, 2022), and others track the registered party affiliation of 143 million individual voters over a decade (Brown et al., 2025), but no source records how anyone voted, because the ballot is secret.

[8] Composition enters at its predetermined 1910 value throughout, so no later vote or move can feed back into the measure. That feedback is the simultaneity Manski (1993) calls the reflection problem.

when the choice of a destination is removed: the post-war housing office assigned newcomers to addresses instead of letting them choose.[9] Together, the two levers pose the same question from the push and the pull side: does choosing where to live sustain the geography, or does the place convert whoever arrives?

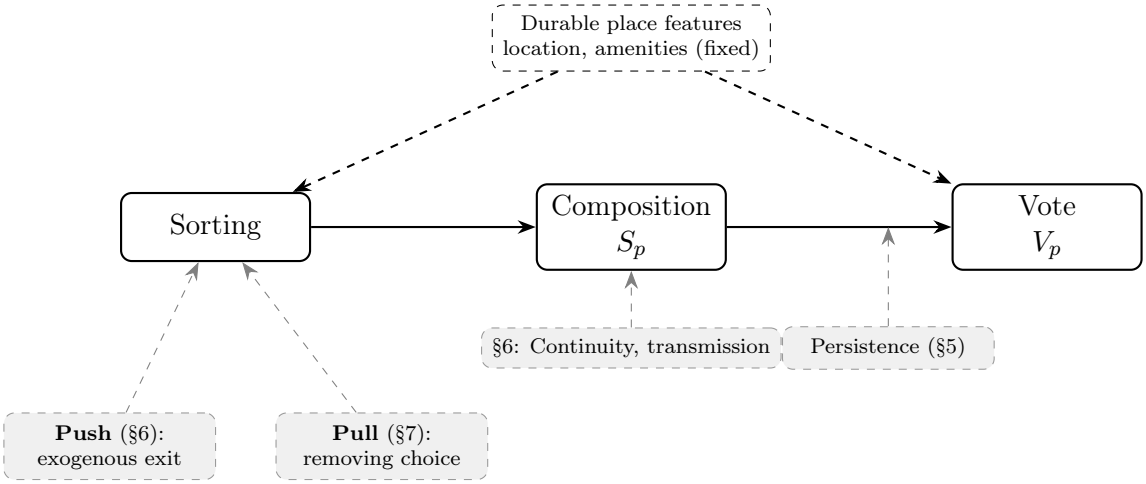

Figure 2: The framework

Notes: Residential sorting builds a precinct's occupational composition S_p, which in turn predicts its vote V_p. Durable place features are predetermined, and dashed arrows mark their possible correlations with sorting and the vote. Section 5 estimates the composition–vote link at the precinct level, controlling for durable place features. Using individual records, Section 6 asks which process regenerates S_p once a precinct's residents are replaced: continuity (6.1), transmission (6.2), or sorting (6.3). Two exogenous levers act on the choice margin. Owner change exogenizes whether a household leaves (the push), holding fixed where it lands. Administrative placement removes the choice of where (the pull), and Section 7 uses it to turn sorting off: where the post-war housing office placed newcomers instead of letting them choose, the composition–vote gradient of Section 5 flattens.

The two levers test whether selection, rather than the place itself, sustains the composition–vote gradient, but cannot reveal what households select on. A precinct's durable features, the residents they gather, the rents they command, and the workplaces within reach move together, leaving four accounts observationally equivalent: households drawn to similar households, households drawn to the amenities, households sorting on what they can afford, and households settling near where they work. The fourth account needs no homophily at all: households that settle near their work cluster by occupation without any pull toward similar

[9]The housing office allocated dwellings address by address. The 1933 precinct boundaries on which we measure composition are an administrative construct that played no part in the assignment and had no post-war function.

neighbors. Including amenities in the choice model does not separate these accounts, because the amenities we observe are a subset of the durable bundle and are collinear with composition. The design can show that choice sustains the geography, but cannot distinguish which feature of a place that choice responds to.

5 Persistence and Composition

5.1 The Persistence Fact

Munich kept the same political geography from the late German Empire to the Federal Republic. Figure 3 ranks each precinct by its vote share within an election year and plots the 1912 ranking against the ranking in two elections from each later regime: 1919 and 1932 under Weimar, 1949 and 1961 under the Federal Republic. For the SPD, precincts cluster around the 45-degree line. A precinct's 1912 SPD rank still predicts its 1961 rank with a slope of 0.83, half a century and three regime changes later.[10] The Catholic bloc looks less stable on the same 1912 baseline, its slope falling to about 0.32 by 1961. That baseline predates the reconstitution of the Zentrum as the BVP in 1919, so comparisons from 1912 cross a break in the party system and understate the bloc's persistence.[11]

What recurs is the political geography, which precincts lean Social Democratic and which lean Catholic. By the time elections resumed in 1949, the war and its aftermath had reshuffled who lived where. The puzzle is why the vote map held its shape while most of the electorate that produced it was replaced.

[10] The ordering persists further still: the 1912 SPD map continues to predict the left bloc's vote in the 2023 Bavarian state election, 111 years and five regime changes later (Appendix B).

[11] The Catholic slope roughly doubles when the baseline moves to 1919, the first election the BVP contested, rising from about 0.32–0.36 to between 0.59 and 0.73. The liberal vote persists just as strongly. Its rank ordering holds with a slope between 0.70 and 0.82 at the 1919 anchor, even as the liberal share fell below 2% by the early 1930s (Appendix A).

5.2 Composition

For the vote map to persist through a turnover of voters, it must track a feature of the precinct itself. Our conceptual framework develops a link from sorting to composition to the vote, with durable place features as a potential confounder. We define a precinct's composition as the mean HISCAM of its household heads, $S_p = N_p^{-1} \sum_{i \in p} \text{HISCAM}_i$, standardized across precincts. Holding S_p at its predetermined 1910 value, we ask whether pre-war composition predicts how a precinct later voted,[12]

$$\bar{V}_{p,t} = \alpha + \beta\, S_p^{1910} + X_p'\delta + \mu_{d(p)} + \varepsilon_{p,t}, \tag{1}$$

where $\bar{V}_{p,t}$ is precinct p's average vote share in period t (Weimar or post-war), X_p the pre-war amenities, and $\mu_{d(p)}$ a district fixed effect. Standardizing S_p makes β the change in vote share per standard deviation of 1910 composition (Table 1). The Weimar panel tests the gradient close to when composition is measured. The post-war panel tests it on votes cast once the population had turned over.

The SPD's working-class base predicts a negative gradient, with higher-status precincts voting less SPD, and the status gradient in confessional voting predicts a positive one for the Catholic bloc. With district fixed effects alone (columns 1 and 4), a one-standard-deviation increase in 1910 composition lowers the SPD share by 5.7 percentage points in the Weimar elections and 4.1 after the war. Conversely, it raises the BVP/CSU share by 3.8 and 3.1 points, both significant at the one-percent level.[13] The gradient remains robust when including the full pre-war amenity set (columns 2 and 5) and the precinct's own 1912 vote (columns 3 and 6). If composition merely stood in for a precinct's prior partisanship,

[12] We anchor the voting correlations at 1912 but read composition and amenities from the 1910 directory, the last address-level Munich source with census demographics. The one-edition gap predates every move we study and does not affect the persistence estimates.

[13] Relative to period mean vote shares in the estimation sample, the gradient is about twice as steep in Weimar as after the war. A one-standard-deviation increase in 1910 composition moves the SPD vote by 22% of its Weimar mean (0.26) against 12% of its post-war mean (0.35), and the Catholic vote by 15% of its Weimar mean (0.26) against 7% of its post-war mean (0.45). The link survives the war but weakens.

holding the 1912 vote fixed would absorb it, yet composition still moves the post-war SPD and BVP/CSU shares by 1.9 and 2.3 points per standard deviation, both at the 1%-level.[14]

These regressions show that a precinct's pre-war composition predicts its vote a generation and a war later, but they cannot resolve the 'places converting-versus-choosing' ambiguity. The converting side, the one the literature has emphasized, runs through institutions that outlived the war (Acemoglu et al., 2001; Dell, 2010) or transmission by the residents who stayed (Bisin and Verdier, 2001); the choosing side is residential sorting (Tiebout, 1956; Bayer et al., 2007). Telling them apart means tracing three processes that could rebuild a precinct's occupational composition once its residents are replaced: continuity, transmission, and sorting.

6 The Channels

We consider three main channels through which a precinct's occupational composition can survive several electoral cycles, political regimes, and two world wars. [15] First, incumbent residents may simply remain in their precincts and continue to vote as before. Second, those residents who stay within a precinct may transmit political norms to newcomers who eventually replace them. Third, movers may sort into precincts that match their own type, reconstituting the local occupational composition. Parties and civic organizations could also sustain partisan identities across generations, independent of who lives where. Institutions are a natural fourth candidate, but the Nazi seizure of power dissolved the parties and their associational networks in 1933, sixteen years before free elections returned, making durable institutional transmission implausible.[16]

[14] A precinct's 1910 composition also predicts its Weimar and post-war composition within district, and net of amenities (Appendix Table G.1), so fixing S_p at 1910 tracks a feature that survives the population turnover rather than imposing one.

[15] All following tests below draw on the annual run of Weimar editions (1919–1933), with one pre-war-to-post-war pair added in Section 6.1 to bound retention across the war itself.

[16] The exception is the Catholic Church, whose parish organization survived both the regime and the war. A parish, however, is anchored to its place: whatever influence it exerts on newcomers is a feature of the precinct, not a carrier that leaves with its residents. Churches accordingly enter the analysis as part of the

The first two channels operate through the existing stock of residents and capture transmission, whereas the third operates through the flow of entrants capturing sorting. We bound the first two and find both too weak to carry persistence. Retention of a precinct's residents from one year to the next is low, and tenancies are too short to transmit local norms to newcomers. In contrast, a mover's occupation type significantly predicts which precinct they enter, even controlling for distance and amenities. We find that this sorting is strong enough to rebuild the compositional map across several decades.

6.1 Continuity

The simplest carrier of persistence is residents themselves. If the same people stay, composition almost mechanically continues with them. We measure this by precinct retention (θ_p), the share of residents at $t-1$ still present at t, constructed by tracking each household head across consecutive address-book editions. A resident is retained if they reappear anywhere inside the precinct's 1933 boundaries in the following edition. It is the ceiling on the continuity channel since a precinct can reproduce its politics through the same people only as far as those same people remain. To bound name-matching uncertainty from both sides, we report a conservative lower bound and a loose upper bound, bracketing the tradeoff between false positives and false negatives inherent in automated linkage (Bailey et al., 2020).[17]

Our conservative lower bound requires an exact full name and the same leading HISCO digit. The HISCO leading digit identifies the major occupational group, distinguishing, for example, professionals from service and production workers. This bound is strict by design as it treats spelling variations as non-matches and hence understates retention. The loose upper bound allows a full-name Jaro-Winkler distance of at most 0.10 and the leading HISCO digit to differ by at most two. This accommodates plausible occupational mobility across adjacent groups as well as minor spelling variations.

durable amenity bundle, and Section 7 tests whether such place features convert whoever arrives (Table D.5 reports the parish case directly).

[17] Matching statistics are in Appendix G.

During the Imperial and Weimar decades, retention runs between 0.43 and 0.62 under the conservative bound and up to about 0.87 under the loose one (Panel a in Figure 4). Taking the conservative numbers, a precinct sheds between a third and half of its residents from one year to the next.[18] However, election cycles span several years, so the share of a precinct's cohort that remains across a full electoral period is substantially lower. The one major drop comes across the Second World War, when retention falls to about 0.07 between 1935 and 1953. Only one in fourteen of a precinct's pre-war residents is still in the same precinct 18 years later.[19]

Panel (b) of Figure 4 breaks down the precinct retention into stayers who remain at the same address and movers who change address within the precinct. The retention rate is driven almost entirely by stayers, so when a precinct's composition changes it is because newcomers arrive from outside, not because sitting residents reshuffle within it. One might worry that residents optimized around precinct lines, but those lines are administrative boundaries drawn in 1933, after the moves we observe, so no earlier mover could have selected around a boundary that did not yet exist.

Continuity of residents is therefore real but thin: those who stay are rooted in place, but few remain for long, and the stock of residents collapses across the war.

6.2 Transmission

Even a year-to-year retention rate of about 50% says little about whether a few long-tenured incumbents exist to transmit local norms. The relevant question may then be not how many residents survive from one edition to the next, but instead how long *some* residents stay in

[18]This turnover is large compared to the U.S. context because until 1951 German law essentially did not permit individual ownership of a single flat within a multi-unit building. The Wohnungseigentumsgesetz (WEG, Condominium Act) was enacted only in 1951. Before it, under §§93–94 BGB whoever owned the land owned everything firmly attached to it, "in particular buildings," so a building was a legally inseparable component of its parcel and one could not carve out a single flat as a standalone object of ownership. Urban residents were therefore overwhelmingly renters, and renters move far more than owner-occupiers.

[19]The other two drops, from 1917 to 1919 and in 1933, can be explained by the directories themselves. No directory was published in 1918, so the pair spans two years, which lowers matching rates. The 1933 edition is the first to list residents in six columns rather than four, and it abbreviates first names to initials at higher rates, making matching on full names more difficult.

a precinct. Within the Weimar Republic, the stock of long-tenured incumbents turns out to be thin.[20]

Figure 5 asks how long the renters recorded in 1933 had already been living in their precinct. For each previous year, it shows the share of those 1933 renters who were already present in the same precinct at that earlier date.[21] The share falls the further back we look, reflecting the high turnover in the rental market. Even a few years before 1933, most of the eventual renters had not yet arrived. In 1930, only 30% (51% upper bound) of the 1933 renters were already in their precinct, and by 1928 fewer than 23% (40% upper bound) were. In 1912, the anchor year of our voting patterns, only 3% (8% upper bound) of the 1933 renters had been in the same precinct. Transmission requires incumbents present long enough to pass on local norms, and the pattern it must explain holds across the whole city and for both blocs (Figure 3). A pattern that recurs everywhere is unlikely to depend on a carrier this scarce, so transmission is unlikely to carry the map either.

6.3 Sorting

Given that continuity and transmission are unlikely to carry persistence, the remaining candidate is sorting: where moving households choose to go. We define a mover as a household head who reappears in the next edition at a new address outside their former precinct, linked under the conservative matching rule of Section 6.1. We model a mover's destination choice with a conditional logit, the canonical discrete-choice estimator for residential sorting (McFadden, 1974; Bayer et al., 2007). For each mover i in year t, the choice set is all precincts with observed composition in t, including the mover's own origin precinct, and i selects one. Identification comes entirely from variation across the precincts in a mover's choice set: any attribute constant across that set drops out, so only the alternative-specific type-match and

[20]The war destroys the incumbent stock almost entirely as fewer than one in fourteen pre-war residents remain afterward.

[21]Each earlier year is matched directly against the 1933 cohort, so a resident who appears in 1928 and 1933 but not in 1930 still counts as present in 1928. This avoids deflating the estimates because of gaps in intermediate editions.

distance shift the choice.[22]

A mover i in year t draws indirect utility from each precinct j in the choice set $\mathcal{J}_t$,

$$u_{ijt} = \gamma\, m_{ij} + \rho\, d_{ij} + X_j'\beta + \mu_{d(j)} + \epsilon_{ijt}, \qquad \text{with} m_{ij} = -\,|\tau_i - S_j|, \tag{2}$$

where the type-match m_{ij} is the negative distance between mover i's own occupational status τ_i and destination precinct j's mean composition S_j, d_{ij} is distance from the origin, X_j are pre-war amenities, $\mu_{d(j)}$ a destination-district fixed effect. The mover chooses the precinct of highest utility, which gives the standard conditional-logit choice probabilities over the choice set $\mathcal{J}_t$.[23] A positive γ means movers favor destinations close to their own type, i.e., homophilous sorting in the sense of Schelling (1971).

The parameter γ is not a destination treatment effect but reflects the destination composition, which itself is an equilibrium outcome of who already lives there. Without further identifying variation, a positive γ could reflect a durable precinct quality correlated with composition rather than a genuine pull toward own-type neighbors (Bayer et al., 2007).

The first threat is that a mover responds to a composition their own arrival helps determine (Manski, 1993). We fix composition at its predetermined 1910 value, which predates every move we study, so this channel is closed by construction. The second threat is a time-invariant amenity correlated with 1910 composition.[24] The logit absorbs this by adding LASSO-selected physical amenities and destination-district fixed effects, which absorb every district-constant unobservable and identify γ from within-district variation.[25]

[22] We estimate the model as a binary fixed-effects logit over the expanded choice set, which coincides with the conditional-logit estimator given one chosen alternative per mover and a large alternative set.

[23] $\Pr(j \in \mathcal{J}_t) = \dfrac{\exp\big(\gamma\, m_{ij} + \rho\, d_{ij} + X_j'\beta + \mu_{d(j)}\big)}{\sum_{k\in\mathcal{J}_t} \exp\big(\gamma\, m_{ik} + \rho\, d_{ik} + X_k'\beta + \mu_{d(k)}\big)}$.

[24] The field's formal fix builds an internal instrument from the equilibrium structure of the sorting model (Bayer and Timmins, 2007), requiring that structure to be imposed. We substitute predetermination, fixed effects, and the two design-based tests below; Kuminoff et al. (2013) surveys the problem and its solutions.

[25] The conditional logit imposes independence of irrelevant alternatives (IIA): the relative odds of choosing between any two precincts do not depend on the rest of the choice set. IIA would be violated if precincts within the same district were closer substitutes for one another than precincts in different districts. We test this below with a district/precinct nested decomposition (Table 3).

6.3.1 Movers Sort to Type

Table 2 reports conditional logit estimates of destination choice among Weimar movers (1919–1933).[26] With mover fixed effects alone, a mover is significantly more likely to choose a precinct whose occupational composition resembles their own type, net of a large and negative distance gradient. That pull could still run through an amenity correlated with composition rather than through type itself, so we add the fifteen pre-war physical amenities that best predict the 1910 composition.[27] Controlling for them leaves the type-match coefficient intact, though unobserved amenities, rents, and proximity to work remain in the bundle (Section 4). Because a mover may also respond to a composition their own arrival helps determine, Panel B fixes it at its predetermined 1910 value, and the pull again holds. A mover is about 27% more likely to choose a precinct one standard deviation better matched to their type, a premium worth 440–580 meters of extra travel.[28]

One concern is that sorting operates at the district rather than the precinct level, since districts were the units residents knew by name while precinct lines existed only to organize the count. Movers would then choose a district for its type and spread randomly across its precincts. We test this directly by decomposing the destination choice into two sequential stages, first among Munich's 30 districts and then among the precincts of the chosen district (Table 3). Both the between-district estimate and the within-district one are positive and highly significant. With 62% of the precinct-level type-variance lying within districts, movers do not merely select a district but pick the type-matching precinct *within* it. A separate

[26] Weimar is the estimation window because it offers annual editions and the owner-change instrument. Sorting is not confined to it: the same logit estimated on individuals linked across 1910–1919, 1935–1953, and 1953–1961 leaves the type-match coefficient positive and significant in every period, at a level roughly a third lower after the war than before (Appendix F).

[27] The selected amenities are the distance to the nearest café, household density, pub density, the mean maximum floors, the distance to the nearest association, the distance to the nearest pub, address fragmentation, the number of schools, the rear-building share, the number of churches, total population, the number of post offices, bakery density, the distance to the nearest school, and the distance to the nearest bath house; these are the fifteen amenities of Table E.4.

[28] The marginal rate of substitution between type-match and distance is $\hat{\gamma}\,\frac{\mathrm{SD(match)}}{|\hat{\rho}_{\mathrm{dist}}|} = 0.27 \times \frac{0.88}{0.94} \approx 0.26$ log-meters; at the mean (median) move of about 1,989 (1,506) m, $\exp(0.26) - 1 \approx 0.29$ implies roughly 580 (440) m.

concern is that $\hat{\gamma}$ reflects a mover's homesickness for a place resembling the one they left rather than a pull toward their own type: the estimate is unchanged when the origin precinct is excluded from the choice set, and movers who change occupation between editions move toward precincts matching their new type, not their old one (Appendix E.1).

Across precincts, the mean type of a precinct's in-movers correlates with its 1910 composition at about $\rho = 0.8$ (Figure 7). Newcomers who never lived in a precinct reproduce its pre-war type-mix, rebuilding the composition that Section 5 showed to predict the vote.

6.3.2 The Push: Owner Change

The conditional logit is estimated only on movers who chose to leave their precinct. If residents move only when they find a precinct that matches their type, the estimated match effect overstates the homophily a representative mover would exhibit. We isolate that possibility by comparing sorting among residents pushed out by a change in their building's owner with sorting among those who left voluntarily. An owner change realigns the tenure relationship for sitting tenants for legal and financial reasons orthogonal to where any given tenant would prefer to live.

The owner change is a strong instrument for displacement. Table E.1 shows that an owner change raises the probability of household turnover by 2.3–3.1 percentage points on a baseline turnover rate of 19.4% during the Weimar period. The estimate survives precinct×year, street, address, and censored first-treatment-period-only specifications. To rule out differential pre-trends, we exploit the staggered timing of ownership changes across buildings in a difference-in-differences design, comparing turnover at treated buildings against the same buildings before the change and against buildings that had not yet experienced one. Using the Callaway and Sant'Anna (2021) estimator, the event study (Figure 6) shows no significant differences between buildings with and without an owner change across five pre-periods. At impact, tenant retention drops by approximately 2.2 percentage points. Turnover reverts to near zero two years after the owner change.

Table 4 runs the same homophily logit separately on movers from owner-change and control buildings across annual Weimar pairs (50,642 control-building vs. 3,875 owner-change-building movers), and finds the pushed movers sorting more strongly in every specification. But most movers from an owner-change building would have moved regardless, so this comparison mixes the movers the owner change induced with a majority of ordinary ones. Scaling the difference in destination match between the two building types by the difference in their moving rates isolates the induced movers alone.[29] Their destinations match their own type at -0.88 standard deviations of type-distance, close to the -0.98 among movers from control buildings and well above the -1.20 those same induced movers would reach absent any sorting. Even tenants who moved only because their building changed hands sorted onto the pre-war type-map.

Taken together, this section demonstrates that rental turnover is high, so people themselves do not carry a precinct's political preferences forward. Transmission of local political norms is unlikely given the short duration people stay within a precinct. Sorting is the most likely remaining explanation, as newcomers move toward precincts whose occupational composition matches their own. This result holds at both the district and precinct level and survives even when isolating tenants an owner change forced to move. That movers sort to type may imply that places also convert newcomers. If they do, the voting map would persist with or without the sorting. The following section removes the choice of moving to a preferred destination and thereby tests that reading.

[29]The owner change raises the probability of an observed cross-precinct move from 3.9 to 4.7%, so about one in five of the observed movers from these buildings moved because of it. That scaling is the Wald estimand: with Z the owner change, D an observed cross-precinct move, and m the destination type-match, $\frac{E[mD|Z=1]-E[mD|Z=0]}{E[D|Z=1]-E[D|Z=0]}$ recovers the induced movers' mean match, provided an owner change displaces rather than retains tenants and shifts whether a tenant leaves, not where voluntary leavers go (Angrist and Pischke, 2009). The sample conditions on re-finding the household in the following edition, so tenants pushed out of Munich, or out of the linked sample, lie outside the estimand. The benchmark below is each induced mover's expected match over the destinations other same-year movers chose. The induced movers' match of -0.88 carries a 95% confidence interval of $[-1.03, -0.73]$. The 3.9 and 4.7% base rates are the share of linked movers whose destination has measurable 1912 composition, a narrower margin than the raw turnover rate in Table E.1 and Figure 6; the two are not competing estimates of the same effect.

7 Removing Choice

Sections 5 and 6 establish that sorting, not continuity or transmission, regenerates a precinct's composition, and that composition predicts its vote. Sorting evidence alone cannot separate two accounts of why the vote tracks composition. Residents may carry their political preferences with them and choose precincts to match, or places may convert whoever moves in. Administrative placement lets us tell the two apart. By assigning residents to addresses without letting them choose, it holds place fixed while removing the choice margin. If the composition–vote gradient survives placement undiminished, places are doing the work; if it flattens, the people are.

7.1 Administrative Placement: The Post-War Wohnungsamt

The post-war housing shortage placed housing allocation under the municipal housing office, the Wohnungsamt. Under Kontrollratsgesetz Nr. 18, the office ranked housing seekers on a need-based urgency list, the *Dringlichkeitsliste*, and assigned them to specific addresses through billeting orders (*Einweisungsverfügungen*) (Alliierter Kontrollrat, 1946; Landeshauptstadt München, Sozialreferat, 2011).[30] By 1948 over 93,000 applications stood on file against roughly 12,000 assignments a year (Landeshauptstadt München, Sozialreferat, 2011). The statute barred preferential treatment by social or financial standing, so a neighborhood preference was not a criterion the office could act on. The allocable stock was assembled over the post-war years from requisitioned, purchased, repaired, and newly built housing, none of it chosen by the residents placed in it (Egner, 2014). The placed households themselves raise a first concern: forced migration shifts preferences durably (Becker et al., 2020), so a gradient that flattens under placement could reflect who arrived rather than the removal of choice. The households the office ranked, however, were largely the city's own. A *Zuzugssperre* (ban on in-migration) closed Munich to new arrivals from August 1945 through

[30] Appendix C documents the legal framework, the categories of administratively allocated housing, and the threats to the identifying assumption.

the early 1950s. The refugees and expellees who did enter, 10.3% of the population at the 1950 census, were managed by a separate refugee administration, the *Flüchtlingsverwaltung*, and housed primarily in peripheral settlements on repurposed military infrastructure (Appendix C). The remaining concern is the stock's location, since nothing in the institution guarantees that requisitioned, repaired, and newly built housing avoided precincts of a particular pre-war character.

We measure placement by the administrative share, i.e., the share of a precinct's 1953 residents living in administratively allocated housing (government, nonprofit, or cooperative), read from owner identity in the address book.[31] The office's application and placement records were largely destroyed in the 1960s, so we observe each precinct's exposure to placement rather than the individual households the office assigned. The comparisons that follow are therefore between precincts with more and less administrative housing, not between assigned and self-selected households.

Munich-wide, the administrative share averages 9.8% across precincts, with district-level shares ranging from 2% to more than a third of all housing, some of it predating the Wohnungsamt.[32] The newly-built periphery tends to have more administrative housing than the center, a variation our district fixed effects absorb. Yet even where the stock predates the office, its function changed with the war: before the war people applied to live there, and after the war they were assigned there. The office could also requisition space in privately owned buildings, which the address book still records under a private owner, so the administrative share understates placement and the flattening we estimate is a lower bound.

We verify that the administrative share does not correlate with pre-war characteristics that drive persistence by regressing it on the 32 pre-war amenities with district fixed effects

[31]The administrative share is a lower bound on administrative placement. A household billeted into a privately owned building under the housing-management law's forced apartment-sharing provisions is recorded under the private owner, so we observe placement only in administratively owned housing.

[32]Precincts count ballots and are not units of housing allocation, so the district is the most likely level at which allocation operated. Of 336 precincts, 48 hold none and 288 a positive share. We winsorize the administrative share at the 90th percentile (0.33) so that a few high-placement precincts do not dominate the estimates. The results hold when they are left in.

and then adding 1910 composition and the 1912 SPD vote (Table 5). Neither composition nor the vote enters significantly, and the within-district R^2 holds at 0.183 with the amenities alone and 0.184 once both are added. A within-district LASSO over the 32 amenities, 1910 composition, and the 1912 vote keeps only three amenities and drops both persistence variables, at a cross-validated R^2 between 0.03 and 0.06.

Still, the administrative share is measured in 1953 and was not assigned by lottery. To mitigate remaining concerns, we use the change in the administrative share from 1910 to 1953, ΔAdmin, as a continuous moderator of the 1910 composition–vote gradient.

7.2 The Pull Factor

A place that converts whoever arrives reproduces the 1910 type-map however its residents came, so the gradient should not vary with placement. If instead residents sorting to type rebuild the map, the gradient should flatten where placement replaced that choice. Figure 8 plots the standardized gradient from 1910 composition to three post-war outcomes, estimated separately in precincts where administrative ownership did not grow over 1910–1953 and in those where it did.[33] In all three outcomes the gradient attenuates toward zero. The standardized gradient linking the 1910 composition to its post-war counterpart falls from 0.87 to 0.64 of a standard deviation, from no-growth precincts to those where administrative ownership grew, as does its link with SPD and CSU voting patterns. The two subgroup gradients differ significantly for composition, the SPD vote, and the CSU vote alike (Appendix Tables D.1 and D.2).

Reproduction of the type-map thus weakens where choice was restricted, first in composition and then in the vote that composition predicts, the pattern sorting implies and place-based conversion does not. The administrative share is an ownership object that no

[33]The figure plots jointly estimated subgroup slopes. The continuous interaction model that formally tests whether the two gradients differ ($\text{Status}_{1910} \times \Delta\text{Admin}$) is in Appendix D, Tables D.1 and D.2. The interaction is significant in every outcome, and the vote interactions stay significant once a precinct's own 1912 vote is held fixed. Appendix Figure D.1 shows the same subgroup gradients hold under district-clustered, wild-cluster-bootstrap, and spatial-HAC standard errors, so the result is not an artifact of the clustering or inference method.

lottery assigned, so we read the contrast as descriptive moderation rather than a second identification.

Part of the composition flattening may be mechanical: if the urgency ranking tracked occupational standing, the placed households would lower the mean status of any precinct they entered. The urgency ranking did not track occupational standing, and a bombed-out doctor ranked beside a bombed-out worker. With the office's records destroyed, how far the placed households moved a precinct's mean status stays unmeasured. A precinct's vote share, however, is not an average of its residents' status scores. The vote gradient flattens with placement as well, and no averaging of the placed households can produce that.

New Arrivals That attenuation is also directly visible in who newly arrives. If placement overrides sorting, the newcomers entering a precinct should track its prior occupational type-map less closely where the housing office assigned more units. Table 6 regresses a precinct's newcomer mean HISCAM on its prior composition, its 1953 administrative share, and the interaction of the two. With district fixed effects and the amenity controls, the interaction is negative in the post-war window: a higher administrative share weakens the tie between a precinct's prior composition and the type of those who move in (column 3, -0.382, significant at the 5% level). The same specification on a pre-Wohnungsamt placebo, regressing 1933 newcomers against the 1910 map before the housing office existed, returns essentially zero (column 6, -0.046). That placebo window is not free of institutional housing, since cooperative and non-profit landlords already owned buildings in Munich before 1933. What separates the two windows is how that housing was allocated: a household joined a building cooperative by choice, while the post-war housing office assigned dwellings to households it was legally obliged to house (Appendix C).

The Vote Moves with the Map If newcomers bring their preferences with them, the vote should follow the composition wherever it goes; if places convert their residents, the vote should stay tied to the place's durable features even where the composition moved. We

therefore regress the change in a precinct's vote share from 1912 to the post-war average on the change in its standardized composition from 1910 to 1953 (Appendix Table D.3). Long-differencing removes every time-invariant characteristic of the precinct, durable amenities and fixed political culture alike, which is precisely the bundle the conversion account relies on; what remains to correlate the two changes is who moved in and out. A precinct that moved one standard deviation up the composition ranking moved 3.8 to 4.4 percentage points away from the SPD ($p < 0.001$), matching the level gradient of Section 5. The CSU change gradient is positive and significant at half its level counterpart, consistent with a confessional vote that follows class composition less tightly than the SPD vote does. Interacting the composition change with administrative growth leaves the link unchanged (SPD 0.015, $p = 0.85$; CSU -0.026, $p = 0.56$). Whatever population arrived, the vote follows the composition that arrived.[34]

War Destruction Heavily bombed precincts received more administratively allocated housing, so destruction rather than placement could drive the attenuation. The interaction survives controlling for war destruction as a level (Appendix D). It also survives within bins of destruction intensity, where the status gradient is free to differ across bins. Our destruction measure counts lost households, however, and bombing may have reshaped neighborhoods in ways that household loss does not record. That channel stays uncontrolled, so the destruction tests support the pull without excluding every alternative.

Surviving Institutions The parish network survived the regime and the war, so churches could sustain the gradient as a treatment rather than as a feature movers sort on. If parishes convert whoever arrives, the gradient that survives placement should be preserved near churches. The prediction is strongest for the CSU, whose milieu survived the regime while

[34]The sharpest case is the expellees: their own party, the BHE (*Block der Heimatvertriebenen und Entrechteten*), ran in the 1950s elections on an expellee platform outside the class blocs. Munich's expellee share stayed well below the national average because destroyed cities took in few (Historisches Lexikon Bayerns, 2024). The BHE's precinct vote share is correspondingly small, which bounds any mechanical dilution of the SPD and CSU shares, and it is uncorrelated with the growth in administrative ownership.

the SPD's was destroyed. Neither pattern appears (Table D.5). Proximity to a 1910 church leaves the gradient unchanged where residents chose their precincts. Where administrative ownership grew, the flattening is, if anything, stronger near churches, for both parties. Had more than a fifth of the total gradient been preserved near churches, the design would have detected it (Appendix D).

8 Conclusion

Munich's political geography outlived the people who made it. Wartime destruction and turnover left fewer than 1% of the 1912 residents at the same address by 1953, yet precinct rankings from before the First World War still predict their post-war counterparts. We show that residential sorting is sufficient to sustain this persistence, a mechanism that requires only a durable feature of place against which each new cohort can sort. In Munich that feature is the occupational composition itself: newcomers chose precincts matching their own type and so rebuilt, year after year, the composition that predicts each precinct's vote. Part of what the literature credits to surviving institutions or transmitted culture is therefore selection.

Our setting differs from the existing literature in two ways. First, we can unpack observed aggregate persistence into individual residential choices, following movers across the city and showing how their matching with existing residents recreates the observed persistence. Second, we can show that the persistence tracks residential choice: it weakens wherever the housing office, not the household, chose the destination. Yet our design observes only realized moves, so it cannot identify what drew a mover to the destination: similar neighbors, shared amenities, affordable rent, or proximity to work.

The implications for future research on persistence are straightforward. Because sorting alone can sustain persistence, ruling out sorting cleanly requires a setting where it is shut off completely, so that whatever persists must come from culture or institutions.

References

Acemoglu, Daron, Simon Johnson, and James A. Robinson, "The Colonial Origins of Comparative Development: An Empirical Investigation," *American Economic Review*, 2001, *91* (5), 1369–1401.

Acharya, Avidit, Matthew Blackwell, and Maya Sen, "The Political Legacy of American Slavery," *The Journal of Politics*, 2016, *78* (3), 621–641.

Alesina, Alberto and Nicola Fuchs-Schündeln, "Good-Bye Lenin (or Not?): The Effect of Communism on People's Preferences," *American Economic Review*, 2007, *97* (4), 1507–1528.

Alliierter Kontrollrat, "Gesetz Nr. 18: Wohnungsgesetz," 1946. Amtsblatt des Kontrollrats in Deutschland, Nr. 5, March 31, 1946. Full text at `https://www.verfassungen.de/de45-49/kr-gesetz18.htm`.

Angrist, Joshua D. and Jörn-Steffen Pischke, *Mostly Harmless Econometrics: An Empiricist's Companion*, Princeton: Princeton University Press, 2009.

Bailey, Martha J., Connor Cole, Morgan Henderson, and Catherine Massey, "How Well Do Automated Linking Methods Perform? Lessons from US Historical Data," *Journal of Economic Literature*, 2020, *58* (4), 997–1044.

Bauer, Franz J., *Flüchtlinge und Flüchtlingspolitik in Bayern 1945–1950*, Stuttgart: Klett-Cotta, 1982.

Bayer, Patrick and Christopher Timmins, "Estimating Equilibrium Models of Sorting Across Locations," *The Economic Journal*, 2007, *117* (518), 353–374.

_ , Fernando Ferreira, and Robert McMillan, "A Unified Framework for Measuring Preferences for Schools and Neighborhoods," *Journal of Political Economy*, 2007, *115* (4), 588–638.

Bazzi, Samuel, Martin Fiszbein, and Mesay Gebresilasse, "Frontier Culture: The Roots and Persistence of "Rugged Individualism" in the United States," *Econometrica*, 2020, *88* (6), 2329–2368.

Becker, Sascha O. and Hans-Joachim Voth, "From the Death of God to the Rise of Hitler," CESifo Working Paper 10730, CESifo 2023.

_ , Irena Grosfeld, Pauline Grosjean, Nico Voigtländer, and Ekaterina Zhuravskaya, "Forced Migration and Human Capital: Evidence from Post-WWII Population Transfers," *American Economic Review*, 2020, *110* (5), 1430–1463.

Bisin, Alberto and Thierry Verdier, "The Economics of Cultural Transmission and the Dynamics of Preferences," *Journal of Economic Theory*, 2001, *97* (2), 298–319.

Bleakley, Hoyt and Jeffrey Lin, "Portage and Path Dependence," *Quarterly Journal of Economics*, 2012, *127* (2), 587–644.

Braun, Sebastian T. and Timur Öztürk, "From Weimar to West Germany: Confessional Cleavages, Party Fragmentation, and Far-Right Voting," Kiel Working Paper 2324, Kiel Institute for the World Economy 2026.

Brown, Jacob R., Enrico Cantoni, Ryan Enos, Vincent Pons, and Emilie Sartre, "Sources and Extent of Rising Partisan Segregation in the U.S. – Evidence from 143 Million Voters," Working Paper 33422, National Bureau of Economic Research 2025.

Bühler, Mathias and Navid Sabet, "Female enfranchisement and the conservative voting gap: Evidence from gender-separated precinct returns," *European Journal of Political*

Economy, 2026, *93*, 102828.

Callaway, Brantly and Pedro H.C. Sant'Anna, "Difference-in-Differences with multiple time periods," *Journal of Econometrics*, 2021, *225* (2), 200–230. Themed Issue: Treatment Effect 1.

Cantoni, Davide and Noam Yuchtman, "Historical Natural Experiments: Bridging Economics and Economic History," in Alberto Bisin and Giovanni Federico, eds., *Handbook of Historical Economics*, Elsevier, 2021, pp. 213–241.

_ **, Felix Hagemeister, and Mark Westcott**, "Persistence and Activation of Right-Wing Political Ideology," 2020. Working paper.

Cantoni, Enrico and Vincent Pons, "Does Context Outweigh Individual Characteristics in Driving Voting Behavior? Evidence from Relocations within the United States," *American Economic Review*, 2022, *112* (4), 1226–1272.

Chetty, Raj and Nathaniel Hendren, "The Impacts of Neighborhoods on Intergenerational Mobility I: Childhood Exposure Effects," *The Quarterly Journal of Economics*, 2018, *133* (3), 1107–1162.

Chyn, Eric and Kareem Haggag, "Moved to Vote: The Long-Run Effects of Neighborhoods on Political Participation," *Review of Economics and Statistics*, 2023, *105* (6), 1596–1631.

Conley, Timothy G., "GMM estimation with cross sectional dependence," *Journal of Econometrics*, 1999, *92* (1), 1–45.

Dahl, Christian Møller, Torben S. D. Johansen, and Christian Vedel, "Breaking the HISCO Barrier: Automatic Occupational Standardization with OccCANINE," *arXiv preprint arXiv:2402.13604*, 2024. EHES Working Paper No. 255.

Damm, Anna Piil and Christian Dustmann, "Does Growing Up in a High Crime Neighborhood Affect Youth Criminal Behavior?," *American Economic Review*, 2014, *104* (6), 1806–1832.

Davis, Donald R. and David E. Weinstein, "Bones, Bombs, and Break Points: The Geography of Economic Activity," *American Economic Review*, 2002, *92* (5), 1269–1289.

Dell, Melissa, "The Persistent Effects of Peru's Mining Mita," *Econometrica*, 2010, *78* (6), 1863–1903.

Egner, Björn, "Wohnungspolitik seit 1945," *Aus Politik und Zeitgeschichte*, 2014, *64* (20–21), 13–19.

Fouka, Vasiliki and Hans-Joachim Voth, "Collective Remembrance and Private Choice: German–Greek Conflict and Behavior in Times of Crisis," *American Political Science Review*, 2023, *117* (3), 851–870.

Haus der Bayerischen Geschichte, "Wiederaufbauatlas: München," 2024. `https://hdbg.eu/wiederaufbau/orte/detail/muenchen/2`.

Heblich, Stephan, Alex Trew, and Yanos Zylberberg, "East-Side Story: Historical Pollution and Persistent Neighborhood Sorting," *Journal of Political Economy*, 2021, *129* (5), 1508–1552.

Historisches Lexikon Bayerns, "Flüchtlinge und Vertriebene," in "Historisches Lexikon Bayerns," Bayerische Staatsbibliothek, 2024. `https://www.historisches-lexikon-bayerns.de/Lexikon/Fl%C3%BCchtlinge_und_Vertriebene`.

Kohl, Sebastian, *Homeowner Nations or Nations of Tenants: How Historical Institutions in Urban Politics, Housing Finance and Construction Set Germany, France and the US*

on Different Housing Paths, Max Planck Institute for the Study of Societies, 2017.

Kuminoff, Nicolai V., V. Kerry Smith, and Christopher Timmins, "The New Economics of Equilibrium Sorting and Policy Evaluation Using Housing Markets," *Journal of Economic Literature*, 2013, *51* (4), 1007–1062.

Lambert, Paul S., Richard L. Zijdeman, Marco H. D. Van Leeuwen, Ineke Maas, and Kenneth Prandy, "The Construction of HISCAM: A Stratification Scale Based on Social Interactions for Historical Comparative Research," *Historical Methods*, 2013, *46* (2), 77–89.

Landeshauptstadt München, Sozialreferat, *100 Jahre Wohnungsamt 1911–2011*, Landeshauptstadt München, 2011.

Leeuwen, Marco H. D. Van, Ineke Maas, and Richard Zijdeman, *HISCO: Historical International Standard Classification of Occupations*, Leuven: Leuven University Press, 2002.

Manski, Charles F., "Identification of Endogenous Social Effects: The Reflection Problem," *The Review of Economic Studies*, 1993, *60* (3), 531–542.

Martin, Gregory J. and Steven W. Webster, "Does residential sorting explain geographic polarization?," *Political Science Research and Methods*, 2020, *8* (2), 215–231.

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

Ochsner, Christian and Felix Roesel, "Migrating Extremists," *The Economic Journal*, 2020, *130* (628), 1135–1172.

Rädlinger, Christine, *GEWOFAG 1928–2003*, Munich: Selbstverlag (GEWOFAG), 2003.

Redding, Stephen J. and Daniel M. Sturm, "Neighborhood Effects: Evidence from Wartime Destruction in London," Working Paper 32333, National Bureau of Economic Research 2024.

Schelling, Thomas C., "Dynamic Models of Segregation," *Journal of Mathematical Sociology*, 1971, *1* (2), 143–186.

Siegfried, André, *Tableau politique de la France de l'Ouest sous la Troisième République*, Paris: Armand Colin, 1913.

Spenkuch, Jörg L. and Philipp Tillmann, "Elite Influence? Religion and the Electoral Success of the Nazis," *American Journal of Political Science*, 2018, *62* (1), 19–36.

Tiebout, Charles M., "A Pure Theory of Local Expenditures," *Journal of Political Economy*, 1956, *64* (5), 416–424.

Voigtländer, Nico and Hans-Joachim Voth, "Persecution Perpetuated: The Medieval Origins of Anti-Semitic Violence in Nazi Germany," *The Quarterly Journal of Economics*, 2012, *127* (3), 1339–1392.

9 Figures and Tables

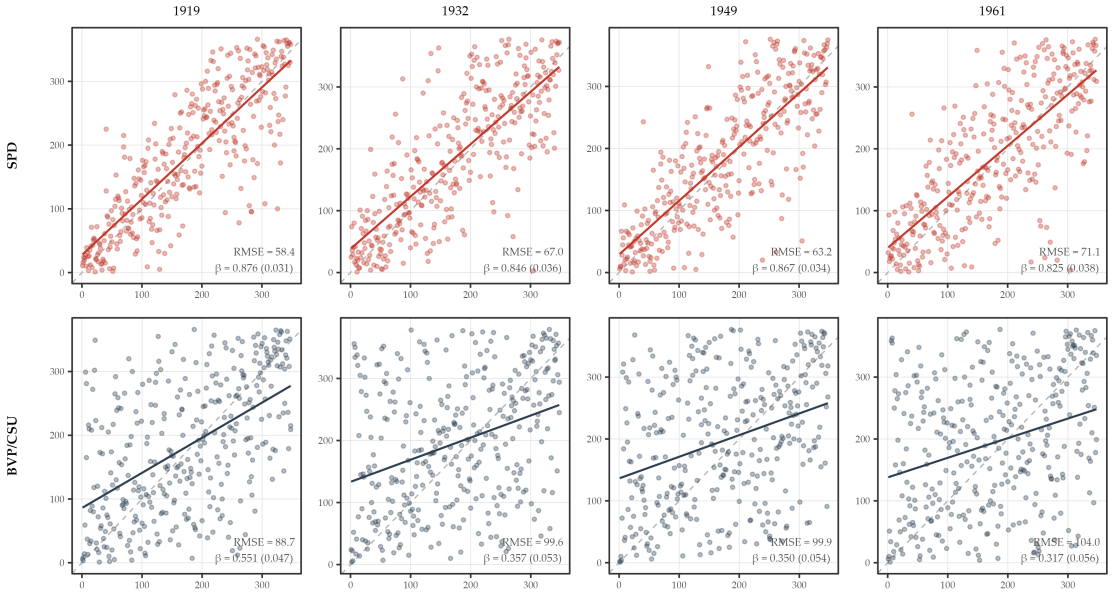

Figure 3: Rank Persistence of Precinct Vote Shares

Notes: Each panel plots the within-year percentile rank of a precinct's vote share in 1912 (x-axis) against its rank in the indicated election year (y-axis). Top row: SPD (Social Democrats). Bottom row: BVP/CSU (Bavarian People's Party and its post-war successor, the Christian Social Union). Columns span one election from each regime: 1919 (early Weimar), 1932 (late Weimar), 1949 (first post-war Federal election), and 1961. Dashed line: 45-degree line (perfect persistence). Solid lines: OLS fits. Each panel reports the root-mean-squared rank displacement (RMSE) and the OLS slope (β) with its standard error. Up to 348 precincts harmonized to 1933 boundaries; variation in N reflects missing returns in individual elections.

Table 1: Occupational Composition and the Persistent Spatial Vote

	SPD			BVP/CSU		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Weimar Republic (1919–1932)						
Composition, 1910	-0.057***	-0.040***	-0.032***	0.038***	0.026***	0.030***
	(0.004)	(0.006)	(0.006)	(0.006)	(0.009)	(0.009)
Lagged vote (1912)			0.176***			0.619***
			(0.036)			(0.071)
Within R^2	0.553	0.708	0.743	0.270	0.553	0.645
N	335	297	297	335	297	297
Panel B: Post-WWII (1949–1961)						
Composition, 1910	-0.041***	-0.024***	-0.019***	0.031***	0.021***	0.023***
	(0.004)	(0.005)	(0.005)	(0.003)	(0.004)	(0.004)
Lagged vote (1912)			0.122***			0.295***
			(0.035)			(0.071)
Within R^2	0.367	0.610	0.629	0.308	0.540	0.572
N	334	297	297	334	297	297
Covariates		✓	✓		✓	✓
Lagged 1912 vote			✓			✓
District FE	✓	✓	✓	✓	✓	✓

Notes: OLS regressions of period-average precinct vote share on *Composition, 1910*: the precinct-mean HISCAM in 1910, standardized. Columns 1–3 report the SPD; columns 4–6 the Catholic party (BVP in Panel A, CSU in Panel B). All columns include district (Stadtbezirk) fixed effects; columns 2–3 and 5–6 add the full pre-WWI amenity covariate set; columns 3 and 6 add the precinct's 1912 vote share for the same party. Within-district R^2. Robust (HC1) standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

(a) Precinct Retention θ_p *(b) θ_p decomposition*

Figure 4: Bounding Continuity: Address-Level Forward Tracking

Notes: The figure tracks every household head in one address-book edition forward to the next, measuring what share of a precinct's residents at $t-1$ reappear in the same precinct at t. A resident counts as retained if the same individual is found anywhere inside the precinct's 1933-harmonized boundaries in the following edition. Precinct retention θ_p is the population-weighted average across precincts. Panel (a) plots θ_p with a band bracketing match stringency. The lower bound (conservative) requires exact first and last name plus the same leading HISCO digit. The upper bound (loose) accepts a full-name Jaro–Winkler distance of at most 0.10 provided the HISCO leading digit shifts by no more than two. Both approaches apply a sequential cascade. The 1935 — 1953 pair spans eighteen years and is computed on the subsample of 1933 addresses that survived into the post-war book (92.1% of addresses, covering 95.0% of 1933 residents). It is an upper bound on cross-war continuity because residents of destroyed addresses are excluded from the denominator. Retention under the conservative bound is between ≈ 0.43–0.62 in the decades of the Empire and Weimar Republic. Even from one year to the next, almost one third to a half of residents do not stay within a precinct. This rate falls to $\theta_p \approx 0.07$ across WW2, and bounds how much of the persistent spatial vote pattern the same residents can carry. Panel (b) decomposes the conservative θ_p into same-address stayers and within-precinct movers: retention is almost entirely same-address staying. The horizontal axis is ordered by address-book pair.

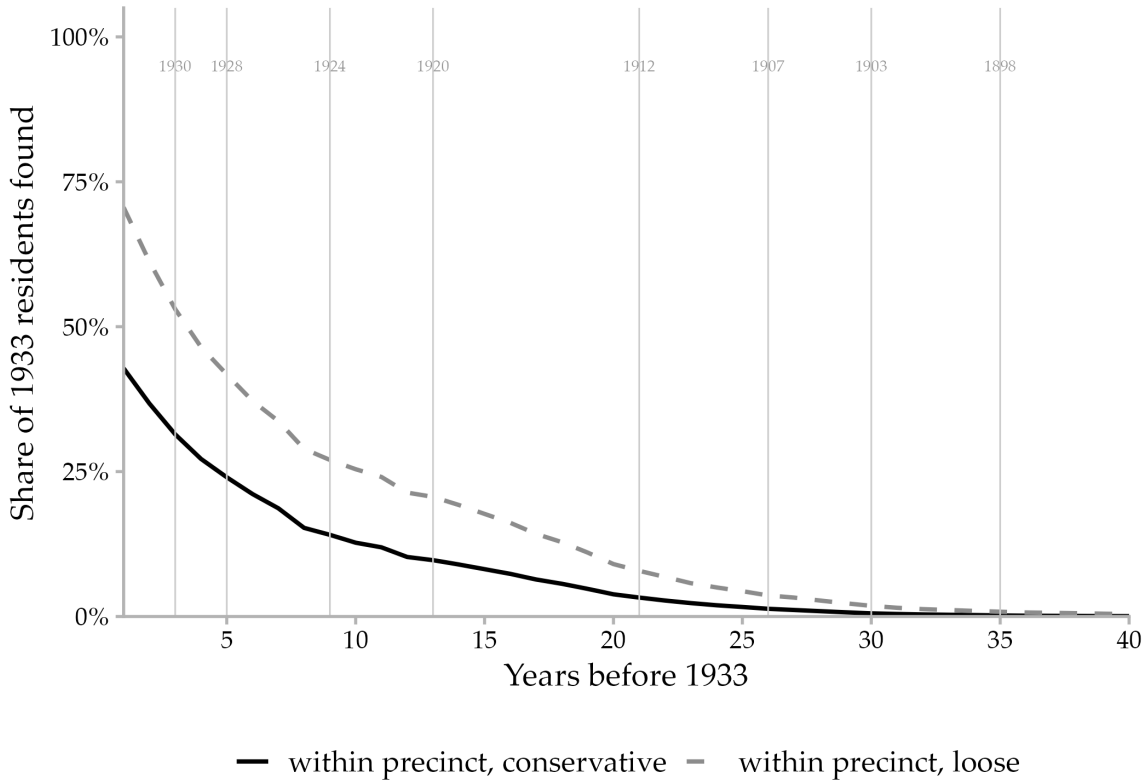

Figure 5: The Transmission Substrate Is Thin: Backward Residence Trace from 1933

Notes: Each panel traces a cohort of renters backward from the anchor year. For each address-book edition prior to that year, the figure plots the share of anchor-year precinct residents who can be found in the same precinct τ years earlier (direct match, no chaining; same-precinct includes both same-address stayers and address-changers within the precinct boundary). The denominator is all renters in 1933 with occupation and geocode ($N = 171{,}044$). Owner-occupiers are excluded throughout; they are the one within-city source of deep residential continuity, so their exclusion is conservative: the figure understates the stock of long-tenured residents relative to an all-residents count. Two match stringencies bracket the result. Both curves decay steeply as more years pass, with the conservative line falling faster. The modal renter in both cohorts cannot be traced back more than a few years: most residents of any anchor-year precinct are recent arrivals. The share who had been in the same precinct long enough to accumulate the tenure, which a teaching relationship requires, is thin.

Table 2: Conditional Logit Estimates of Mover Destination Choice

	(1)	(2)	(3)	(4)	(5)
Panel A: Same-period S_j					
Type match $(-\lvert\tau_i - S_j\rvert)$	0.303***	0.323***			0.338***
	(0.012)	(0.011)			(0.010)
Type match $(-\lvert\tau_i - S_j\rvert) \times S_j$			0.262***	0.251***	
			(0.007)	(0.007)	
log distance	-0.926***	-0.939***	-0.926***	-0.938***	-0.931***
	(0.014)	(0.014)	(0.014)	(0.014)	(0.016)
N	17,300,296	17,300,296	17,300,296	17,300,296	17,300,296
Panel B: Fixed 1910 S_j^{1910}					
Type match $(-\lvert\tau_i - S_j\rvert)$	0.258***	0.274***			0.283***
	(0.013)	(0.012)			(0.011)
Type match $(-\lvert\tau_i - S_j\rvert) \times S_j$			0.218***	0.210***	
			(0.007)	(0.007)	
log distance	-0.927***	-0.941***	-0.927***	-0.939***	-0.933***
	(0.014)	(0.014)	(0.014)	(0.014)	(0.016)
Amenity controls		✓		✓	✓
Destination-district FE					✓
Mover FE	✓	✓	✓	✓	✓
N	16,908,584	16,908,584	16,908,584	16,908,584	16,908,584

Notes: Conditional logit over the full precinct choice set, using annual Weimar pairs (1919–1933). The outcome is the chosen destination precinct. τ_i is the mover's HISCAM score, S_j is the destination precinct's mean HISCAM. Columns (1) and (2) use the homophily match $-\lvert\tau_i - S_j\rvert$; columns (3) and (4) interact that match with S_j, $(-\lvert\tau_i - S_j\rvert) \times S_j$. All specifications include mover fixed effects. Columns (2) and (4) add LASSO-selected amenity controls (10-fold CV, `lambda.1se`, regressing 1910 precinct composition on 32 physical amenities). Panel A uses same-period composition (pre-move year); Panel B uses fixed 1910 pre-war composition, identical across all Weimar pairs, testing whether movers sort on the durable type-map. Column (5) adds destination-district (Stadtbezirk) fixed effects to the amenity homophily specification, so the type-match coefficient is identified purely from within-district variation across precincts. N is Movers × Destination Precincts. Table E.3 reports the results for the origin precinct dropped from the choice set. Standard errors clustered by origin precinct. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

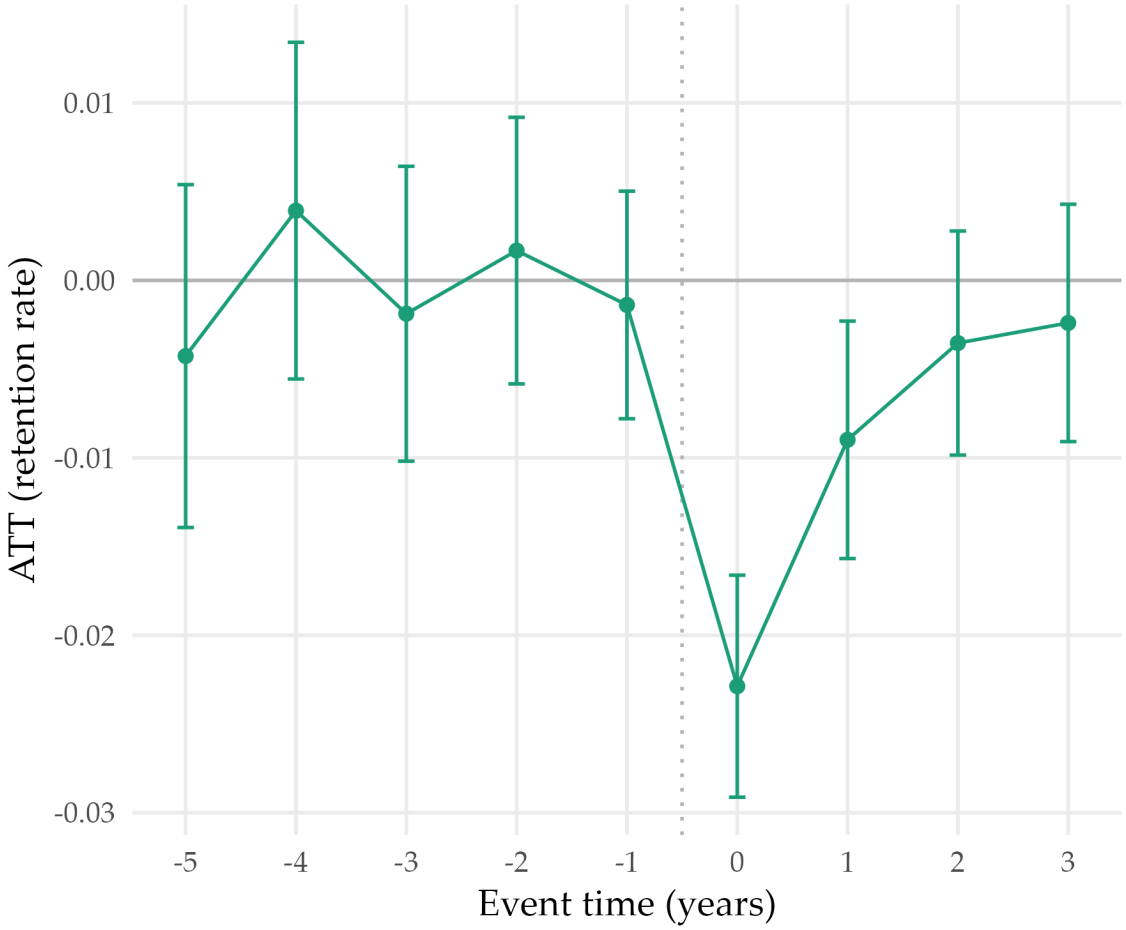

Figure 6: Owner Change Displaces Tenants during the Weimar Republic

Notes: Callaway and Sant'Anna (2021) event study of tenant retention around a building's first owner change using Weimar address books (1919–1933) and not-yet-treated units as control group. Standard errors are clustered by address and 95% confidence intervals reported in the figure. The outcome is one-period tenant retention. Pre-trends are flat across the five pre-periods. The owner change induces a drop of about 2.2 percentage points at $t = 0$ that shrinks to about one point at $t = 1$ and is near zero from $t = 2$. This is the first stage for the push instrument: owner change is an exogenous shock to whether a tenant leaves (not where they land), used to test whether forced movers sort to type like voluntary movers.

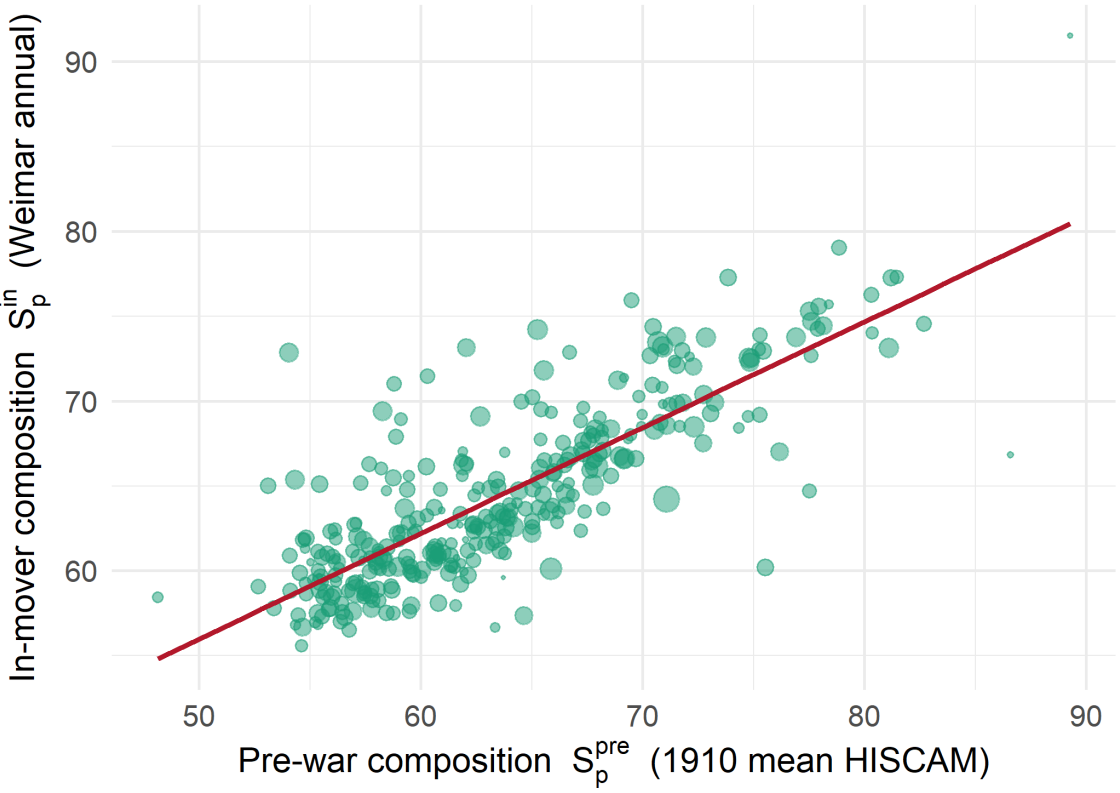

Figure 7: In-Movers Reconstitute the Pre-War Type-Mix

Notes: Each point is a precinct. The horizontal axis is the precinct's 1910 pre-war composition (mean HISCAM); the vertical axis is the mean HISCAM of its Weimar *in-movers* (residents who arrived from another precinct). The strong positive association ($\rho = 0.81$, in-mover-weighted 0.79) shows that the newcomers refilling a precinct reproduced its pre-war type-mix, the composition object that connects the sorting mechanism back to the persistent vote. This is a relationship between two composition measures and contains no vote variable. Point size is in-mover count; the line is the OLS fit.

Table 3: Sorting at Both Levels: Nested District–Precinct Decomposition

	Between-district	Within-district
Type-match γ	0.214^{***}	0.329^{***}
	(0.020)	(0.012)
SD(type-match)	0.888	0.835
Std. effect $\gamma\times$SD	0.190	0.275
Mover FE	✓	✓
N	48,080	60,983

Notes: Nested decomposition of the homophily logit into two sequential stages: between Munich's 30 districts, which district the mover picks, and within the chosen district, which of its precincts the mover picks. The within-stage logit conditions on the chosen district, so type-match is measured only on within-district variation in composition. γ is the homophily coefficient $-|\tau_i - S_j|$; the standardized effect is $\gamma \times \text{SD}(\text{type-match})$. Both stages show positive, significant sorting; the within-district standardized effect exceeds the between-district effect. Sorting at the precinct level is not reducible to district-level selection. Mover fixed effects included. Standard errors clustered by origin district. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Forced Movers Sort Similarly to Voluntary Movers: Annual γ-Contrast, Weimar

	Voluntary (control bldg.)	Pushed (owner-change bldg.)	$\Delta\gamma$
Panel A. Same-period S_j, no amenities			
Type-match γ	0.298***	0.366***	0.068
	(0.012)	(0.032)	
Panel B. Same-period S_j, + LASSO amenities			
Type-match γ	0.320***	0.375***	0.056
	(0.011)	(0.032)	
Panel C. Fixed 1910 S_j^{1910}, + LASSO amenities			
Type-match γ	0.266***	0.373***	0.106
	(0.012)	(0.032)	
Panel D. Same-period S_j, + LASSO amenities + dest.-district FE			
Type-match γ	0.335***	0.387***	0.052
	(0.010)	(0.031)	
N	50,642	3,875	

Notes: The same conditional logit as Table 2, estimated separately on movers leaving owner-change buildings ("pushed," $N = 3{,}875$) and on movers leaving control buildings ("voluntary," $N = 50{,}642$), using annual Weimar pairs (1919–1933). γ is the homophily type-match coefficient $-|\tau_i - S_j|$ and $\Delta\gamma = \gamma_{\text{push}} - \gamma_{\text{vol}}$. The owner change exogenizes *whether* a tenant leaves, not *where* they land, so this contrast tests the pull (sorting), holding the push exogenous. Panel A uses same-period S_j, no amenity controls. Panel B uses same-period S_j and LASSO amenity controls. Panel C uses fixed 1910 pre-war S_j^{1910} and LASSO amenity controls. Panel D uses same-period S_j, LASSO amenity controls and destination-district fixed effects (which absorb district-level demand shocks correlated with ownership change). Forced movers sort more strongly than voluntary movers in every specification: $\Delta\gamma$ is positive throughout, from 0.052 (Panel D) to 0.106 (Panel C). Standard errors clustered by origin precinct.* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5: What Predicts Administrative Placement? Pre-War Correlates of the 1953 Admin Share

	(1) OLS	(2) OLS	(3) LASSO
Average precinct status, 1910	-0.003	0.000	–
	(0.011)	(0.013)	
SPD vote share, 1912		0.045	–
		(0.097)	
Amenities selected by the LASSO			
Population			-0.007
Building height			-0.004
Distance to cafe			0.001
Pre-war amenities (32 controls)	✓	✓	
District FE	✓	✓	✓
Within R^2 (OLS)	0.183	0.184	
Cross-validated R^2 (λ_{1se} / λ_{min})			0.032 / 0.060
Mean dep. var.	0.096	0.096	0.098
N	317	317	319

Notes: The outcome is the precinct's 1953 Wohnungsamt-allocated housing share, winsorized at the 90th percentile. Columns 1 and 2 are OLS regressions on pre-war characteristics with district (Stadtbezirk) fixed effects, so the coefficients and the within-R^2 use within-district variation alone. *Average precinct status, 1910* is the standardized precinct-mean HISCAM; *SPD vote share, 1912* is the precinct's pre-war SPD share. Both columns also include the 32 pre-war physical amenities (the sorting model's candidate set, Table E.4), which alone give a within-R^2 of 0.183. Column 3 is a within-district LASSO (Frisch–Waugh: the outcome and all candidates are demeaned by district, then `cv.glmnet`, $\alpha = 1$, ten-fold) over all 34 pre-war candidates: the two persistence variables (composition and the 1912 vote, shown above) and the 32 amenities. Given a free choice, it discards both persistence variables (marked –) and keeps only the 3 amenities listed, shown with their penalized LASSO coefficients (admin-share points per standard deviation); the LASSO yields penalized point estimates, not standard errors. The low cross-validated R^2 (reported in the table) is the point: within district, almost nothing about a precinct's pre-war character predicts where housing was later administratively placed. Robust standard errors for the OLS columns in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 6: Administrative Placement Changed Who Sorted In: Newcomer Composition and the Prior Type-Map

	Newcomers 1953			Newcomers 1933 (placebo)		
	(1)	(2)	(3)	(4)	(5)	(6)
Prior composition × Admin share	-0.519***	-0.461***	-0.382**	-0.443	-0.259	-0.046
	(0.146)	(0.142)	(0.160)	(0.324)	(0.306)	(0.218)
Prior composition	0.789***	0.755***	0.713***	0.742***	0.666***	0.684***
	(0.020)	(0.027)	(0.034)	(0.029)	(0.033)	(0.045)
Admin share	33.289***	29.294***	23.530**	28.840	16.967	2.656
	(9.168)	(8.975)	(10.361)	(19.406)	(18.543)	(13.394)
Prior map	1933	1933	1933	1910	1910	1910
LASSO amenities			✓			✓
District FE		✓	✓		✓	✓
R^2	0.868	0.786	0.803	0.749	0.657	0.728
N	335	333	317	336	334	317

Notes: OLS regressions of *newcomer* precinct-mean HISCAM (raw points) on the precinct's prior composition, its 1953 administrative share, and their interaction. A newcomer is a person (linked by household identifier across address-book editions) present in a precinct in the later year but not in that same precinct in the earlier year. Columns 1–3 take 1953 newcomers against the 1933 composition; columns 4–6 take 1933 newcomers against the 1910 composition, a placebo window before the Wohnungsamt existed. *Admin share* is the 1953 Wohnungsamt-allocated housing share, winsorized at the 90th percentile. The negative interaction shows that newcomers reproduce the prior occupational type-map more weakly where the administrative share is higher: administrative placement changed who sorted in. The interaction is significant in the real window (columns 1–3) and insignificant in the placebo (columns 4–6), so it is not a pre-war sorting trend. Columns 3 and 6 add the 15 LASSO-selected pre-war amenities; columns 2–3 and 5–6 add district fixed effects. Standard errors clustered at the precinct in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

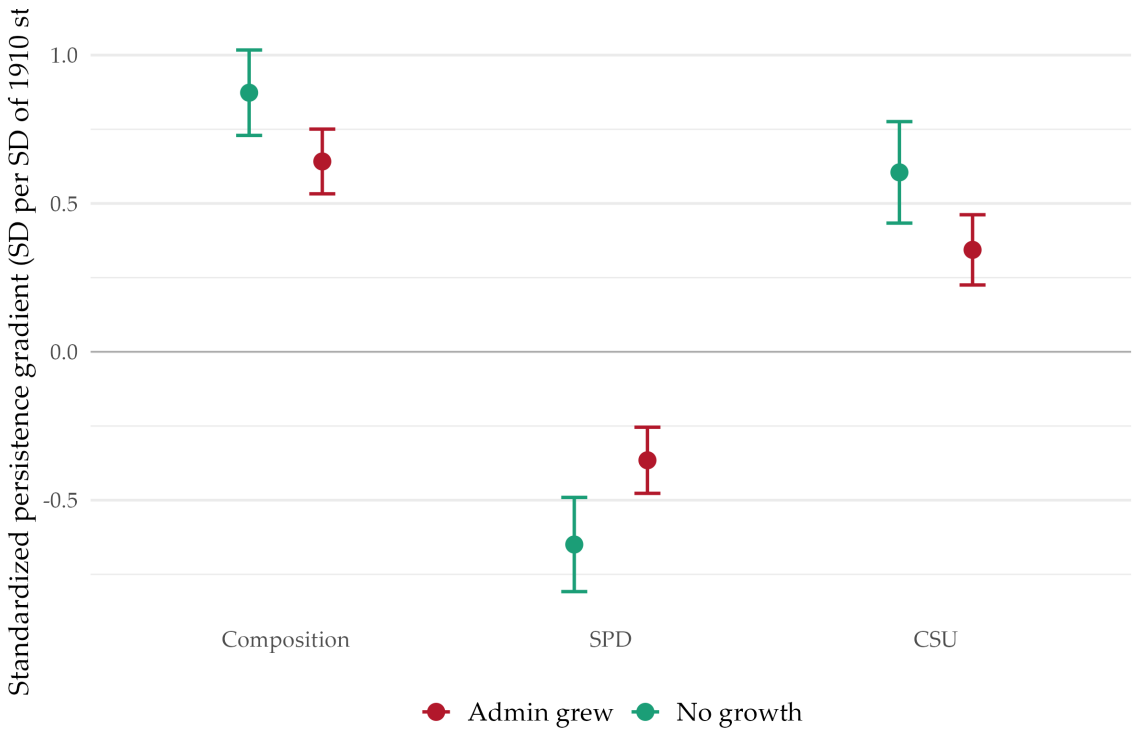

Figure 8: Administrative Placement Attenuates Persistence Across Composition and Both Votes

Notes: The figure plots the standardized gradient linking a precinct's 1910 occupational composition (standardized precinct-mean HISCAM) to three post-war (1949–1961) outcomes: occupational composition, the SPD vote, and the CSU vote. Each gradient is estimated separately in precincts where the share of administratively owned housing (government, non-profit-housing, and cooperative) did not grow over 1910–1953 (*No growth*, $n = 89$) and in those where it did (*Admin grew*, $n = 246$). The split is the sign of ΔAdmin = the 1953 administrative share minus its 1910 value, so it isolates the post-war build-out of administratively allocated housing net of the pre-war ownership map; computing ΔAdmin requires a valid 1910 baseline, so the sample here (335 precincts) is one smaller than the 336 with a valid 1953 administrative share alone. The two subgroup gradients are estimated jointly in a single split-sample regression of the standardized outcome on the 1910 composition interacted with the subgroup indicator, controlling for the 15 LASSO-selected pre-war amenities (Table E.4) and district fixed effects; bars are 95% confidence intervals from robust standard errors. Outcomes are standardized so the three gradients share a common axis, one standard deviation of the outcome per standard deviation of 1910 composition. In every outcome the gradient is flatter where administrative ownership grew: where post-war housing was administratively placed rather than self-sorted, the pre-war type-map is reproduced more weakly. The attenuation is individually significant in all three outcomes; the corresponding regressions, with the continuous $\text{Status}_{1910} \times \Delta\text{Admin}$ interaction, are in Tables D.1 and D.2. The administrative share is an ownership object rather than a pre-assigned instrument, so this is descriptive moderation consistent with placement switching sorting off. The timing evidence is the newcomer-sorting test of Table 6, where the interaction is significant in the post-war window (column 3) and indistinguishable from zero in the pre-Wohnungsamt window (column 6).

Online Appendix

Not intended for publication.

This appendix contains the following sections:

A. Re-Anchoring the Persistence Baseline: the rank persistence anchored at 1919, the first universal-suffrage election, and at 1893, the earliest election in the panel.

B. Persistence into the Present: the 1912 spatial ordering carried to the 2023 Bavarian state election.

C. Post-War Housing Allocation in Munich: institutional detail, identifying assumptions, and threats.

D. Administrative Placement and Persistence: the continuous ΔAdmin interaction estimates behind the split-sample figure, the change-on-change link between composition and the vote, and the church-proximity heterogeneity of the moderation.

E. The Grain of Residential Sorting: the amenity LASSO, the owner-change first stage, the district-grain push contrast, the origin-in-choice-set check, and the occupational-change exhibit.

G. Summary Statistics.

A Re-Anchoring the Persistence Baseline

Figure 3 anchors the rank persistence of precinct vote shares at 1912. This appendix moves the anchor in both directions. Section A.1 re-anchors at 1919, the first election under the post-war franchise, and adds the Liberals as a third row. Section A.2 re-anchors at 1893, the earliest election in the panel, and stretches the horizon to 68 years. The three anchors together establish that the persistence fact is a property of the political geography, not of the choice of baseline year.

A.1 Anchoring at 1919

The 1912 baseline spans three breaks at once. Between 1912 and the first Weimar election the franchise expanded to women and moved to proportional representation, the Catholic-conservative vote regrouped from the Zentrum into the BVP, and the socialist vote fragmented as the USPD and then the KPD broke from the SPD. A slope measured from 1912 mixes the persistence of the spatial ordering with these one-time shifts. The 1919 election is the first held under universal suffrage and proportional representation and the first in which the BVP ran, so anchoring there draws the baseline and outcome elections from the same franchise and party system. The slope then measures how stable the spatial ordering of the vote was rather than the size of the 1918–19 institutional break.

The SPD's slope barely moves with the anchor, running from about 0.82–0.84 on the 1919 baseline against about 0.83–0.87 on the 1912 baseline. The 1912 SPD share bundles the later SPD with the USPD and KPD that split from it after the war, yet separating those parties at 1919 leaves the slope almost unchanged, so the spatial ordering of the socialist vote survived the fragmentation intact. The Catholic bloc moves the most. Its slope roughly doubles, from about 0.32–0.36 on the 1912 baseline to about 0.59–0.73 on the 1919 baseline. The 1912 baseline counts the Zentrum-to-BVP regrouping as geographic change, understating the bloc's persistence. The liberals persist as strongly as the SPD, their slope holding between

about 0.70–0.82 even as the liberal vote share fell below 2% by the early 1930s. The spatial ordering of their support outlasted the party itself.

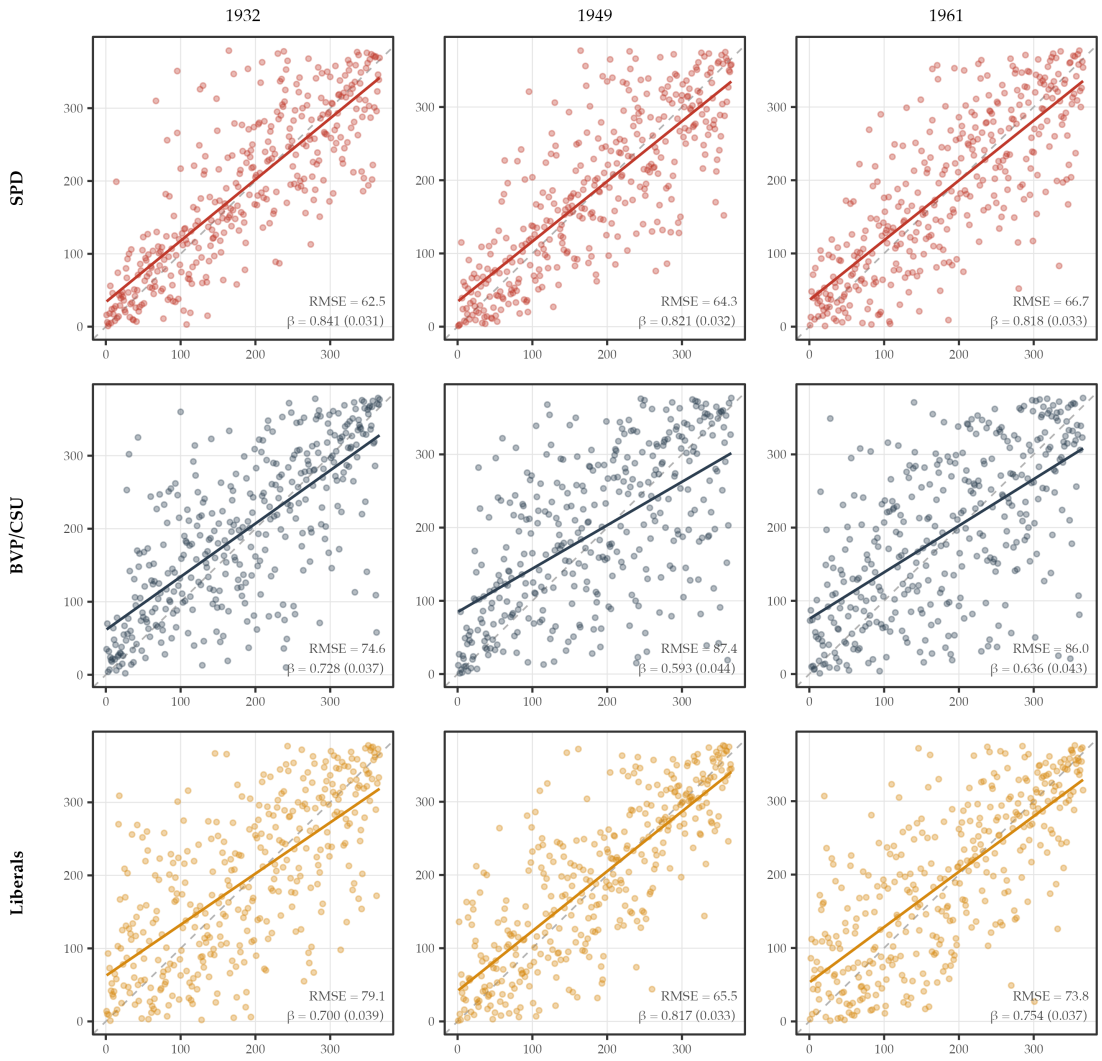

Figure A.1: Rank Persistence of Precinct Vote Shares, Anchored at 1919

Notes: Each panel plots the within-year percentile rank of a precinct's vote share in 1919 (x-axis) against its rank in the indicated election year (y-axis). Rows: SPD, BVP/CSU (Bavarian People's Party and its post-war successor, the Christian Social Union), and the Liberals. Columns span one election from each regime: 1932 (late Weimar), 1949 (first post-war Federal election), and 1961. 1919 is the first election under universal suffrage and proportional representation, and the first contested by the BVP, which makes it the first baseline comparable to the post-1919 outcomes. Dashed line: 45-degree line (perfect persistence). Solid lines: OLS fits. Each panel reports the root-mean-squared rank displacement (RMSE) and the OLS slope (β) with its standard error. Up to 366 precincts harmonized to 1933 boundaries; variation in N reflects missing returns in individual elections.

A.2 Anchoring at 1893

The 1893 election is the earliest in the panel, 68 years before the last outcome. Anchoring there stretches the baseline in the direction opposite to Section A.1. Instead of holding the franchise and party system fixed, the 1893 anchor crosses every break at once: the franchise expansion, the party regroupings, and both wars. The 1893 returns record three parties, the SPD, the Zentrum, and the National Liberals. The National Liberals are not the left-liberal lineage that the Liberal series tracks from 1898 onward, so Figure A.2 shows the SPD and the Catholic bloc only. The returns also cover only the 269 precincts of the 1893 city, so within each panel both years are ranked among the precincts observed in both. To trace the decay election by election, the columns add 1912 and 1919 ahead of the three outcome years of Figure A.1.

The SPD ordering decays smoothly with the horizon. Its slope falls from 0.76 against 1912 to 0.69 against 1919, 0.63 against 1932, and 0.59 against 1961. The decline per decade is roughly constant, and the crossings of the franchise break and of the war cost no more than the peacetime intervals around them. The Catholic bloc fades faster. Its slope is 0.41 against 1912, within the Zentrum era, and falls to 0.21–0.23 after the Second World War. Both pre-war anchors cross the Zentrum-to-BVP regrouping, which Section A.1 shows roughly halves the measured slope: the same post-war outcomes anchored at 1919 give 0.59–0.73.

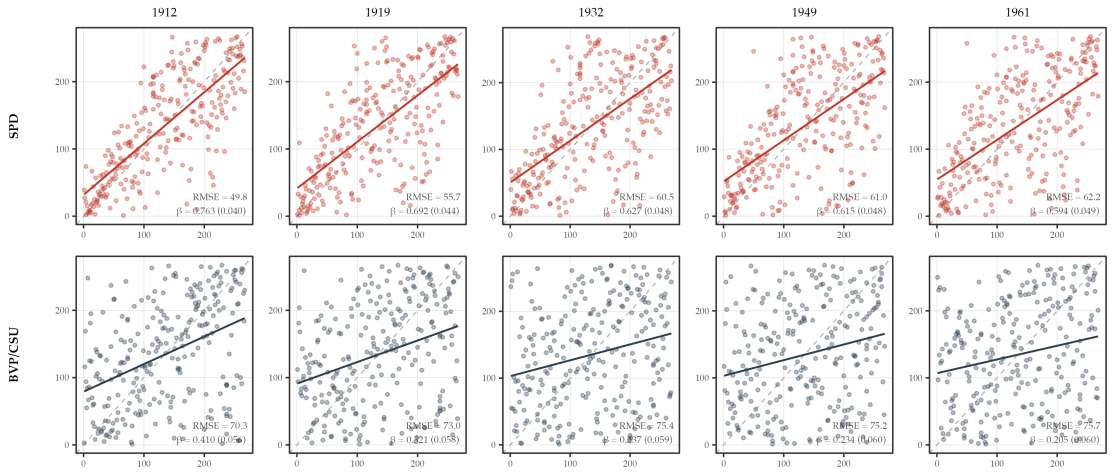

Figure A.2: Rank Persistence of Precinct Vote Shares, Anchored at 1893

Notes: Each panel plots the rank of a precinct's vote share in 1893 (x-axis) against its rank in the indicated election year (y-axis). Rows: SPD and BVP/CSU (the Zentrum in 1893 and 1912, the Bavarian People's Party in the Weimar columns, the Christian Social Union after the war). The 1893 returns record the SPD, the Zentrum, and the National Liberals; the National Liberals are not the left-liberal lineage of the Liberals row in Figure A.1, so that row is omitted. Columns: 1912 (the main baseline), 1919 (the first universal-suffrage election), 1932, 1949, and 1961. The 1893 returns cover the 269 precincts of the 1893 city (harmonized to 1933 boundaries); within each panel, ranks are computed among the precincts observed in both elections, so the two axes share a scale. Dashed line: 45-degree line (perfect persistence). Solid lines: OLS fits. Each panel reports the root-mean-squared rank displacement (RMSE) and the OLS slope (β) with its standard error.

A.3 Vote Shares, Anchored at 1919

Rank persistence is ordinal: a precinct can hold its position in the distribution while its vote share moves substantially, or the reverse. Figure A.3 repeats the 1919-anchored comparison in levels, plotting each precinct's 1919 vote share (x-axis) against its share in the same three elections (y-axis). SPD and BVP/CSU persist in levels about as strongly as in ranks: slopes of 0.46–0.53 for SPD and 0.48–0.77 for BVP/CSU, with root-mean-squared deviations of 5–6 percentage points. The Liberals are the exception. Their 1932 level-slope collapses to 0.05, against a rank slope of 0.70 in the same cell: the party's vote fell to a few percent almost everywhere in the Nazi landslide, so relative position survived while the level gap it measured nearly disappeared. The gap reopens once the party revives after the war (0.36 in 1949, 0.21 in 1961).

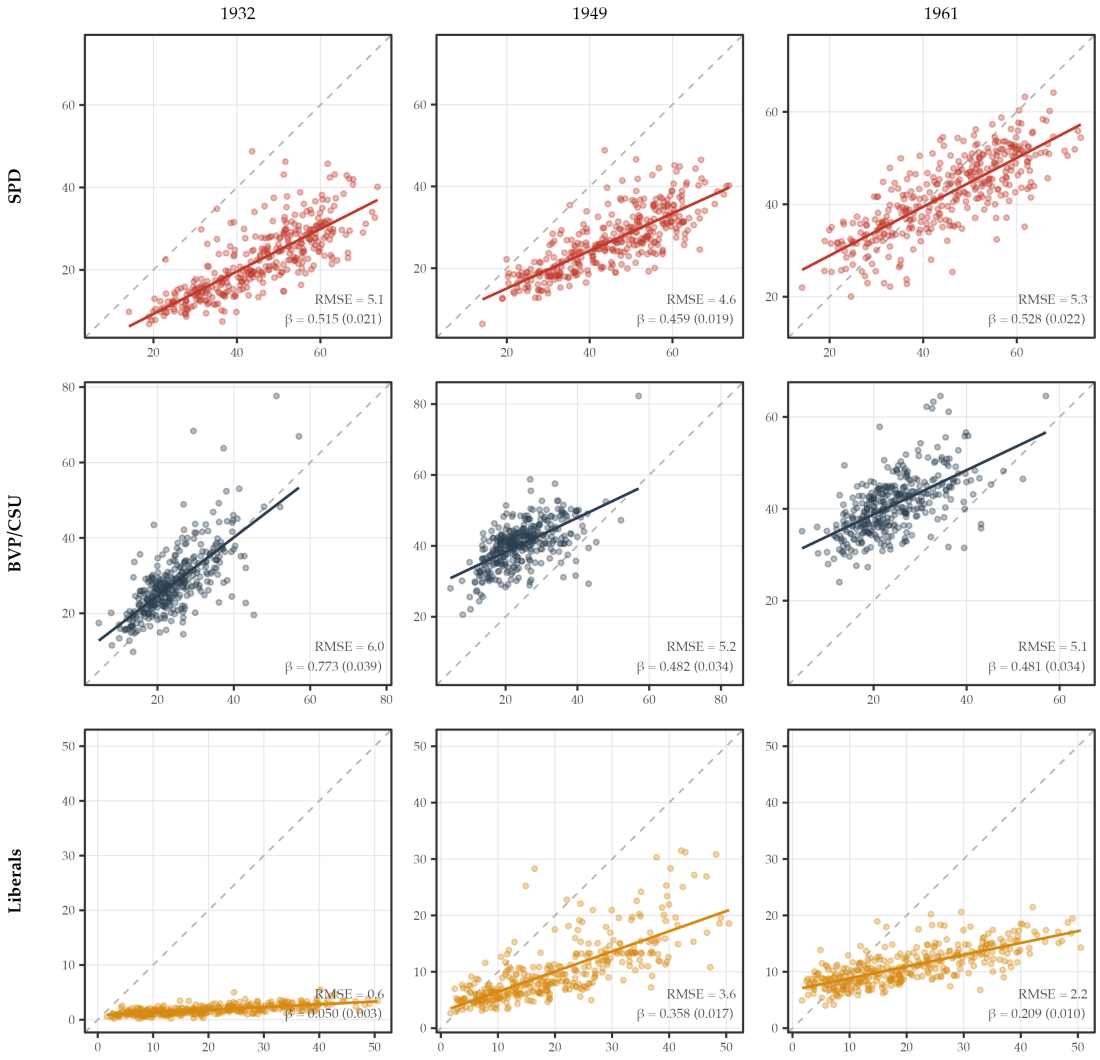

Figure A.3: Vote-Share Persistence of Precincts, Anchored at 1919

Notes: Each panel plots a precinct's vote share in 1919 (x-axis, percentage points) against its share in the indicated election year (y-axis). Rows: SPD, BVP/CSU (Bavarian People's Party and its post-war successor, the Christian Social Union), and the Liberals. Columns span one election from each regime: 1932 (late Weimar), 1949 (first post-war Federal election), and 1961. Dashed line: 45-degree line (no change in vote share). Solid line: OLS fit. Each panel reports the root-mean-squared deviation (RMSE, in percentage points) and the OLS slope (β) with its standard error. Axes are scaled to a common range within each panel so the 45-degree line is a true diagonal; the range therefore differs by party. Up to 366 precincts harmonized to 1933 boundaries; variation in N reflects missing returns in individual elections.

B Persistence into the Present

The persistence the paper documents extends through 1961. The 2023 Bavarian state election (Landtagswahl) lets us carry it forward six decades more. We map each 2023 polling precinct to a 1933 precinct through the address. The published street directory assigns every house number to a 2023 precinct, and the 1933 address book assigns the same addresses to 1933 precincts, so the two precinct systems meet at the building. Re-aggregating the 2023 party-list votes onto the 1933 units, weighted by the number of 1933 households at each address, reproduces Figure 3 for the modern vote (Figure B.1).

The correlation over more than a century is weaker than one documented until 1961. Part of the decline could lie in the party labels rather than the places, so the right column of Figure B.1 compares ideological blocs. The 1912 Social Democratic vote has since split across several parties of the left. The left bloc, SPD plus Grüne, tracks the 1912 SPD map at 0.34, above the SPD alone, consistent with the left vote holding its location as it spread across parties (upper panel). The right shows no such recovery. The conservative bloc, CSU plus Freie Wähler plus AfD, tracks the 1912 BVP map at 0.13, no higher than the CSU alone (lower panel). The geography that organized the Catholic vote before the First World War does not organize the modern right.

Three caveats bound the comparison. First, the directory is published for federal elections (Bundestagswahlen) as well, which could add comparison years through 1977, though a fifty-year data-protection seal on the records blocks anything more recent without special permission. Second, a majority of the 2023 ballots (54%) were cast by post and carry no precinct of residence, so the figure uses in-person votes alone, and postal voters need not resemble those who voted in person. Third, the bloc assignments are judgment calls, the place of the AfD most of all. With these caveats, Figure B.1 bounds the paper's claim in time. The spatial structure of the vote, sharp through mid-century, has faded by 2023, less for the left than the right.

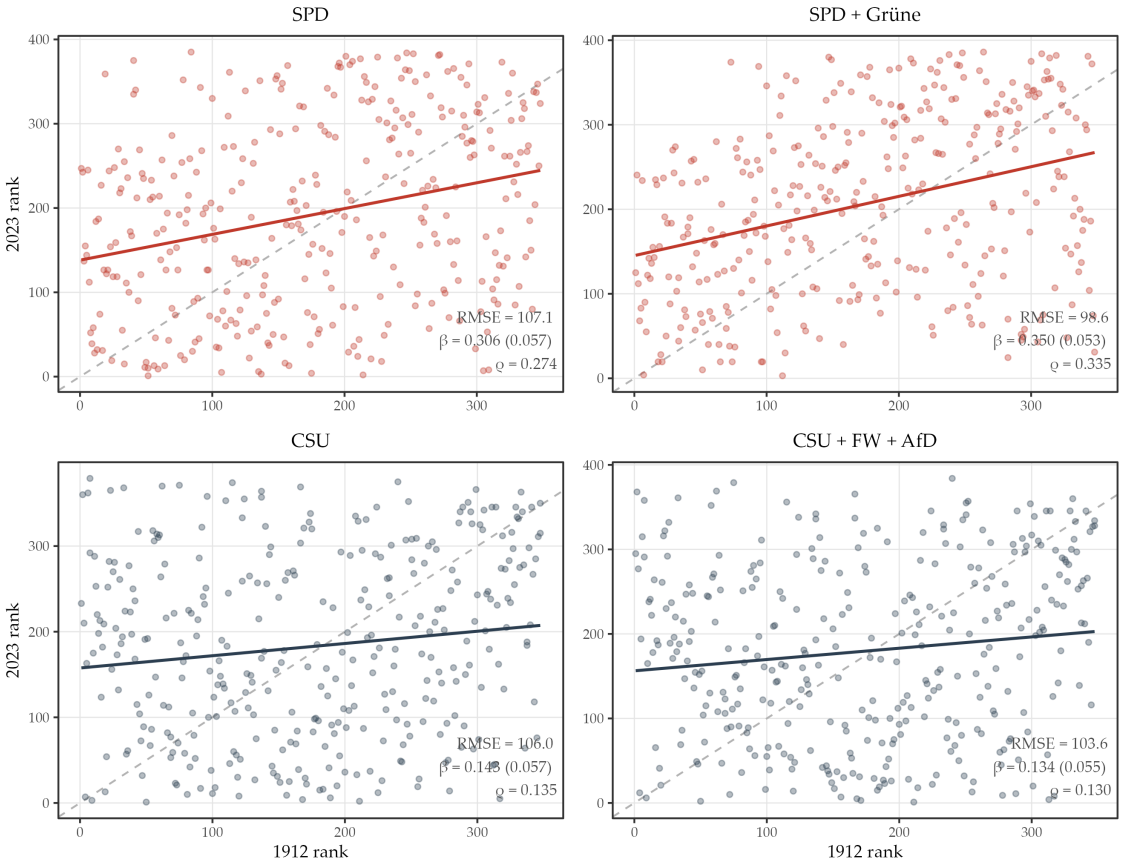

Figure B.1: Rank Persistence of the Precinct Vote, 1912 to 2023

Notes: Each point is a 1933 precinct. The horizontal axis is its 1912 vote rank and the vertical axis its 2023 rank. The 2023 party-list votes (Zweitstimmen) are re-aggregated onto the 1933 precinct units through shared addresses and weighted by 1933 households (in-person votes only). Top row: the SPD lineage. Bottom row: the BVP lineage. Left column: the single successor party (SPD; CSU). Right column: the ideological bloc (left, SPD plus Grüne; right, CSU plus Freie Wähler plus AfD). Each panel reports the rank-rank slope and the Spearman correlation ρ. Dashed line: the 45-degree line.

C Post-War Housing Allocation in Munich

The administrative-placement contrast (Figure 8) requires that residents in administratively allocated housing were placed independently of the pre-war characteristics that predict the vote. Two conditions must hold: assignment was driven by need and availability, not by neighborhood characteristics, and applicants could not select into neighborhoods that matched their political type. This appendix documents the institutional basis for both.

Assignment was need-based and availability-driven

Munich's Wohnungsamt allocated housing through a ranked *Dringlichkeitsliste* (urgency list) ordering all applicants by severity of need (Landeshauptstadt München, Sozialreferat, 2011). Under Kontrollratsgesetz Nr. 18 (1946), applicants were assigned to specific addresses through *Einweisungsverfügungen* (billeting orders) whenever a vacancy was identified (Alliierter Kontrollrat, 1946). The priority hierarchy was codified in Article VIII: victims of National Socialism first, then the disabled and elderly, large families, and essential workers. Article VIII(1)(d) prohibited preferential treatment based on social or financial standing. By 1948, over 93,000 applications competed for roughly 12,000 annual assignments, and refusal of an assigned unit cost one's place on the list. From 1946, most of the existing housing stock fell under this compulsory management, with a free market largely confined to housing exempted by statute, such as new construction completed after the 1948 currency reform; Munich retained the controls, gradually dismantled from 1960, until 1968 (Egner, 2014).

The assignment rule matched applicants to vacancies, not applicants to neighborhoods. What determined where a person lived was which unit happened to become available when their urgency ranking was reached (Kohl, 2017). Administrative courts upheld these decisions in 98.7% of appeals (1950–1952), so applicants had effectively no recourse to override their assignment (Landeshauptstadt München, Sozialreferat, 2011). A *Zuzugssperre* (ban on in-migration) further constrained self-selection into Munich itself. From August 1945 through

the early 1950s, only Soviet-zone refugees, Berlin residents, and ethnic German expellees could enter the city.

Assignment was independent of the pre-war amenity bundle

The amenities that predict the vote (pub density, church proximity, housing density, factory distance) reflect pre-war neighborhood character. The Wohnungsamt's criteria (urgency, family size, housing loss, employment status) are orthogonal to that bundle by construction. Nothing in the priority hierarchy references the physical characteristics of the assigned neighborhood. Table 5 confirms this within district: pre-war composition and the pre-war vote do not predict a precinct's 1953 administrative share.

Two institutional features reinforce the independence. First, reconstruction under Stadtbaurat Karl Meitinger preserved the medieval street plan and rebuilt destroyed buildings in place (Haus der Bayerischen Geschichte, 2024), so the pre-war amenity landscape was largely intact when the Wohnungsamt made its assignments. Second, the office assigned applicants across the entire city, housing refugees and bombed-out families wherever intact structures existed while actively preventing spatial concentration (Historisches Lexikon Bayerns, 2024). The resulting distribution of administrative housing reflects the geography of surviving and reconstructed buildings, not the geography of political preference.

Threats to the identifying assumption

Two institutional features could compromise this independence. Under §4 of the Erstes Wohnungsbaugesetz (1950), housing cooperatives whose statutes restricted membership could allocate units outside the Wohnungsamt system. Confessional cooperatives existed in Munich, the Katholisches Siedlungswerk (1,816 units, 1949–1959) largest among them, but they were small against the 76,000 units managed by the secular, city-owned GEWOFAG and GWG. Publicly subsidized cooperative units remained under Wohnungsamt allocation regardless of membership (Rädlinger, 2003). Separately, the Wohnraumbewirtschaftungsgesetz (§12(4))

permitted post-allocation apartment swaps with little friction. This biases against the contrast rather than toward it. If placed residents swapped into neighborhoods matching their type, the gradient should survive in administratively allocated precincts, not flatten. That it flattens makes the reading conservative.

Historical background

Munich's post-war housing crisis was among the most severe in Western Europe. Allied bombing destroyed 45% of the city's building substance and 27–31% of its housing stock (roughly 81,500 of 260,000 pre-war apartments); only 2.1% of buildings survived completely undamaged (Haus der Bayerischen Geschichte, 2024; Landeshauptstadt München, Sozialreferat, 2011). The population, down to about 479,000 in May 1945, rebounded to 769,000 by the end of 1946, 824,000 at the 1950 census, and past one million by December 1957. By 1950 Munich had regained its pre-war population but not its pre-war stock, a deficit of over 80,000 apartments (Landeshauptstadt München, Sozialreferat, 2011). The legal framework evolved through four statutes. Kontrollratsgesetz Nr. 18 (March 1946) established compulsory housing management across all occupation zones (Alliierter Kontrollrat, 1946). The Erstes Wohnungsbaugesetz (April 1950) shifted emphasis toward construction and carved out membership-restricted cooperatives. The Wohnraumbewirtschaftungsgesetz (March 1953) translated the framework into federal law. The Abbaugesetz (June 1960) began a stepwise dismantling; Munich retained full controls until December 31, 1968 (Egner, 2014). Munich's largest providers were city-owned and secular: GEWOFAG (founded 1928, ∼36,000 units) and GWG (founded 1918, ∼40,000 units), both allocating through the Wohnungsamt (Rädlinger, 2003). Refugees and expellees, 10.3% of the population by the 1950 census, were managed through a separate Münchner Flüchtlingsverwaltung and housed primarily in peripheral settlements on repurposed military infrastructure; no ethnic enclaves formed (Historisches Lexikon Bayerns, 2024; Bauer, 1982). Figure C.1 traces this build-out in the address books: the city and the charitable housing associations displace private land-

lords at the head of the ownership distribution between 1910 and 1933, and the stock the Wohnungsamt later allocated was largely in place before the war.

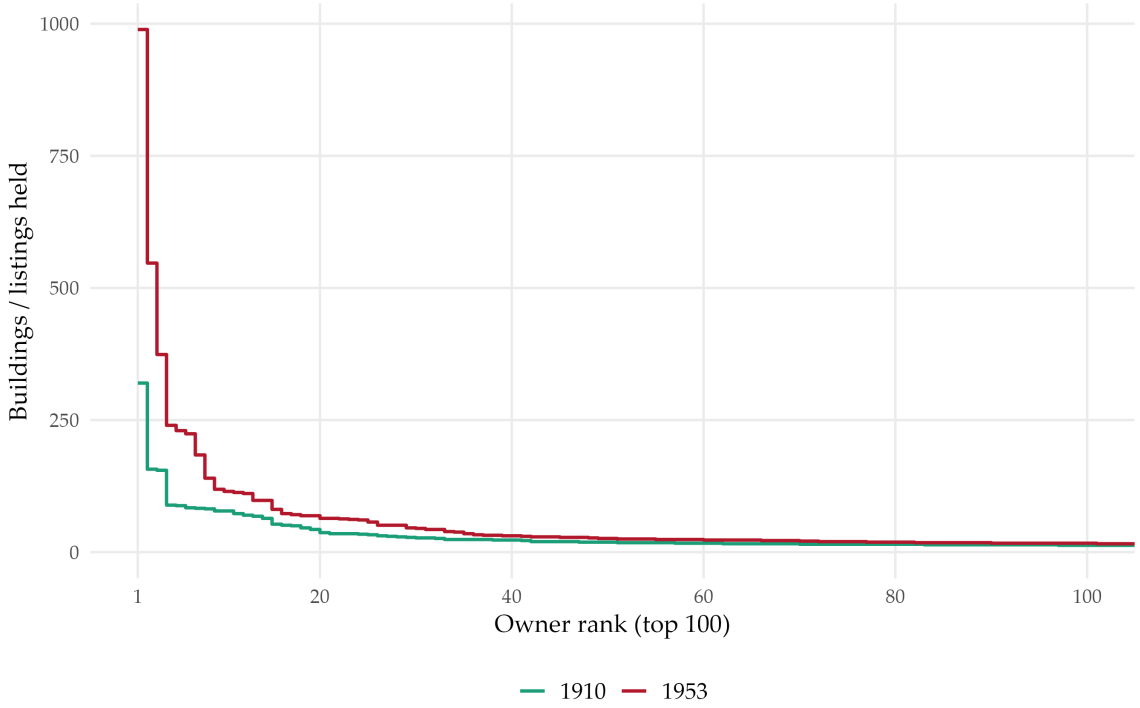

Figure C.1: The municipal and social-housing build-out, 1910–1953

Notes: Rank-size curves of building ownership in the 353 precincts (1933 boundaries) present in both address-book editions. Each step plots an owner's rank against the number of buildings it holds; only the top 100 owners are shown. The head of the distribution shifts from private landlords and breweries toward the city and the charitable housing associations: the Stadt München rises from 320 buildings in 1910 to 989 by 1953, and the Gemeinnützige Wohnungsfürsorge AG, absent from the 1910 head, becomes the second-largest owner with 547. The largest five owners are the Stadt München (320), the Spatenbräu (157), the Bavarian state fisc (155), and the private landlords Wagner (89) and Schmid (88) in 1910; and the Stadt München (989), the Gemeinnützige Wohnungsfürsorge AG (547), the Land Bayern (374), the Bundesrepublik (240), and the Löwenbräu (230) in 1953. The Wohnungsamt's post-war role was to *allocate* this largely municipal and cooperative stock, not to build it. Building-listing counts are raw for private owners and lightly de-duplicated for the largest institutions.

D Administrative Placement and Persistence: Regression Estimates

Figure 8 summarizes, how the growth in administrative ownership attenuates the reproduction of the 1910 type-map across composition and both votes. This appendix reports the underlying regressions with the continuous interaction. Table D.1 gives the composition outcome (post-war precinct-mean HISCAM on its 1910 value, both standardized, interacted with ΔAdmin, the 1910–1953 change in the administrative ownership share). Table D.2 gives the SPD and CSU vote outcomes, mirroring the specification of Table 1 and adding the administrative interaction and, in the final column of each party, the precinct's own 1912 vote. The interaction flattens the gradient toward zero in every outcome: negative where the 1910 gradient is positive (composition, CSU) and positive where it is negative (SPD). The interaction is significant in every outcome; the vote interactions remain significant once the precinct's own lagged 1912 vote is controlled, so the result holds outcome by outcome and is reinforced by the uniform direction across all three. ΔAdmin nets out the time-invariant pre-war ownership map, so the interaction loads on the post-war build-out. Table 6 carries the timing evidence, showing the newcomer-composition interaction is significant after the Wohnungsamt existed and null before it. The administrative share is an ownership object, so these are descriptive moderations rather than a second identification.

Table D.1: Growth in Administrative Ownership and the Reproduction of the Occupational Type-Map, Post-WWII

	(1)	(2)	(3)
$\text{Status}_{1910} \times \Delta$Admin share	-0.814***	-0.689***	-0.717***
	(0.171)	(0.228)	(0.223)
Average precinct status, 1910	0.765***	0.723***	0.733***
	(0.044)	(0.055)	(0.056)
ΔAdmin share (1953−1910)	-0.023	-0.072	-0.072
	(0.250)	(0.271)	(0.268)
War destruction (HH loss)			-0.171
			(0.171)
LASSO amenities		✓	✓
War destruction			✓
District FE	✓	✓	✓
Within R^2	0.530	0.606	0.607
N	333	317	317

Notes: OLS regressions of post-war (1953/1961 mean) precinct-mean HISCAM, standardized. *Average precinct status, 1910* is the standardized precinct-mean HISCAM in 1910; both outcome and regressor are standardized, so coefficients are in standard-deviation units. Δ*Admin share* is the 1910–1953 change in institutional (government, non-profit-housing, cooperative) ownership, the post-war build-out net of the pre-war ownership map, entered without winsorizing. The negative interaction shows the 1910 type-map is reproduced more weakly where administrative ownership grew most. Column 2 adds the 15 LASSO-selected pre-war amenities; column 3 adds war destruction (the 1933–1949 net loss of households). Within-district R^2. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table D.2: Growth in Administrative Ownership and the Persistent Spatial Vote, Post-WWII

	SPD				CSU			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\text{Status}_{1910} \times \Delta$Admin share	0.066***	0.059**	0.045*	0.050**	-0.049***	-0.044**	-0.047**	-0.051**
	(0.017)	(0.025)	(0.025)	(0.024)	(0.013)	(0.021)	(0.021)	(0.020)
Average precinct status, 1910	-0.047***	-0.039***	-0.025***	-0.027***	0.036***	0.028***	0.030***	0.031***
	(0.004)	(0.005)	(0.005)	(0.005)	(0.003)	(0.004)	(0.004)	(0.004)
ΔAdmin share (1953−1910)	0.030	0.028	0.017	0.017	0.000	0.004	0.024	0.023
	(0.022)	(0.025)	(0.023)	(0.022)	(0.020)	(0.023)	(0.020)	(0.019)
Lagged vote, 1912			0.192***	0.185***			0.423***	0.409***
			(0.033)	(0.034)			(0.069)	(0.067)
War destruction (HH loss)				0.025**				-0.026***
				(0.011)				(0.010)
LASSO amenities		✓	✓	✓		✓	✓	✓
Lagged 1912 vote			✓	✓			✓	✓
War destruction				✓				✓
District FE	✓	✓	✓	✓	✓	✓	✓	✓
Within R^2	0.414	0.478	0.549	0.555	0.338	0.412	0.498	0.508
N	333	317	317	317	333	317	317	317

Notes: OLS regressions of period-average post-war (1949–1961) precinct vote share. Average precinct status, 1910 is the standardized precinct-mean HISCAM in 1910. ΔAdmin share is the change in the precinct's institutional (government, non-profit-housing, cooperative) ownership share from 1910 to 1953, the post-war build-out of administratively allocated housing net of the pre-war ownership map; it is signed (ownership can fall) and so is entered without winsorizing. The interaction $\text{Status}_{1910} \times \Delta$ Admin tests whether the post-war growth in administrative ownership attenuates the pre-war status gradient: positive for the SPD (whose gradient is negative) and negative for the CSU (whose gradient is positive), flattening the gradient toward zero in both cases. Columns 2-4 and 6-8 add the 15 LASSO-selected pre-war amenities; columns 3-4 and 7-8 add the precinct's own 1912 vote; columns 4 and 8 add war destruction (the 1933-1949 net loss of households). Within-district R^2. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Change on Change Table D.3 restates the level regressions of Tables D.1 and D.2 as a long difference: the change in a precinct's vote share from 1912 to the post-war average on the change in its standardized composition from 1910 to 1953. Differencing removes every time-invariant precinct characteristic, durable amenities and fixed political culture alike, so only who moved in and out can drive the correlation between the two changes. The composition-change coefficient matches the level gradient for the SPD and is smaller for the CSU, and interacting it with ΔAdmin leaves both estimates unchanged, so the change link does not depend on the administrative build-out either. A level correlation could in principle reflect some durable precinct trait that happens to track both composition and the vote; the

change specification rules that out, since a fixed trait cannot explain a change.

Table D.3: The Vote Moves with the Map: Vote Changes and Composition Changes, 1912–1961

	ΔSPD			ΔCSU		
	(1)	(2)	(3)	(4)	(5)	(6)
ΔComposition (1910–1953)	-0.044***	-0.038***	-0.038***	0.008*	0.014***	0.015***
	(0.009)	(0.009)	(0.009)	(0.005)	(0.005)	(0.005)
War destruction (HH loss)			-0.004			0.011
			(0.024)			(0.013)
LASSO amenities		✓	✓		✓	✓
War destruction			✓			✓
District FE	✓	✓	✓	✓	✓	✓
Within R^2	0.103	0.248	0.248	0.012	0.189	0.191
N	332	317	317	332	317	317

Notes: OLS long-difference regressions. The outcome is the change in a precinct's vote share from 1912 to the post-war period average (1949–1961); the regressor is the change in its standardized occupational composition from 1910 to 1953, each year's precinct-mean HISCAM standardized within its own cross-section. Differencing removes every time-invariant precinct characteristic; district fixed effects absorb district-level trends. Columns 2–3 and 5–6 add the 15 LASSO-selected pre-war amenities (Table E.4); columns 3 and 6 add war destruction (the 1933–1949 net loss of households). Interacting the composition change with ΔAdmin leaves the estimates unchanged (SPD 0.015, $p = 0.85$; CSU -0.026, $p = 0.56$), so the change link does not vary with administrative growth. Within-district R^2. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

War Destruction A final concern is that war destruction, rather than the removal of choice, drives the attenuation: heavily bombed precincts were both more likely to receive administratively allocated housing and rebuilt into a different housing stock, which could reshape composition on its own. Controlling for destruction as a level addresses its effect on outcome levels but not on the status gradient itself. Table D.4 therefore estimates the $\text{Status}_{1910} \times \Delta\text{Admin}$ interaction within bins of war destruction, adding destruction-bin fixed effects together with the bin interacted with both 1910 status and ΔAdmin. The status gradient is then free to differ across destruction intensities, and the interaction is identified only within bins of comparable destruction. It survives in all three outcomes at both ten and twenty bins, so destruction alone does not account for the attenuation.

Table D.4: The Moderation Survives Within War-Destruction Bins

	Composition		SPD		CSU	
	(1) 10 bins	(2) 20 bins	(3) 10 bins	(4) 20 bins	(5) 10 bins	(6) 20 bins
$\text{Status}_{1910} \times \Delta\text{Admin}$ share	-0.953***	-1.264***	0.054**	0.078***	-0.040*	-0.059***
	(0.274)	(0.281)	(0.026)	(0.028)	(0.022)	(0.022)
ΔAdmin share (1953−1910)	1.585*	-0.097	0.010	0.175	-0.009	-0.116
	(0.828)	(0.996)	(0.051)	(0.109)	(0.040)	(0.094)
Destruction bins	10	20	10	20	10	20
Bin × Status_{1910}	✓	✓	✓	✓	✓	✓
Bin × ΔAdmin	✓	✓	✓	✓	✓	✓
LASSO amenities	✓	✓	✓	✓	✓	✓
District FE	✓	✓	✓	✓	✓	✓
Within R^2	0.641	0.698	0.527	0.594	0.488	0.553
N	317	317	317	317	317	317

Notes: Each column re-estimates the $\text{Status}_{1910} \times \Delta\text{Admin}$ moderation of Tables D.1 and D.2 within bins of war destruction, testing whether the attenuation merely reflects that heavily bombed precincts were rebuilt into a different housing stock. Precincts are sorted into 10 (columns 1, 3, 5) or 20 (columns 2, 4, 6) equal-count bins of war destruction (the 1933–1949 net loss of households). Each regression adds destruction-bin fixed effects, the bin interacted with 1910 status (so the status gradient is free to differ across destruction intensities), and the bin interacted with ΔAdmin, leaving $\text{Status}_{1910} \times \Delta\text{Admin}$ as the only term identified within destruction bins. The composition outcome (columns 1–2) is the standardized post-war precinct-mean HISCAM; the vote outcomes (columns 3–6) are post-war (1949–1961) party shares. All columns add the 15 LASSO-selected pre-war amenities and district fixed effects. The interaction keeps its sign and significance across columns. Destruction alone therefore does not account for the moderation. Within-district R^2. Robust (HC1) standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Standard Errors The subgroup gradients in Figure 8 use robust (HC1) standard errors, as in the main tables. Figure D.1 re-estimates the same six point estimates (three outcomes × two subgroups) under three alternative standard errors: clustered by district (Stadtbezirk), the wild cluster bootstrap by district (Rademacher weights, 9,999 replications), and Conley (1999) spatial HAC standard errors with a 1km Bartlett kernel on precinct centroids (the mean address-book coordinates within each precinct). All four standard-error variants use the 333 precincts in districts with at least two precincts, since the wild cluster bootstrap requires clusters with more than one observation. Point estimates are unchanged by construction; confidence intervals widen modestly under the wild bootstrap, the most conservative of the four, but every subgroup gradient remains significant and the no-growth/admin-grew gap remains non-overlapping in all three outcomes under all four standard-error choices.

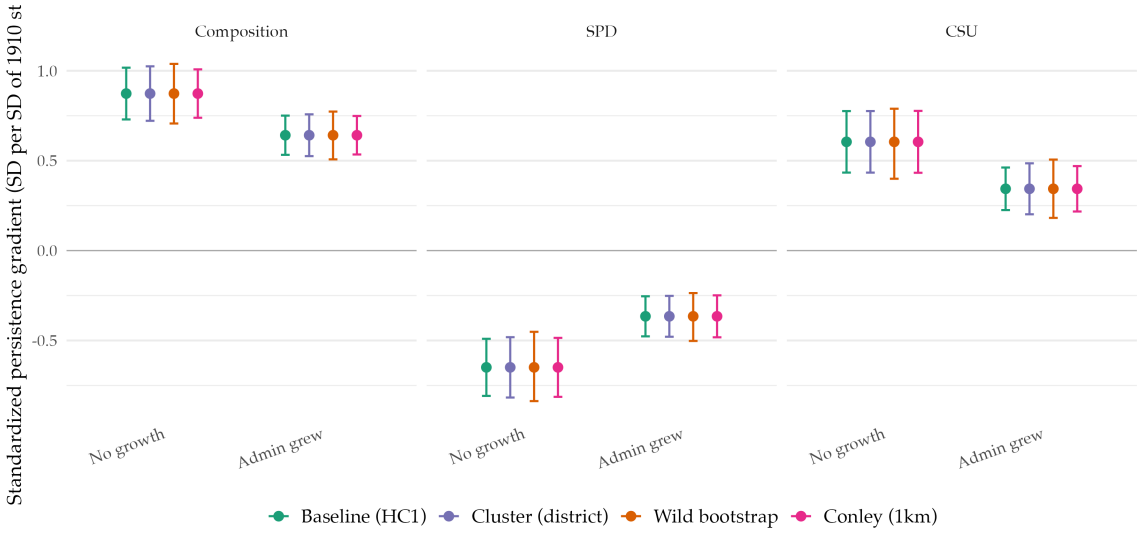

Figure D.1: The Administrative-Placement Gradient Under Alternative Standard Errors

Notes: The same six standardized subgroup gradients as Figure 8 (composition, SPD, and CSU, each split into *No growth* and *Admin grew* precincts by the sign of ΔAdmin), estimated under four standard-error methods: robust (HC1), as in Figure 8; clustered by district (30 clusters); a wild cluster bootstrap by district (Rademacher weights, 9,999 replications); and Conley (1999) spatial HAC with a 1km Bartlett kernel on precinct centroids. Point estimates are identical across methods by construction; only the 95% confidence intervals differ. All four methods agree that the gradient is significantly flatter in precincts where administrative ownership grew, in every outcome.

Heterogeneity by Sex of Household Head The Wohnungsamt's urgency list ranked applicants by severity of need, war widows and bombed-out families among the neediest, so if placement disrupts sorting, the disruption should be visible among female-headed households. Figure D.2 re-estimates the composition gradient of Figure 8 separately by the sex of the household head, with each sex's own 1910 precinct-mean HISCAM predicting its own post-war counterpart. Where administrative ownership did not grow, the two maps persist at similar strength, 0.81 for men and 0.69 for women. Where it grew, the male gradient falls to 0.58 and the female gradient halves to 0.33.

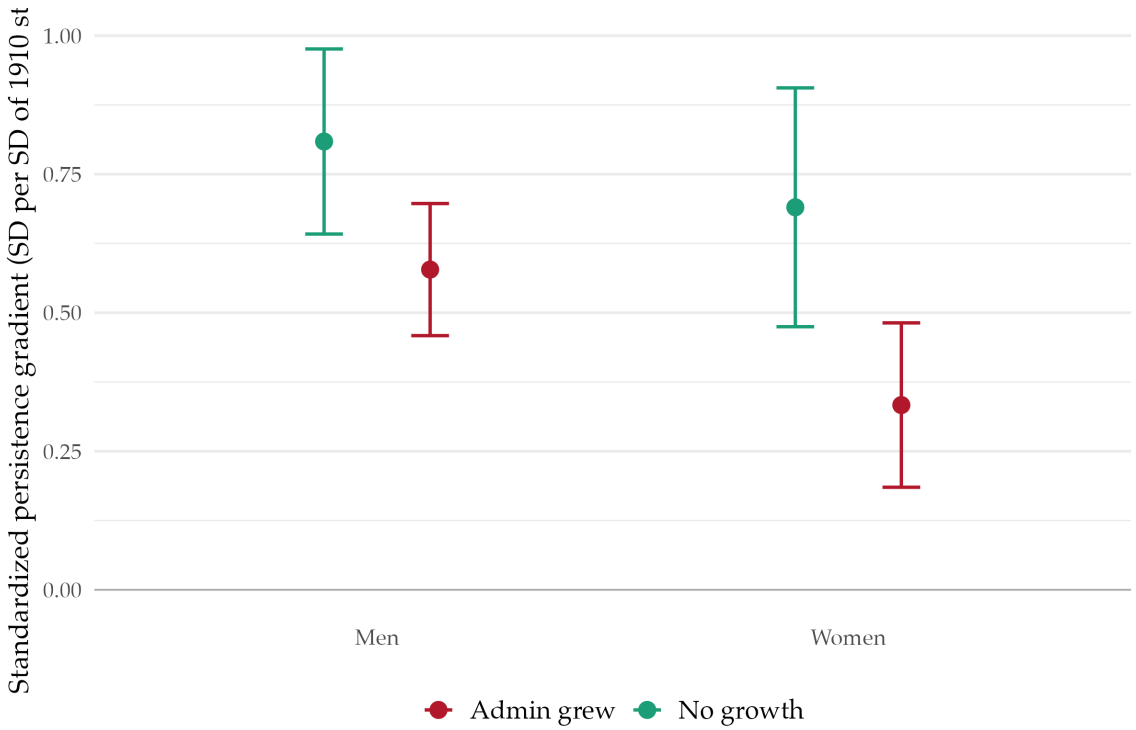

Figure D.2: Administrative Ownership Growth and the Composition Gradient, by Sex of Household Head

Notes: Standardized gradients from the 1910 sex-specific precinct-mean HISCAM to its post-war counterpart (1953/1961 mean), estimated as in Figure 8: jointly estimated subgroup slopes by the sign of ΔAdmin, with the 15 LASSO-selected pre-war amenities and district fixed effects; HC1 standard errors, 95% confidence intervals. Sex is inferred from first names. Precinct means require at least ten heads of the given sex with a HISCAM score Female household heads are a selected population, about half of them are widows. Estimation sample: $N = 331$ precincts (men), 308 (women).

Churches and Parishes as Institutions A church is both an institution and a durable feature of place, so heterogeneity of the gradient by church proximity alone cannot separate the two: a steeper gradient near churches could reflect parish influence on residents or movers sorting toward the parish. The placement moderation separates them, because placement closes the sorting margin. If parishes convert whoever arrives, the gradient that survives placement should be concentrated near churches, above all in the CSU columns, since the Catholic milieu survived the regime while the SPD's did not. Church proximity is the precinct-mean distance to the nearest church in the 1910 register, sign-flipped and standardized so that higher values mean closer. Table D.5 adds it to the vote specification of Table D.2: columns 1 and 4 report the plain gradient, columns 2 and 5 the placement moderation, and columns 3 and 6 the full triple expansion. Proximity alone does not steepen the gradient: without the admin terms, the $\text{Status}_{1910} \times \text{Proximity}$ interaction is 0.032 for the SPD and -0.031 for the CSU, both insignificant and an order of magnitude below the baseline gradient of 0.48. The triples instead carry the sign of the moderation itself, 0.096 for the SPD and -0.082 for the CSU. Where placement removed choice, more of the gradient disappears near churches, the opposite of what parish treatment predicts. The pattern is consistent with churches drawing matching movers, a pull that placement removes.

Table D.5: The Status Gradient, Administrative Ownership Growth, and Church Proximity

	SPD			CSU		
	(1)	(2)	(3)	(4)	(5)	(6)
$\text{Status}_{1910}\times\Delta\text{Admin}\times\text{Church prox. 1910}$			0.096**			-0.082
			(0.047)			(0.053)
Status_{1910}	-0.479***	-0.497***	-0.503***	0.477***	0.492***	0.487***
	(0.050)	(0.045)	(0.048)	(0.050)	(0.047)	(0.052)
$\text{Status}_{1910}\times\Delta\text{Admin}$		0.106***	0.130***		-0.100***	-0.131***
		(0.027)	(0.027)		(0.029)	(0.027)
$\text{Status}_{1910}\times\text{Church prox. 1910}$			0.060			-0.027
			(0.065)			(0.069)
$\Delta\text{Admin}\times\text{Church prox. 1910}$			-0.039			0.084
			(0.047)			(0.054)
ΔAdmin share (1953−1910)		0.048	0.020		0.004	0.035
		(0.036)	(0.032)		(0.042)	(0.037)
Church prox. 1910			0.002			0.069
			(0.068)			(0.077)
District FE	✓	✓	✓	✓	✓	✓
Within R^2	0.365	0.414	0.428	0.307	0.338	0.357
N	332	332	332	332	332	332

Notes: OLS regressions of the standardized period-average post-war (1949–1961) precinct vote share on the standardized 1910 precinct-mean HISCAM, on the 1933-boundary precinct map. ΔAdmin share is the 1910–1953 change in the administratively owned housing share (government, non-profit housing, cooperative), standardized. Church prox. 1910 is the precinct-mean distance from each address to the nearest church in the 1910 register (90 churches), sign-flipped and standardized so that higher values imply shorter distance; the register predates every placement studied. Columns 1 and 4 report the plain gradient, columns 2 and 5 add the placement moderation of Table D.2, and columns 3 and 6 the full triple expansion. The $\text{Status}_{1910}\times$Church prox. term conflates parish influence with sorting toward the parish and carries no institutional reading on its own. Only the triple measures whether the flattening under placement varies with church proximity. Catholic proximity (correlation 0.98 with the any-church measure) yields triples of 0.084 and −0.075; the 1933 register distance and the 1912 Catholic population share yield zeros. All columns carry district fixed effects; the amenity-controlled version of the moderation is in Table D.2. The sample is the 334 precincts with 1910 church distances; the fixed effects drop two precincts that are the only members of their districts. Within-district R^2. Robust (HC1) standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

E The Grain of Residential Sorting

This appendix provides robustness checks for the sorting result of Section 6.3 from several angles: whether it survives at the coarser district grain, whether it reflects genuine type-seeking rather than a mover's origin resembling the destination, and whether movers who change occupation sort toward their new type rather than their old one. It also documents the pre-war amenities the LASSO selects as controls throughout Sections 6 and 7.

Table E.1: Owner change and household head turnover 1919–1933

	(1)	(2)	(3)	(4)	(5)
Owner change	0.032***	0.030***	0.025***	0.024***	0.023***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.003)
Mean dep. var. (no change)	0.192	0.192	0.192	0.192	0.192
At-risk occupants (treated bldgs)	144,654	144,654	144,654	144,654	82,183
Implied induced movers	4,562	4,404	3,682	3,487	1,889
F (owner change)	247	241	173	160	72
First treatment only					✓
Precinct + Year	✓	✓	✓	✓	✓
Street FE		✓	✓	✓	✓
Building FE			✓	✓	✓
Precinct × Year				✓	✓
Within R^2	0.003	0.002	0.002	0.002	0.001
N	155,287	155,227	154,446	154,346	106,727

Notes: The dependent variable is household turnover at a building between adjacent address-book editions: the share of edition-$t-1$ tenants no longer present at t (cohorts of at least three tenants). "Owner change" indicates that the building's owner name-set changed between adjacent editions (Jaro-Winkler similarity below 0.90). Columns add fixed effects progressively: precinct and year (column 1), then street (column 2), then building (column 3), then precinct×year (columns 4–5). Column 5 keeps each building only up to and including its first owner-change wave, so the treatment effect is identified off the first treatment period alone ($t = 0$ is the one post period); never-changers keep all waves, indicated by the "First treatment only" row. The instrument is strong (F-statistic between 70 and 243) but its intensity is modest: the coefficient implies a 2.3–3.1 percentage point increase in turnover on a base rate of 0.194. An owner change significnatly raises the probability that a household leaves. Standard errors clustered by address. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The owner change is a strong instrument for displacement. Table E.1 shows that an owner change raises the probability of household turnover by 2.3–3.1 percentage points on a baseline turnover rate of 19.4% during the Weimar period. The estimate survives precinct×year,

street, address, and censored first-treatment-period-only specifications.

The sorting logit of Table 2 ranges over the precinct choice set, and the precinct is an administrative unit harmonized to the 1933 boundaries. A mover in 1929 neither saw these lines nor knew where they would later fall, so the precinct is not the unit anyone chose over. The city districts (Stadtbezirke) are the coarser, older divisions a contemporary resident did navigate. If sorting were a choice among the thirty districts, then estimating it over an average choice set of 317 precincts would impose a partition finer than the decision and read precision into noise. The resolution is that the precinct is not a choice set at all. It is a fine grid laid over a within-district surface of composition and amenities that a mover perceives at the level of the street and the block, and it is this within-district variation, not the choice of district, that carries the sorting.

Table E.2 re-runs the forced-versus-voluntary contrast of Table 4 at the district level (i.e. the coarser geographical unit), defining a mover as anyone who crosses a district boundary and the type as the district-mean HISCAM. The coefficient is smaller than at the precinct level, as coarsening to thirty districts averages away the within-district type-spread that the previous table shows carries most of the sorting. The same holds at the district grain: $\Delta\gamma$ runs from 0.033 to 0.063 across the three panels, positive in each. Thus, the finding that forced movers sort to type as strongly as voluntary movers is not an artifact of the fine partition.

Table E.2: Forced Movers Sort to Type at the District Level, Weimar

	Voluntary (control bldg.)	Pushed (owner-change bldg.)	$\Delta\gamma$
Panel A. Same-period S_n, no amenities			
Type-match γ	0.191***	0.225***	0.033
	(0.021)	(0.036)	
Panel B. Same-period S_n, + LASSO amenities			
Type-match γ	0.183***	0.225***	0.041
	(0.022)	(0.036)	
Panel C. Fixed 1910 S_n^{1910}, + LASSO amenities			
Type-match γ	0.182***	0.245***	0.063
	(0.023)	(0.040)	
N	44,777	3,303	

Notes: The γ-contrast of Table 4 re-estimated at the district level. A mover is anyone whose destination district (Stadtbezirk) differs from the origin. The choice set is districts rather than precincts. The logit is estimated separately on movers leaving owner-change buildings ("pushed") and control buildings ("voluntary"); γ is the homophily type-match $-|\tau_i - S_n|$ and $\Delta\gamma = \gamma_{\text{push}} - \gamma_{\text{vol}}$. Panels A and B use the same-period composition S_n, Panel C the fixed 1910 pre-war composition; Panels B and C add the district amenities a LASSO ($\lambda_{1\text{se}}$, ten-fold) selects when regressing 1910 district composition on the 32 physical amenities aggregated to the district. Panel D from Table 4 cannot be estimated at the district level: its amenity controls are themselves aggregated to the district, so once destination-district fixed effects enter, the district-level amenity vector is collinear with those fixed effects and drops out. The type-match regressor $-|\tau_i - S_n|$ is unaffected, since it varies across movers within a district through τ_i, but the specification as a whole is no longer identified. The coefficient is smaller than the precinct-level γ of Table 4 because aggregation to districts averages out the within-district type-spread. Forced movers again sort more strongly than voluntary movers in all three panels, as at the precinct level. Standard errors clustered by origin district. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

E.1 Origin Resemblance versus Type-Seeking

One concern may be that γ in Table 2 is actually origin-similarity rather than type-match: movers might seek out precincts that resemble the one they just left, not precincts that match their own type. Two design choices already close much of that gap (composition is fixed at its predetermined 1910 value, and the district fixed effects absorb anything durable about where a mover started that also describes their destination), but neither directly rules out simple homesickness for a familiar type of place. This section adds two checks that speak to that objection without relying on a joint regression of origin-similarity against type-match, which the errors-in-variables problem below would rig against the type-match term regardless of what movers actually do.

The origin precinct never chosen Table 2's choice set includes the mover's own origin precinct as an alternative, even though it is never the chosen destination by construction: a mover is defined as someone observed at a new address outside their former precinct. The origin precinct is also, trivially, the nearest possible alternative and carries a near-maximal type-match, so its presence in a never-chosen row could bias both γ and the distance coefficient toward zero. Table E.3 re-estimates every column of Table 2 with the origin precinct dropped from the choice set. Dropping a never-chosen alternative-row never drops a mover, so the estimation sample is unchanged. The type-match coefficient is essentially unaffected: γ moves by at most 0.63 of its published clustered standard error across all ten columns, most columns moving by under half an SE. The distance coefficient moves more, from about -0.93 to about -1.10 log-meters, because the excluded row was, for every mover, the shortest possible distance in the choice set and a row the model could never rationalize as chosen. Once it is removed, the model recovers a distance penalty that is stronger, not weaker, than the published one. Sorting on type is unaffected by the origin-precinct artifact; if anything, the corrected specification understates how costly distance is relative to a good type-match, so the marginal-rate-of-substitution figure in Section 6.3.2 is conservative.

Table E.3: Sorting Logit with the Origin Precinct Dropped from the Choice Set

	(1)	(2)	(3)	(4)	(5)
Panel A: Same-period S_j, origin dropped					
Type match $(-\lvert\tau_i - S_j\rvert)$	0.304***	0.327***			0.342***
	(0.012)	(0.011)			(0.010)
Type match $(-\lvert\tau_i - S_j\rvert) \times S_j$			0.266***	0.255***	
			(0.007)	(0.007)	
log distance	-1.091***	-1.112***	-1.092***	-1.112***	-1.105***
	(0.018)	(0.018)	(0.017)	(0.018)	(0.020)
N	17,249,303	17,249,303	17,249,303	17,249,303	17,249,303
Panel B: Fixed 1910 S_j^{1910}, origin dropped					
Type match $(-\lvert\tau_i - S_j\rvert)$	0.260***	0.276***			0.286***
	(0.013)	(0.012)			(0.011)
Type match $(-\lvert\tau_i - S_j\rvert) \times S_j$			0.221***	0.212***	
			(0.007)	(0.007)	
log distance	-1.093***	-1.114***	-1.093***	-1.112***	-1.107***
	(0.018)	(0.018)	(0.018)	(0.018)	(0.020)
Amenity controls		✓		✓	✓
Destination-district FE					✓
Mover FE	✓	✓	✓	✓	✓
N	16,858,466	16,858,466	16,858,466	16,858,466	16,858,466

Notes: Table 2 re-estimated with the mover's own origin precinct removed from the choice set; otherwise identical samples, formulas, and columns. Because the origin precinct is never the chosen alternative, dropping it never drops a mover, only the never-chosen row. N is the number of mover-by-destination observations remaining after the drop. Standard errors clustered by origin precinct. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Movers who change occupation move with their type If γ reflected homesickness for a familiar type of place rather than type-seeking, a change in a mover's own occupation between the origin and destination edition should carry no information about where they land: origin resemblance has no channel through which a mover's own occupational change could shift their destination. Type-seeking predicts the opposite: movers whose occupation improved should land in higher-composition precincts than movers whose occupation held steady, who in turn should land higher than movers whose occupation declined. Figure E.1 tests this directly. For Weimar movers linked under the loose matching rule (the only match set in which an occupational change can be observed, since the conservative rule blocks on exact occupational agreement by construction) and typed at both the origin and destination edition, it plots the change in destination-minus-origin composition, on the fixed 1910 map, against the change in the mover's own occupational score, separately for movers whose HISCAM rose by at least ten percent of its own pre-move value and movers whose HISCAM fell by the same margin. Measurement error in a machine-coded occupational string can only attenuate this ordering toward nothing; it cannot manufacture an ordering that is not there. The pattern in the figure is therefore conservative evidence for type-seeking rather than a result inflated by noise, and a null result here would not itself be evidence against sorting, since attenuation alone could produce it.

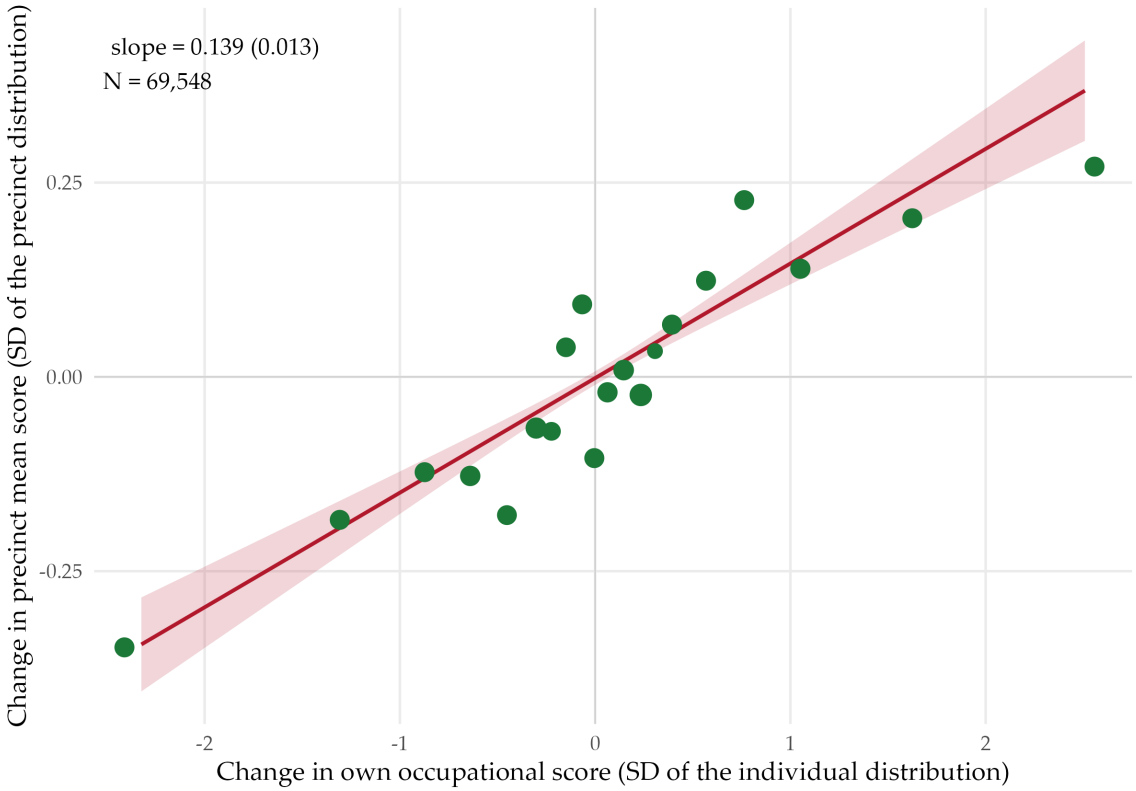

Figure E.1: Movers Who Change Occupation Move Toward Their New Type

Notes: Weimar movers linked under the loose matching rule (1919–1932 annual pairs), typed at both the origin and destination edition. The horizontal axis is the change in the mover's own occupational score (HISCAM) between editions, standardized by its cross-sectional SD among movers; the vertical axis is the change in destination-minus-origin precinct composition on the fixed 1910 map, standardized by its cross-precinct SD, the same object as S_j in Table 2. Points are binned means; the line is a fitted OLS slope with the origin precinct clustered standard error. This design conditions on the mover's own occupational change rather than pitting a type-match term against an origin-similarity term in one regression, which a joint specification cannot settle: the destination's composition is a smooth precinct aggregate while the occupational match is a noisy machine-coded string, so a joint regression would let the smooth aggregate win by errors-in-variables regardless of what movers actually do. The unconditional distribution of destination-minus-origin composition across all movers is not shown, since the stable composition map of Section 5 forces its mean toward zero by construction (inflows balance outflows at every level of S_p), and a centered distribution would be misread as evidence against sorting rather than as the mechanical consequence of a stable map that it is.

The amenity controls in the sorting logit (Table 2) and the conditioning set in the administrative-placement gradient (Figure 8) are not chosen by hand. A LASSO selects them, regressing the 1910 precinct composition on the thirty-two pre-war physical amenities, and the selected bundle is the set of amenities that predict where the high-status residents already lived. Table E.4 lists the fifteen the LASSO keeps, ordered by the size of their standardized shrinkage coefficients. The selection is a variable-selection step rather than an inferential regression, so it carries a cross-validated fit but no standard errors. Distance to a café, household density, and pub density carry the most weight, which describes the dense inner-city blocks where the white-collar residents concentrated; the bundle is a compact physical signature of pre-war status, not a curated list.

Table E.4: Amenities Selected by the Composition LASSO

Amenity	Std. LASSO coef.
Household density	−0.333
Dist. to café (m)	−0.327
Pub density	−0.220
Mean max. floors	+0.196
Dist. to pub (m)	+0.084
Dist. to association (m)	+0.077
Address fragmentation	−0.057
Schools (count)	+0.046
Churches (count)	+0.041
Rear-building share	−0.039
Dist. to school (m)	−0.035
Population	−0.035
Post offices (count)	+0.028
Bakery density	+0.026
Dist. to bath house (m)	+0.014
Selected / candidate	15 / 32
Precincts	319
CV R^2 (1se)	0.311

Notes: The 15 pre-war amenities a LASSO ($\lambda_{1\text{se}}$, ten-fold cross-validation) selects when regressing 1910 precinct composition (precinct-mean HISCAM) on the thirty-two physical amenities, ordered by |coefficient|. This is the amenity bundle used as controls in the sorting logit (Table 2) and as the conditioning set in the administrative-placement gradient (Figure 8). Both the outcome and the amenities are standardized before the LASSO, so the coefficients are standardized shrinkage coefficients, comparable across amenities; LASSO is variable selection, so the table reports the cross-validated fit but no standard errors. Census composition shares and distance-to-center are excluded from the candidate set by construction.

F Sorting Across the War

The sorting logit of Table 2 is estimated on Weimar movers, while the composition persistence documented in Section 5 runs to 1961. That table therefore leaves open whether movers still sorted to type after the war. We settle it by linking individuals across four pairs of directory editions under the conservative matching rule of Section 6.1: 1910–1919, 1919–1933, 1935–1953, and 1953–1961. We then re-estimate equation (2) separately on each set of movers. The third pair spans the bombing and the population replacement, and the fourth follows the decade after it.

Movers sort to type in every period we can link (Figure F.1). The type-match coefficient is 0.36 for 1910–1919 and 0.37 for 1919–1933, then 0.24 across the war and 0.24 for 1953–1961, positive and significant throughout. After the war it sits roughly a third below its pre-war level and does not recover by 1961. The two pre-war pairs span nine and fourteen years against eighteen and eight after the war, so the decline is not a matter of how long each window is open.

Precincts are harmonized to the 1933 boundaries throughout the paper, so the post-war decline could reflect a city that had outgrown that frame. A second line in Figure F.1 re-estimates each period over the actual precincts of the year, so new buildings, suburbs, and newly cut precincts enter the choice set. Munich's frame grows from 386 precincts in 1933 to 759 in 1961. Sorting is stronger on the actual precincts of the year in every period, since a given amount of segregation shows up as more composition variation across a finer choice set. The two lines fall in parallel, so the post-war decline holds on both frames.

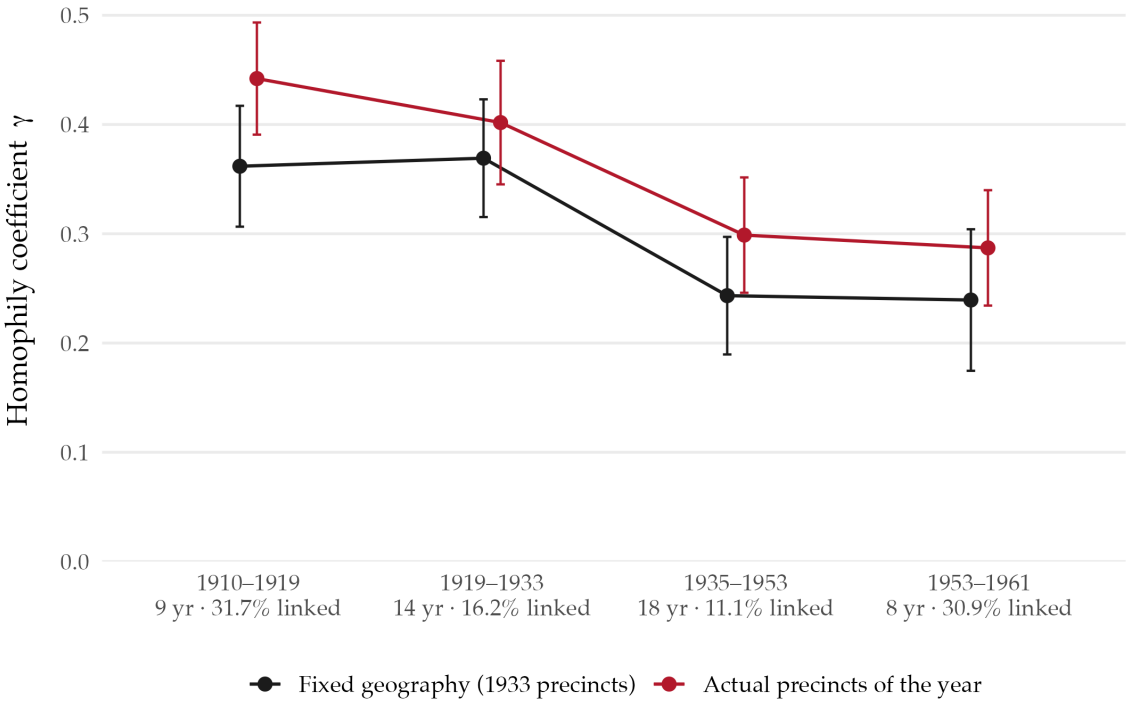

Figure F.1: Sorting to Type in Four Mover Samples, 1910–1961

Notes: Type-match coefficient γ from equation (2), estimated separately on individuals linked across each pair of directory editions, with mover and destination-district fixed effects and the distance control. Bars are 95% confidence intervals, clustered on the destination precinct. The fixed-geography line uses the 1933-harmonized precincts that the rest of the paper uses. The line for the actual precincts of the year uses that year's own precincts (1919, 1933, 1953, 1961), taken from its matched election file. Each such precinct is assigned to the original district most of its addresses fall in. Ties go to the lower district number, and a precinct with no address in any original district goes to the nearest district centroid. Amenity controls enter the fixed-geography estimates only, since the amenity data are built on the 1933 precincts; adding them leaves that line at 0.37, 0.38, 0.26, and 0.25. Link rates fall with the length of the window rather than with the era: 1953–1961 and 1910–1919 link at 30.9 and 31.7%. Their coefficients nonetheless differ by a third, so selection into being linkable does not generate the decline. Restricting to movers whose origin and destination both match a crosswalk coordinate exactly, which drops roughly 40% of the sample, leaves the profile unchanged.

G Summary Statistics

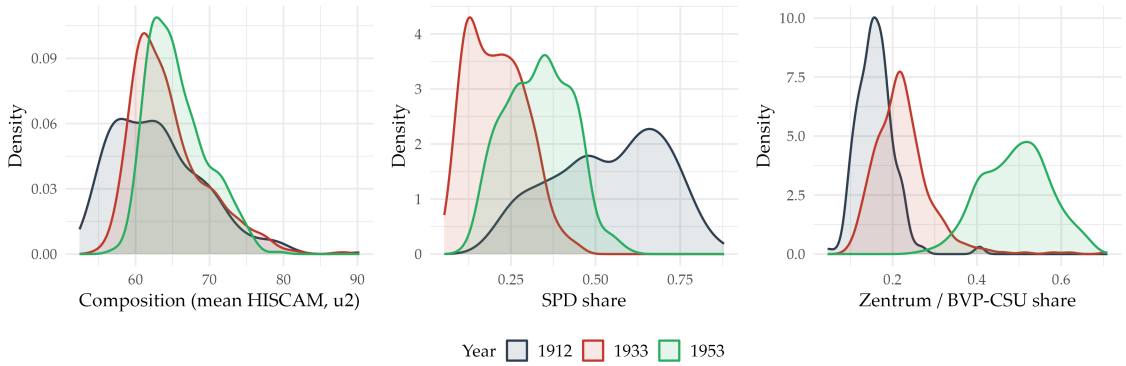

Figure G.1: Precinct Distributions of Composition and Vote Shares, 1912–1953

Notes: Kernel densities across precincts for three elections spaced roughly twenty years apart: 1912 (late Empire), 1933, and 1953 (early Federal Republic). Left: precinct-mean occupational status (HISCAM, universal u2 scale), read for the pre-war point from the 1910 directory matched to the 1912 election. Center: SPD vote share. Right: Zentrum/BVP-CSU vote share. The composition mean drifts modestly upward across the three dates, from 63.0 to 65.6 (about four-tenths of a 1910 standard deviation), and its spread narrows, a small movement set against far larger shifts in both vote distributions. This contrast between a relatively stable occupational structure and a changing vote is the descriptive counterpart to the paper's central claim. The vote label is held fixed while the underlying party changes: the right panel plots Zentrum/BVP in 1912 and 1933 and its post-war successor, the CSU, in 1953, so part of the rightward shift reflects the CSU absorbing a broader conservative coalition than the pre-war Zentrum rather than pure vote gain among a constant bloc. Precincts are harmonized to 1933 boundaries; sample sizes vary by measure and year (composition: 338, 380, 376; vote shares: 348, 380, 376 for 1912, 1933, 1953), reflecting missing returns and the residents-per-precinct threshold for a composition mean.

Table G.1: Occupational Composition Persists Across the War

	Composition, Weimar			Composition, Post-WWII		
	(1)	(2)	(3)	(4)	(5)	(6)
Composition, 1910 (s.d.)	0.827***	0.873***	0.739***	0.659***	0.708***	0.531***
	(0.033)	(0.031)	(0.048)	(0.041)	(0.042)	(0.060)
Within R^2	0.761	0.808	0.857	0.500	0.549	0.681
N	335	297	297	335	297	297
Covariate sample		✓	✓		✓	✓
Covariates			✓			✓
District FE	✓	✓	✓	✓	✓	✓

Notes: OLS regressions of standardized precinct-mean HISCAM (universal u2 scale) in the Weimar (1919–1932, columns 1–3) and post-WWII (1949–1961, columns 4–6) periods on the standardized precinct-mean HISCAM in 1910. Composition is built as in Table 1: occupation-level HISCAM averaged over a precinct's residents, pooling the period's address books. Both sides are z-scored, so the coefficient reports the within-district standard-deviation change in later composition per one standard deviation of 1910 composition. Within each period, columns 1 and 4 use the full sample; columns 2 and 5 restrict to the covariate-complete sample without adding covariates; columns 3 and 6 add the full pre-WWI amenity set on that same sample. The near-identical estimates in columns 1–2 and 4–5 show the persistence coefficient is not an artifact of the covariate sample; the attenuation in columns 3 and 6 reflects the covariates themselves. All columns include district (Stadtbezirk) fixed effects. Within-district R^2. Robust (HC1) standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table G.2: Summary Statistics — Precinct Covariates (I)

	N	Mean	P10	P50	P90
Densities & Counts					
Pub density	352	0.075	0.000	0.062	0.159
Store density	352	0.079	0.000	0.060	0.178
Restaurant density	352	0.017	0.000	0.000	0.051
Café density	352	0.005	0.000	0.000	0.021
Churches (count)	352	0.244	0.000	0.000	1.00
Schools (count)	352	0.844	0.000	0.000	2.00
Factories (count)	352	0.514	0.000	0.000	1.00
Housing					
Rear-building share	352	0.078	0.000	0.064	0.167
Mean max. floors	352	2.55	1.11	2.86	3.53
Address fragmentation	352	0.040	0.000	0.001	0.121
Household density	352	10	4.16	11	16
Religion					
Protestant share	317	0.143	0.072	0.134	0.205
Jewish share	317	0.018	0.002	0.011	0.048
Catholic share	349	0.824	0.742	0.832	0.902
Dist. to church (m)	352	381	171	320	631

Notes: N reports non-missing observations across 352 precincts harmonized to 1933 boundaries. *Densities & Counts:* Densities are the number of establishments (identified by occupation string in address books) divided by total unique establishment addresses in the precinct, averaged across the 1895, 1900, 1905, and 1910 address books. Church locations from the 1933 parish register; school locations from municipal records; factory counts from the 1912 factory survey. *Housing:* Rear-building share is the fraction of address-book entries located in a Rückgebäude (rear courtyard building). Mean max. floors is the average maximum floor number observed per address, a proxy for building height. Address fragmentation is the share of addresses with letter-subdivided house numbers (e.g., 41a, 41b), indicating residential subdivision of plots. Household density is the number of listed persons per unique address. All housing variables from address books, averaged across 1895–1910. *Religion:* From the 1910 Census, aggregated at the district (Stadtbezirk) level.

Table G.3: Summary Statistics — Precinct Covariates (II)

	N	Mean	P10	P50	P90
Distances					
Dist. to school (m)	349	255	80	170	504
Dist. to pub (m)	349	118	28	58	183
Dist. to store (m)	349	119	14	43	225
Dist. to restaurant (m)	349	258	60	145	467
Dist. to café (m)	349	541	91	339	1,239
Dist. to center (m)	349	2,347	875	2,201	3,732
Dist. to factory (m)	349	348	94	227	611
Dist. to tram (m)	349	395	53	180	869
Dist. to association (m)	347	655	163	450	1,292
Dist. to pharmacy (m)	349	477	138	267	958
Dist. to bath house (m)	349	641	223	489	1,174
Dist. to bank (m)	349	1,252	356	1,043	2,254
Dist. to post office (m)	349	513	170	395	948
Services					
Pharmacies (count)	352	0.159	0.000	0.000	1.00
Bath houses (count)	352	0.074	0.000	0.000	0.000
Banks (count)	352	0.040	0.000	0.000	0.000
Post offices (count)	352	0.128	0.000	0.000	1.00
Demographics					
Population	327	1,429	328	1,574	2,239
Male share	325	0.470	0.401	0.477	0.515
Military share	325	0.014	0.000	0.000	0.010
Young share (<15)	317	0.250	0.158	0.269	0.323
Working-age share (15–60)	317	0.599	0.545	0.598	0.673
Born in Munich share	317	0.397	0.340	0.405	0.486
Born in Bavaria share	317	0.478	0.443	0.481	0.534
Foreign-born share	317	0.050	0.030	0.038	0.073

Notes: N reports non-missing observations across 352 precincts harmonized to 1933 boundaries. *Distances:* Mean Haversine distance (in meters) from all geocoded addresses in a precinct to the nearest amenity of each type. Amenity locations from address books (pubs, stores, restaurants, cafés), the 1933 parish register (churches), municipal records (schools), the 1912 factory survey (factories), and the 1900 tramline shapefile (tram stops). Center distance measured to Marienplatz. *Services:* Counts of pharmacies, bath houses, banks, and post offices per precinct, geocoded from the 1914 municipal directory. *Demographics:* Population from the 1910 Census matched to 1933 precincts via address crosswalk (77% match rate). Shares (age, birthplace) from the 1910 Census at the district level.

Table G.4: Match Quality Flags by Step

Step	Outcome	N	Flagged	Share flagged (%)
A. Conservative step				
	Stayer	104	0	0.0
	Moved within precinct	4	0	0.0
	Moved across precinct	7	0	0.0
	Step total	115	0	0.0
B. Loose step				
	Stayer	120	0	0.0
	Moved within precinct	2	0	0.0
	Moved across precinct	13	3	23.1
	Step total	135	3	2.2
Overall		**250**	**3**	**1.2**

Notes: $N = 250$ matches drawn at random from the full matched file for manual quality inspection; the 1938 register year is excluded. The *conservative* step requires exact name match and exact HISCO major group agreement at level 2. The *loose* step relaxes the name criterion to a Jaro–Winkler distance ≤ 0.10 (similarity ≥ 0.90) and is applied to residuals not matched in the conservative pass. A match is *flagged* (`flag = 1`) when it passes the name-distance filter but shows indicators of a likely false positive. Specifically, across-precinct move combined with a large occupational shift across HISCO major groups. All three flagged matches occur in the loose step on cross-precinct movers (23.1% of that cell); stayers and within-precinct movers are unflagged under both steps.