Electoral Risk and Housing Supply*

Vincent Heddesheimer†

August 19, 2026

Click for latest version

Abstract

Why do local politicians restrict housing even when shortages are severe? New construction creates immediate, concentrated political costs, while many benefits arrive later, are diffuse, and are harder to attribute. I argue that electoral risk determines whether these costs deter politicians from approving new housing. I combine administrative records from Germany and the United States with an original survey of more than 5,200 local politicians. In Germany, a difference-in-differences design shows that a sharp increase in electoral risk reduces housing supply, especially multifamily construction. In US cities, municipal permitting rises under the same mayor once term limits remove reelection pressure. The survey shows that voter opposition, fiscal costs, and planning burdens deter projects, while many politicians who regard supply expansion as effective expect voters to prefer other policies. Housing shows how electoral accountability can deter long-term investments whose political costs precede their benefits.

*I thank Christian Baehr, Fiona Bare, Carles Boix, Rafaela Dancygier, Alex Gazmararian, Hanno Hilbig, Miguel Pereira, Florian Sichart, Anna Weissman, Andreas Wiedemann, and audiences at EPSS, EUI, Konstanz, LSE, and Princeton for helpful feedback. I also thank all interviewees for their time and my research assistants for their diligent work. Princeton IRB protocols: #16641 (interviews) and #18295 (politician survey). The research was made possible by funding from the Mamdouha S. Bobst Center for Peace and Justice and the Niehaus Center for Globalization and Governance.

†Princeton University, vincent.heddesheimer@princeton.edu

1 Introduction

Housing is increasingly unaffordable across advanced democracies. In many metropolitan regions, rents and property prices have risen much faster than household incomes, locking younger and lower-income groups out of homeownership and straining renters' budgets. Housing shapes where people can live and work (Molloy, Nathanson and Paciorek, 2022), as well as access to neighborhoods and the distribution of wealth (Ioannides and Ngai, 2025). The affordability crisis also has electoral consequences, changing how renters and homeowners vote (Abou-Chadi, Cohen and Kurer, 2025; Ansell et al., 2022; Held and Patana, 2023; Larsen et al., 2019).

The housing crisis reflects a persistent imbalance between demand and supply. Population growth, urban agglomeration, and labor-market concentration have fueled demand, but new construction has not kept pace (Findeisen, 2022; Gyourko, Mayer and Sinai, 2013). Politics is a central explanation for this weak supply response. As Gyourko and Molloy (2015, p. 1290) conclude, "Regulation appears to be the single most important influence on the supply of homes." Across most advanced economies, local governments retain substantial authority over the zoning, planning, and permitting decisions that restrict or enable construction (Brouwer and Trounstine, 2024; OECD, 2017). These local decisions are politically contested. Research on housing politics has concentrated on the citizen side of that contestation. Homeowners, neighborhood groups, and civic associations mobilize to delay or block development to protect property values (Fischel, 2001), preserve neighborhood character (Trounstine, 2018), or resist projects whose price effects they doubt (Elmendorf, Nall and Oklobdzija, 2025; Hankinson, 2018). That work explains why residents oppose construction and how they mobilize. Whether this opposition restricts supply, however, depends on how politicians respond.

I center local politicians in the politics of housing supply. Resident opposition shapes development, but it does not translate mechanically into permitting decisions. I argue that electoral risk is a distinct determinant of housing supply. It captures the career consequence an officeholder expects from losing political support over a policy choice: the greater the expected consequence, the less willing she is to accept the political costs of construction.

The argument builds on theories of long-term policy investment and electoral pandering (Canes-Wrone, Herron and Shotts, 2001; Jacobs, 2011). New construction creates visible and attributable costs through neighborhood opposition, disruption, and fiscal demands. Its benefits arrive later, are harder to observe, and often accrue to people who are weakly represented in local politics. When the expected electoral cost outweighs a project's political return, a politician who values additional housing still refuses it. The expected electoral cost depends on both the support penalty associated with the project and the career consequence of losing that support. Holding the project and opposition constant, greater electoral risk makes rejection more likely.

I test the argument in German municipalities, where mayors and municipal councils control local land-use planning and shape the permitting process. I construct a panel covering all German municipalities from 1995 to 2023. It links restricted administrative registries of every building permit and completion to the German Election Database (GERDA; Heddesheimer et al., 2025). In a stacked difference-in-differences design, I follow the same re-elected mayor through consecutive terms, distinguishing a change in electoral vulnerability from a change in politician, party, or municipality.

I compare mayors whose first-round support falls by more than ten percentage points with returning mayors whose support remains stable. A sharp loss of support reduces total permits by 12.0 percent and multifamily permits by 16.8 percent. The point estimates follow the argument's contentiousness ordering: they are larger for multifamily than single-family permits and completions, for multi-story than low-rise buildings, and for projects by corporate developers than private households. Differences in local demand, political and fiscal conditions, developer influence, or administrative timing do not explain the decline.

Instead, an original survey of more than 5,200 local politicians reveals the political calculus behind the result. A conjoint experiment shows that voter opposition, fiscal burdens, planning requirements, and delay reduce support for projects. Respondents also distinguish effective from electorally attractive housing policy. Among politicians who regard supply expansion as effective, 42 percent do not rank it among the two policies they believe receive the strongest support from their own voters. For these politicians, electoral incentives conflict with their policy judgments.

The survey identifies the political costs that discourage construction. The argument predicts that these costs should lose force when reelection is no longer at stake. I test this prediction in US cities, where term limits end eligibility for another term in the same office. Following the same mayor across eligible and final terms, I find that permitting rises once reelection is no longer possible: total permits by 17 percent, multifamily permits by 16 percent, and single-family permits by 13 percent.

Throughout the paper, I draw on 28 semi-structured interviews with politicians, planning officials, interest-group representatives, and other participants in German housing policy. The interviews motivate and illustrate the assumptions behind the argument. Appendix A describes recruitment, the interview protocol, respondents, and the use of the material.

I make three contributions. First, I place elected officials at the center of local policy responsiveness. Existing accounts of housing politics explain who opposes development, who participates, and how partisan coalitions shape policy (de Benedictis-Kessner, Jones and Warshaw, 2025; Einstein, Glick and Palmer, 2019; Fischel, 2001). I ask when those pressures bind the officeholder who decides whether policy changes. Holding the officeholder fixed, I show that the same politician changes supply as her electoral vulnerability changes. The comparison separates career incentives from stable differences in ideology, competence, and governing style and shows that responsiveness to organized residents depends on the politician's vulnerability to electoral punishment (Broockman and Skovron, 2018; Müller and Strøm, 1999; Sheffer, Loewen and Lucas, 2024).

Second, I identify electoral risk as the officeholder-specific link between opposition and the supply decision. My argument separates the support penalty attached to a project from the career consequence of losing that support. Housing makes this distinction especially useful because officials can identify who is likely to object, which decision will trigger opposition, and whom residents will blame. The same costly project can therefore be approved at low electoral risk and rejected at high electoral risk. Theories of pandering and long-term investment take the political cost of the unpopular choice as given (Canes-Wrone, Herron and Shotts, 2001; Jacobs, 2011; Maskin and Tirole, 2004). I show how an expected loss of support becomes a binding constraint when the officeholder is electorally vulnerable.

Third, I provide policy-level evidence on a classic argument about accountability and development. Advocates of political insulation held that voter pressure shifts resources from investment toward current consumption (Przeworski and Limongi, 1993; Przeworski et al., 2000). Because that literature compared regimes and aggregate growth, it could not show whether electoral pressure changed local investment decisions. I observe the policy response within democratic government. In Germany, municipal permitting falls under the same mayor after electoral risk rises. In US cities, it rises under the same mayor once a term limit ends reelection pressure. The findings identify a condition under which accountability discourages investment: organized constituents can punish politicians before voters observe and reward the investment's benefits. Housing therefore shows how electoral pressure can deter investments that officeholders themselves value.

2 Electoral Risk and the Politics of Housing Supply

Housing supply poses an intertemporal political problem.[1] The costs of authorizing construction become salient when a project is proposed. Many benefits arrive only after planning, construction, and subsequent moves through the housing market. These benefits often accrue to different people and are harder to connect to the decision that made construction possible. An officeholder can therefore bear an immediate penalty yet lose office before the benefits arrive or voters assign her credit. I argue that electoral risk determines when this imbalance constrains a politician's choice.

2.1 Why housing's political costs arrive first

Additional housing can ease scarcity and widen access to productive labor markets. New units can moderate rents and house prices (Mense, 2025; Molloy, Nathanson and Paciorek, 2022). Municipalities may also gain residents and revenue. Yet the political return to permitting differs from the economic and social return because costs and benefits reach the electorate through different channels.

The costs become concrete as soon as a project enters local politics. Nearby residents may anticipate noise, traffic, infrastructure demands, lower property values, or changes to neighborhood character. Opposition can also reflect distrust of developers, aversion to prospective residents, or doubts that market-rate construction will improve affordability (Broockman, Elmendorf and Kalla, 2024; Elmendorf, Nall and Oklobdzija, 2025; Fischel, 2001; Hankinson, 2018). For the argument developed here, the motive can vary. What matters is the officeholder's expectation that approval will cost political support.

[1] Following Jacobs (2011), I use *intertemporal* to describe how a policy distributes costs and benefits over time. This differs from models of dynamically inconsistent preferences (Strotz, 1955) and from credible-commitment problems in which policymakers later have incentives to reverse an announced policy (Kydland and Prescott, 1977).

Housing approvals make opposition traceable to a political decision. Residents can identify a proposed building, the municipal action that allows it, and the mayor or council members responsible for that action. These links allow voters to assign responsibility and punish those officials (Arnold, 1990). Residents who object can mobilize before construction begins, recruit neighbors, and carry the grievance into the next election. Negative consequences also receive more weight than comparable gains: voters attend more closely to negative political information and assign more responsibility to government for deteriorating than improving conditions (Larsen, 2021; Soroka, 2014; Weaver, 1986). Project-level discretion thus gives opponents a short and visible chain from a proposed change to the officeholder who enabled it (Mildenberger and Sahn, 2025). The expected penalty is larger when a project affects more residents or gives opponents a focal point for organizing.

The benefits follow a longer chain. Planning and construction often take years. New units begin to affect affordability only after tenants or owners move in. Those moves start chains that free older homes and reduce demand elsewhere in the housing market (Bratu, Harjunen and Saarimaa, 2023; Mast, 2023). Residents observe the rent they pay, not the higher rent that would have prevailed without additional supply. Officeholders therefore receive little credit for price increases that new construction prevented (Gailmard and Patty, 2019). Delay weakens attribution further because the approving official may lose office before these effects become visible.

The delay also changes who can take part when a project is decided. Nearby opponents already live in the municipality and share a focal grievance. Beneficiaries include current renters, younger households, prospective workers, and households helped by later moves through the housing stock. Some are dispersed across the city. Others are absent from the current electorate. Small groups with a shared grievance can organize more readily than a larger set expecting diffuse gains (Olson, 1965). Evidence from local planning meetings shows the resulting imbalance: nearby opponents participate at high rates, participants differ from the wider electorate, and the balance of comments is associated with project outcomes

(Einstein, Glick and Palmer, 2019; Sahn, 2025).

A German mayor described the missing constituency directly. When a few residents object, council members want to help them, but "the people who don't have an apartment, they don't come. They don't yet know that they will live there" (Interview F4, November 11, 2025). Many future beneficiaries neither live in the municipality nor know which project will house them when approval is contested. A project can therefore be blocked before the households it would serve know where they will live or can organize in its favor (Fernandez and Rodrik, 1991).

Housing decisions therefore give politicians an unbalanced account of their electoral effects. Opposition comes from identifiable current residents and attaches to a specific project. Many gains accrue later to people who cannot organize before the decision or reward the officeholder who approved it.

The menu of housing policies can sharpen this imbalance further. A politician under pressure to address high rents can choose an instrument that delivers visible relief to current residents instead of one that adds supply slowly. Rent regulation gives current tenants an immediate and identifiable benefit even though it can misallocate existing homes and induce landlords to withdraw rental units (Diamond, McQuade and Qian, 2019; Glaeser and Luttmer, 2003). Supply expansion instead asks current residents to accept a project now in exchange for a market outcome that arrives later. A representative of a German landlord association contrasted the two timelines. Construction effects become measurable "only in the very long run," he said. A four-year electoral focus instead favors rent regulation, which is immediately visible and easier to sell to voters (Interview D3, December 4, 2024). A short political horizon can therefore shift policy toward immediate relief rather than new supply.

Together, these features create immediate political pressure to restrict supply. They do not explain why similar politicians facing similar pressure make different choices. Voter preferences determine the support penalty. The remaining question is when that penalty binds the decision-maker.

2.2 When electoral risk binds the supply decision

Housing supply has the temporal structure of what Jacobs (2011) calls policy investment: officeholders incur political costs now to pursue gains that emerge later. Jacobs argues that such investments require expected long-run returns, institutional capacity, and enough electoral safety to bear the current costs. Each condition matters for housing, but the electoral-safety condition takes a particularly concrete form. A proposed development identifies both the residents who may mobilize against it and the officials who can be blamed for approving it. Politicians may be unsure exactly how much support is at stake, but they can anticipate the source and direction of the penalty before they act. The central question is whether the officeholder can afford that expected loss.

The expected-returns condition combines a reason to build with a belief that construction will deliver it. That reason need not equal social welfare. A politician may value housing access or the growth and municipal revenue that construction can bring (Peterson, 1981). Doubts that a proposed housing development will advance these goals reduce its expected return. If the expected gains do not exceed the project's non-electoral burdens, lower electoral risk cannot make approval attractive.

Institutional capacity begins with authority over the decision. Land-use authority is principally local in many advanced democracies.[2] Authority within local government remains a scope condition because mayors, councils, and planning offices can share control. Administrative capacity also affects how difficult a project is to approve. Lower electoral risk cannot increase supply when the relevant officeholder cannot act.

[2]Across 32 countries, the OECD finds that land-use planning is predominantly a local task. German municipalities prepare local plans, UK local authorities decide planning applications, and US states generally delegate land-use control to local governments. Higher levels still set legal frameworks and may retain authority over major projects (OECD, 2017).

Once a politician wants a project and has the authority to advance it, electoral risk can bind the decision. Consider an officeholder who expects approval to cost political support. If she expects to survive that loss, she may proceed. If the same loss now threatens her career, she may refuse the same proposal even though neither the project nor the opposition has changed. The expected support penalty belongs to the project and the political response it provokes. *Electoral risk* describes the officeholder's perceived career consequence of losing that support. Keeping the two separate is central to the argument.

Electoral risk is an ex ante assessment of what a loss of support would cost the politician's career, not the eventual election result or another name for electoral safety. It combines *electoral sensitivity*, how much a support loss changes the chance of retaining a valued political career, with the *career stake*, the value of that career relative to no longer holding it. Vote margins, current popularity, time until the next election, and term-limit status can shape these quantities, but none measures risk alone. A binding term limit removes the value of another term in the same office, though concern for one's party, reputation, or higher office may preserve part of the stake. A city councillor who was not standing again captured this distinction when he described himself as having *Narrenfreiheit*, or "free rein," because he was not running (Interview F2, December 4, 2024).

A compact decision rule records the choice. Conditional on authority over the project, the politician's net value of approving project m at date t is defined as

$$\Delta_{imt} = B_{imt} - C_{imt} - r_{it} p_{imt} > 0. \tag{1}$$

The term B_{imt} is the value, at the time of the decision, of the policy gains the politician expects approval to produce. It excludes electoral credit, which instead offsets the support loss in p_{imt}. The term C_{imt} collects burdens that do not arise from losing political support, including planning difficulty, infrastructure commitments, administrative work, and the personal cost of managing conflict. Limited administrative capacity raises these burdens. Before electoral considerations enter, the politician must expect the gains from approval to exceed these burdens, or $B_{imt} > C_{imt}$.

The term $p_{imt} \geq 0$ is the support the politician expects to lose from approving project m, net of any support she expects to gain. This penalty varies across projects. Prior research finds stronger opposition to apartments and taller buildings than to single-family or low-rise housing (Larsen and Nyholt, 2024; Trounstine, 2023). Reactions also depend on who builds. Voters may oppose corporate developers and institutional investors (Broockman, Elmendorf and Kalla, 2024; Dancygier and Wiedemann, 2025), while public housing can provoke concerns about municipal costs and prospective residents (Anderson et al., 2025; Dancygier, 2010; Hilbig and Wiedemann, 2026).

The term $r_{it} \geq 0$ translates a unit of lost support into a perceived career consequence. Writing $r_{it} = s_{it}v_{it}$ separates electoral sensitivity, s_{it}, from the career stake, v_{it}. The project determines how much support is at stake. The politician's electoral position and career stake determine what that loss would cost her.

An increase in r_{it} can also lead politicians to shrink or delay a project rather than reject it. At the aggregate level, the proposal set must not change enough to offset this threshold effect.

Figure 1 makes the comparison concrete for a project the politician would otherwise want to approve. The project, the expected opposition, and every non-electoral burden remain fixed. Only the politician's electoral risk changes. At low risk, the anticipated loss of support is tolerable and the project is approved. As risk rises, the career cost attached to the same loss of support rises with it. At the threshold $r^* = (B - C)/p$, that cost exhausts the project's net value. Beyond the threshold, the same proposal is rejected. A project with a larger support penalty would reach this point sooner, but the figure deliberately holds that penalty constant.

The argument shares an information structure with theories of electoral pandering, in

Figure 1: Perceived electoral risk and the approval decision

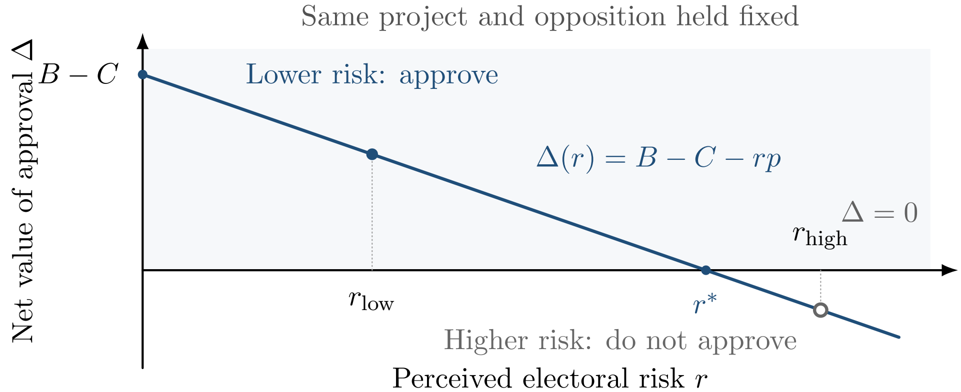

Notes: The figure plots $\Delta = B - C - rp$, the net value of approval, against perceived electoral risk for one project with $B > C$ and $p > 0$. The terms B, C, and the expected support penalty p are held fixed. At r_{low}, net value is positive and the politician approves. At r_{high}, the same project yields negative net value and is not approved. The decision changes at $r^* = (B - C)/p$. Only perceived electoral risk varies; opposition and all project attributes remain fixed.

which reelection incentives induce an official to follow public demand rather than her own policy judgment when voters cannot evaluate consequences before the election (Canes-Wrone, Herron and Shotts, 2001; Maskin and Tirole, 2004). Housing fits that structure: opposition is visible before benefits can be assessed. My mechanism does not require the private signal those models assume. Politicians need not know better than voters. They act on beliefs about the response they will face. Such beliefs predict support for long-term investments even when simple election proximity does not (Sheffer, Loewen and Lucas, 2024).

Among politicians considering comparable projects and able to act, electoral risk explains why the same support penalty binds some officeholders more than others. Voter preferences and organized opposition set the expected support penalty, p. Party and ideology can shape both B and p, while administrative capacity and project difficulty affect C. Holding these factors fixed, greater electoral risk raises the career cost of approval and makes rejection more likely.

2.3 Observable implications and scope conditions

The argument produces three observable implications and two conditions that determine where they should apply.

Implication 1: electoral risk. A politician should approve fewer projects when losing political support becomes more dangerous for her career. This prediction holds the proposal, opposition, and her reasons for building fixed. In the decision rule, an increase in r lowers the net value of approval whenever $p > 0$. The same costly project can therefore be approved at low risk and rejected at high risk.

Implication 2: the support penalty. Electoral risk should matter more for projects with a larger expected support penalty, p_{imt}. A given increase in r_{it} makes a larger penalty harder to bear.

Implication 3: the career stake. Electoral risk should weaken whenever continued office no longer depends on voter support. A binding term limit is one such case: by removing another term in the same office, it lowers v_{it} and makes costly projects more attractive. Concern for party reputation or higher office can leave part of the career stake intact, so the effect should be clearest among politicians whose electoral careers actually end.

Necessary condition: a reason to build. Electoral risk matters only when the politician expects a project's gains to exceed its non-electoral burdens. Lower risk cannot make a politician value approval where she otherwise prefers rejection. In the decision rule, this condition is $B > C$.

Scope condition: timing and authority. The argument applies most strongly when approval creates an immediate and attributable support penalty while its public benefits arrive later or remain hard to credit. It weakens when opposition is unlikely to cost support or when the officeholder lacks control over the decision. Housing need not be unique for this logic to apply.

3 Electoral Risk and Housing Supply in Germany

This section tests whether greater electoral risk reduces housing supply. I ask whether a sharp loss of electoral support changes subsequent permitting. I also examine whether the response is concentrated in more contentious forms of construction and assess competing explanations for the decline.

3.1 Municipal authority over housing supply

I test the electoral-risk argument in Germany, where municipalities exercise substantial authority over housing supply. Within the federal framework of the *Baugesetzbuch*, municipalities decide where housing may be built and at what density. Councils adopt preparatory and binding land-use plans, while directly elected mayors lead the municipal administration, help set the development agenda, and oversee the preparation and implementation of those plans.[3] Building permits translate these planning and administrative choices into authorized housing. They are the outcome closest to municipal action, whereas completions also depend on subsequent financing, construction costs, and implementation delays.

Mayors occupy a particularly important position in this process. They head the administration that prepares proposals, coordinates implementation, and helps determine which development questions receive sustained political attention. Councils formally adopt land-use plans, but mayors direct much of the work that brings those plans forward and puts them

[3] The allocation of administrative tasks varies across states and municipality types. The common institutional core is municipal planning authority under Sections 1–2 of the *Baugesetzbuch*. Sections 30–35 govern whether proposed construction is permissible under local plans and in areas without a binding plan.

into effect. Mayors are also directly elected executives whom residents can hold responsible for local development. Turnout averages 66 percent across recorded council elections and 59 percent across recorded mayoral contests, while 72 percent of mayoral races with observed candidate counts feature more than one candidate.[4] Administrative leadership and direct electoral accountability therefore make mayors the central actors for testing whether electoral pressure shapes housing supply.

Germany also provides leverage for distinguishing the electoral channel from institutions central to leading US accounts of housing politics. A majority of German households rent, household wealth depends less on owner-occupied property, and German homeowners do not consistently treat rising house prices as good economic news (Reisenbichler and Koenig, 2025). Public participation is channeled through defined stages of plan-making rather than the repeated discretionary project hearings emphasized in US research (Einstein, Glick and Palmer, 2019). These differences reduce the centrality of homeowner wealth protection and hearing-room participation, but they do not remove conflict over development. Residents can contest new construction through formal consultation, litigation, and local referenda, while elected officials remain responsible for the local decisions that determine supply. Germany therefore combines strong municipal authority with an institutional setting in which the electoral response of officeholders is especially visible.

3.2 Data

Testing the argument requires observing both politicians' electoral risk and the housing decisions made as their electoral position changes. I therefore combine local election returns and longitudinal mayor records with administrative data on permits and completions.

[4]Turnout and contestation figures are the author's calculations from GERDA. The contestation rate uses the 48,550 nonsuperseded mayoral races for which the first-round candidate count is observed.

Housing supply: I construct an annual panel of German municipalities from the restricted Building Permits Statistics (*Baugenehmigungsstatistik*) and Building Completions Statistics (*Baufertigstellungsstatistik*) for 1995–2023. The underlying registries cover the universe of permitted and completed construction. I harmonize the approximately 625,000 municipality-year records in these extracts across municipal boundary changes. The housing analysis contains 318,365 municipality-years across 10,981 municipalities and 9.52 million permitted dwelling units: 4.44 million in single-family buildings and 5.08 million in multifamily buildings. The permit registry records issued permits, not applications or rejected projects. The outcome therefore measures authorized dwellings in the municipality, not the mayor's project-level approvals or the share of proposed projects that receive approval. Because applications are unavailable, the analysis cannot distinguish fewer submissions from a more restrictive municipal process. The registry identifies building height, developer type, living space, and whether permitted housing was completed. These fields allow the analysis to compare projects with different expected political costs.

Local elections: I link the housing supply data to local elections through GERDA, a nationwide database of German local election returns that I built with coauthors (Heddesheimer et al., 2025). GERDA records more than 73,000 municipal council elections, including turnout and party vote and seat shares, and roughly 49,000 mayoral contests with first-round candidate returns. The council results provide measures of the winning margin between the two largest parties, council partisanship, fragmentation, and split government. From the mayoral election data, I built a mayor panel linking candidate-level election records and official term and re-election counters to the same individual across contests, and then expanding each officeholding spell into annual municipality-mayor records. The resulting longitudinal panel links 115,974 municipality-years to 13,137 individual mayors, including 6,633 observed in a second or later term. Appendix B reports data coverage, variable construction, and the mayor-linkage procedure.

Other municipal data: For covariate adjustment, moderator analyses, and tests of alternative explanations, I add official municipal population, area, migration, and land-use statistics from the Federal and State Statistical Offices. Municipal employment records from the Federal Employment Agency measure the shares working in real estate and construction. I also use rent and purchase-price indices constructed from real-estate listings (Ahlfeldt, Heblich and Seidel, 2023), income and purchasing-power series from empirica regio and RWI, and the BBSR classification of municipalities within major-city regions. Official 2011 and 2022 census records provide homeownership and vacancy rates. I measure local fiscal capacity with *Wegweiser Kommune* data on tax revenue and debt, and organized opposition with a municipality-year panel of housing-related *Bürgerbegehren* and referenda from the *Datenbank Bürgerbegehren* and Mehr Demokratie e.V. (Institut für Demokratie- und Partizipationsforschung, 2025).

3.3 Empirical strategy

Identifying the causal effect of electoral risk on housing supply is difficult because electoral risk is not randomly assigned. A mayor who receives less support may serve in a municipality with weaker housing demand, stronger local opposition, or less administrative capacity. A cross-municipality association would then combine electoral risk with these underlying differences. The relevant counterfactual is how permitting would have changed in the same municipality under the same mayor had electoral support remained stable. I approximate this counterfactual by measuring changes in the incumbent's first-round support across successful re-elections and comparing the corresponding changes in housing supply.

The empirical design follows re-elected mayors through consecutive terms. Its unit is a two-term episode for one mayor in one municipality. I define an episode as treated when the mayor's first-round vote share falls by more than ten percentage points from the election that begins the baseline term to the re-election that begins the next term.[5] I define an episode as stable when the change is smaller than five points in either direction. Episodes that satisfy neither definition do not enter the main comparison. More formally, let e index a two-term episode, with 0 denoting the baseline election and 1 the focal re-election. Let V_{e0} and V_{e1} be the mayor's first-round vote shares in those elections, measured as proportions. Treatment is defined as

$$\Delta V_e \equiv V_{e1} - V_{e0}, \qquad D_e = \begin{cases} 1 & \text{if } \Delta V_e < -0.10, \\ 0 & \text{if } |\Delta V_e| < 0.05. \end{cases} \tag{2}$$

I impose four conditions for an election sequence to enter the design. First, the same person must win both elections and remain mayor of that municipality between them. Second, the elections must be three to nine years apart.[6] The wider window admits early elections and scheduling irregularities while excluding long interruptions that would not represent adjacent terms in office. Third, the party label must remain unchanged so that the measured loss of support does not also capture a party switch or a new electoral alliance. Fourth, the mayor must face the same number of first-round candidates in both elections, which prevents a change in the size of the candidate field from mechanically shifting the incumbent's vote share.

The treatment captures a sharp deterioration in a returning mayor's demonstrated electoral support. Electoral position combines the cushion a mayor currently holds with the swings her support has shown in the past (Kayser and Lindstädt, 2015). A sharp loss moves

[5] I use first-round results because they record support before runoff coordination and candidate withdrawal can reshape the race (Arnold, 2018; Granzier, Pons and Tricaud, 2023). I use the incumbent's own vote share rather than the winner–runner-up margin because 818 of the 1,471 episodes have only one candidate in both elections, leaving the runner-up margin undefined. Own share also measures personal support rather than only the nearest challenger's proximity.

[6] Ninety-eight percent of eligible episodes are five to seven years apart.

both terms. Surveyed German mayors assess their own safety by the share they hold.[7] A large loss narrows her cushion and makes another support loss more consequential. These mayors lose substantial support but still win reelection and remain in office. The treatment therefore captures an increase in electoral risk through its electoral-sensitivity component, s. The design asks whether housing policy changes after this deterioration in the same officeholder's position. These conditions yield 262 treated episodes and 1,209 stable-support episodes across 1,183 municipalities. Appendix Figure C.1 displays every episode in calendar time and reports the contribution of each state.

I estimate the effect by pooling the 262 treated and 1,209 stable-support episodes in a stacked difference-in-differences design. Each episode enters once as a separate two-term panel. The model compares the average pre-to-post change in treated episodes with the corresponding change across the full stable-support group. All stable-support episodes jointly estimate one comparison path, and every treated episode is evaluated against that pooled path. No episode is duplicated in forming the stack.

For each episode e, I create an episode-specific panel containing every observed year in the two adjacent terms. I center the panel on the focal re-election year E_e and append all panels while retaining the episode identifier (Baker, Larcker and Wang, 2022). This structure keeps distinct re-election episodes separate when a municipality contributes more than one. Let $r_{et} = t - E_e$ denote relative time and $\text{Post}_{et} = \mathbf{1}\{t > E_e\}$. I exclude the election year because it combines months governed under each electoral condition, so the post-election period begins at $r_{et} = 1$.

[7] In the politician survey described in Section 4, self-assessed electoral safety among the 298 mayors who can be linked to their municipality's most recent mayoral election rises by a third of a standard deviation for every ten percentage points of own first-round share. Conditional on where a mayor currently stands, how far she fell to reach that position adds nothing measurable, so a loss operates through the position it produces. The treated–stable difference in post-election support corresponds to roughly half a standard deviation of self-assessed safety. Appendix C.2 reports these estimates and what the survey cannot show.

For housing outcome h, I estimate the following Poisson model on raw permit counts:

$$\begin{aligned}\mathbb{E}\big[Y_{et}^{h} \mid \mathbf{X}\big] = \exp\big\{&\alpha_e + \lambda_{s(e)t} + \rho_h \text{Post}_{et} + \tau_h\big(D_e \times \text{Post}_{et}\big) \\ &+ \gamma_h\big(Z_e \times \text{Post}_{et}\big) + \delta_h\big(V_{e0} \times \text{Post}_{et}\big) + \log(\text{Population}_{et})\big\}.\end{aligned} \tag{3}$$

The episode fixed effect α_e absorbs characteristics of the mayor and municipality that remain constant across the two terms. The term $\lambda_{s(e)t}$ denotes state-by-calendar-year fixed effects, which absorb shocks shared by municipalities in the same state and year. The coefficient ρ_h captures the mean post-election change among stable-support episodes. The difference-in-differences coefficient τ_h captures the additional post-election change among treated episodes.

Following Chen and Roth (2024), I estimate permit counts with Poisson quasi-maximum likelihood. Log current population enters only as an exposure offset, with its coefficient constrained to one, so τ_h describes a proportional change in the permit rate. I interact baseline vote share (V_{e0}) with the post-election period because treatment is defined by a change from that starting point. The variable Z_e is log population density measured in the last observed pre-election year.[8] Its interaction allows dense and less dense municipalities to experience different post-election shifts in permitting. The quantity $100[\exp(\tau_h) - 1]$ is the estimated percentage change in outcome h after a sharp loss of support. Standard errors are clustered by municipality, which also accounts for dependence when a municipality contributes more than one episode.

A causal interpretation of τ_h requires conditional parallel trends. Absent the loss of electoral support, permitting in treated episodes must have followed the same path as permitting in stable-support episodes after accounting for episode fixed effects, state-by-calendar-year fixed effects, and post-election shifts associated with baseline support and pre-election density. Two concerns remain. A municipality-specific shock could lower the incumbent's support and independently reduce subsequent housing demand or permitting. Alternatively,

[8]Measuring density before the focal re-election avoids conditioning on density changes that permitting decisions may themselves cause.

unusually high permitting during the baseline term could reduce support and then return to its previous level.

To assess pre-election differences and trace the response, I estimate a dynamic extension of Equation 3. The model compares treated and stable-support episodes in each relative year. It bins endpoints at five years, uses the year before re-election as the reference period, and excludes the election year. The lead coefficients show whether permitting was already diverging before re-election, although failing to reject zero cannot establish the counterfactual trend (Roth et al., 2023). Appendix Section C.4 reports the full specification. Because these relative-time comparisons cannot show whether permitting rose when the baseline term began, Appendix Figure C.10 compares that term with earlier municipal history.

3.4 Electoral risk reduces housing supply

Table 1 reports the pooled post-election difference-in-differences estimate for each permit outcome. The estimates compare the proportional pre-to-post change in treated episodes with the corresponding change in stable-support episodes. A sharp loss of support reduces total permits by 12.0 percent and multifamily permits by 16.8 percent. The single-family point estimate is smaller, 7.3 percent, and less precise. This ordering matches the prediction that electoral risk most strongly constrains projects likely to provoke visible opposition.[9]

Figure 2 traces the annual estimates to show when permitting changes and whether treated and stable-support episodes were already on different paths before re-election. Appendix Figure C.3 first reports the corresponding unadjusted group means. The adjusted pre-election estimates oscillate around zero. Joint tests cannot reject equal paths for total (p = .426), single-family (p = .761), or multifamily permits (p = .486). After re-election, the multifamily estimate is negative in every year and the decline persists across the new term. The annual estimates are less precise than the pooled result because only 206 to 261

[9]The direct multifamily–single-family difference is imprecise (Appendix Table C.3).

Table 1: Electoral risk reduces housing permits

	Total permits	Single-family permits	Multifamily permits
Vote-share fall × post-election	-0.128^{***}	-0.076^{*}	-0.184^{**}
	(0.049)	(0.043)	(0.086)
Effect (percent)	-12.0%	-7.3%	-16.8%
95% CI (percent)	[-20.1, -3.2]	[-14.8, 0.9]	[-29.7, -1.4]
Joint pre-trend p-value	0.426	0.761	0.486
Episode-year observations	14,170	14,170	14,170
Treated episodes	262	262	262
Stable-support episodes	1,209	1,209	1,209
Municipalities	1,183	1,183	1,183

Notes: Estimates of τ_h from Equation 3. Treatment follows Equation 2. Models include episode and state-by-calendar-year fixed effects, the post-election interactions in Equation 3, and the current-population offset. Standard errors are clustered by municipality. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

treated episodes contribute at each post-election horizon and municipal permit counts vary from year to year. These dynamics reveal no differential pre-election pattern and support the conditional-parallel-trends interpretation.

The argument predicts a larger response for projects that politicians expect to cost more support. Figure 3 reports estimates for completions, building form, and developer type. Total completions fall by 13 percent. The point estimates are larger for multifamily than single-family completions and for multi-story than low-rise houses, although the differences are imprecise.

Developer identity yields a similar pattern. Permits fall more for units built by corporate developers, and especially by public bodies and non-profit organizations, than for units built by private households. The estimates are imprecise, so the ordering remains suggestive. The ordering is nevertheless consistent with the project-level sources of opposition summarized in Section 2: opposition to developers and institutional investors, budget concerns surrounding public housing, and expectations about prospective residents.

Figure 2: Permitting falls after the incumbent loses electoral support

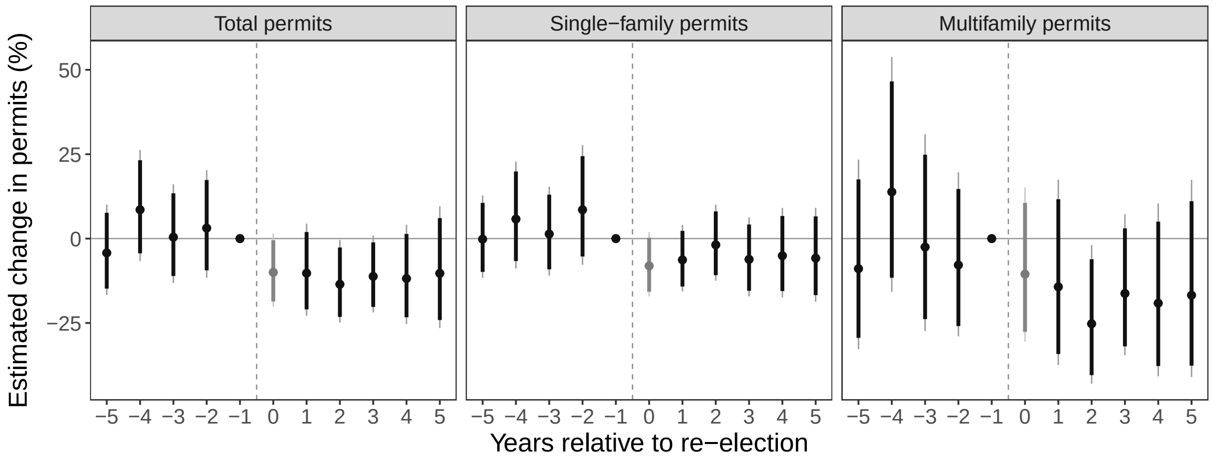

Notes: Estimates of τ_{kh} from the dynamic specification in Appendix Equation 5. The year before re-election $(t-1)$ is the reference period, and end points combine all earlier or later observations. The gray election-year estimate is excluded from the estimating equation because it combines months before and after re-election. Thin bars show 95 percent confidence intervals and thick bars show 90 percent confidence intervals. Standard errors are clustered by municipality. Confidence intervals are asymmetric because the log-scale estimates are transformed to percentage changes.

3.5 Additional results

I conduct additional analyses to examine whether housing supply responds to the direction and size of the electoral change, to other measures of electoral safety, and to conditions before re-election.

First, I reverse the treatment to identify re-elected mayors whose vote share rises by more than ten percentage points. Relative to stable-support episodes, the point estimates are positive for total, single-family, and multifamily permits, but none is statistically distinguishable from zero (Appendix Figure C.4). These estimates are consistent with supply rising when electoral risk recedes, but they remain suggestive rather than evidence of a symmetric response. I next split vote-share losses into ranges. The largest declines occur after losses between ten and fifteen points. The single-family estimate is positive for losses between five and ten points, while larger-loss bands are negative but increasingly uncertain. I also vary the minimum loss that defines treatment from five to twenty points. Total and

Figure 3: Electoral risk and theory-predicted housing supply outcomes

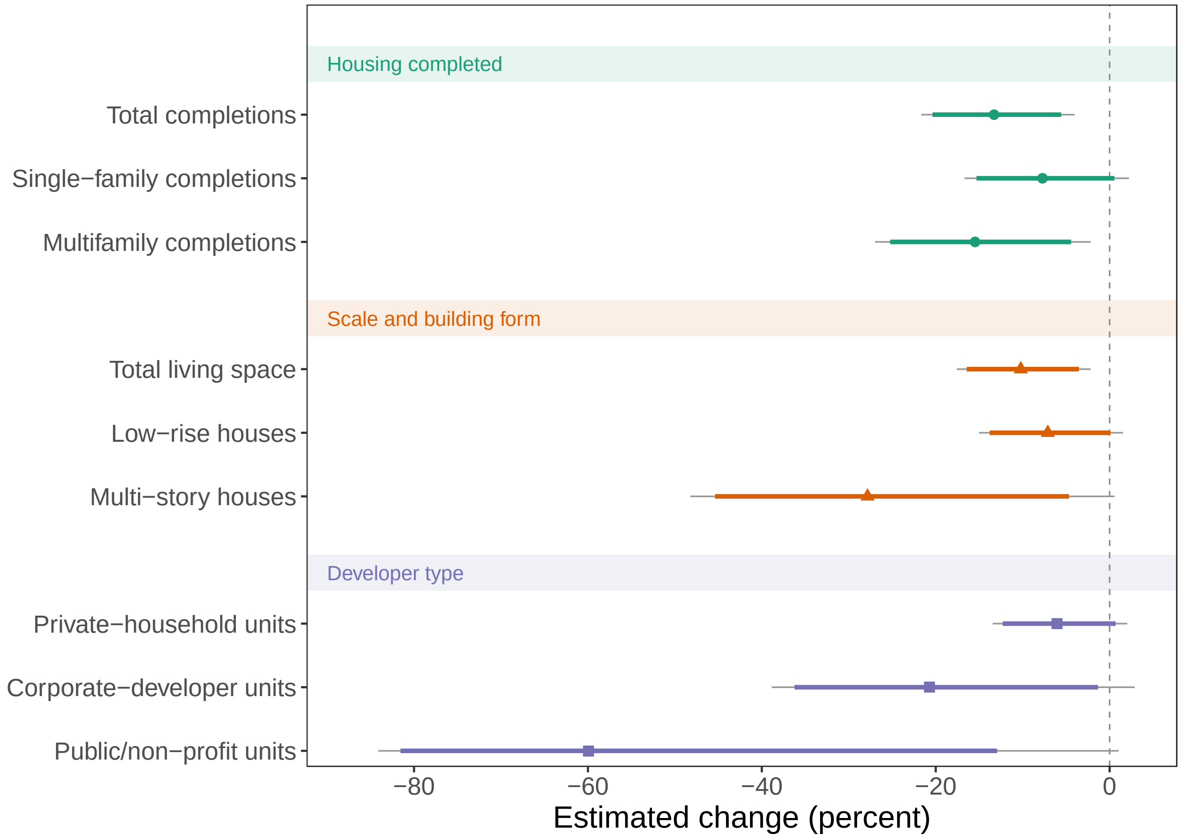

Notes: Each point shows $100 \times [\exp(\widehat{\tau}_h) - 1]$ from a separate estimate of Equation 3. Treatment follows Equation 2. Thin bars show 95 percent confidence intervals and thick bars show 90 percent confidence intervals. Standard errors are clustered by municipality.

multifamily estimates remain negative throughout, although only the ten-point definition yields estimates for both outcomes that are significant at the five-percent level (Appendix Figures C.5 and C.6).

Second, I estimate continuous fixed-effects models that impose fewer eligibility conditions and use much broader municipality panels. The mayoral panel contains about 111,000 municipality-years in 4,538 municipalities across 13 states. A ten-percentage-point loss in the mayor's first-round vote share predicts 2.1 percent fewer total permits, 1.2 percent fewer single-family permits, and 3.4 percent fewer multifamily permits (Appendix Table C.4).

I also measure the leading council party's safety using its vote-share margin over the second-largest named party. This measure captures the electoral position of the body that formally adopts local land-use plans. The council panel contains about 206,000 municipality-years in 7,388 municipalities across Germany. A ten-point narrower council margin predicts 4.9 percent fewer total permits, 2.4 percent fewer single-family permits, and 7.1 percent fewer multifamily permits.

In both broader panels, declining electoral safety predicts fewer permits, and the multifamily estimate is larger than the single-family estimate. These continuous specifications extend the evidence to both centers of local housing authority, but they do not reproduce the discrete comparison in the stacked design. Their smaller magnitudes are not directly comparable with the stacked estimates because they average associations across all observed electoral movements, while the stacked design isolates sharp losses relative to stable support in a restricted sample of re-elected mayors.

Finally, I examine whether the mayoral effects vary with conditions measured before re-election. The estimates show no gradient across quintiles of the mayor's starting vote share. A broader set of political, demographic, housing-market, and fiscal moderators also produces few systematic differences. Multifamily permitting falls more where pre-election tax revenue is low, which suggests that limited fiscal capacity may amplify the response to electoral risk (Appendix Section C.8). I otherwise find no systematic evidence that the result varies across political settings or with rent levels, rent trends, vacancy rates, or migration profiles.

3.6 Robustness checks

I assess how the result changes with baseline-support overlap, inference, episode construction, and an alternative difference-in-differences estimator.

The stacked sample is geographically concentrated. Bavaria contributes 73 percent of treated episodes and 78 percent of stable-support episodes. The less restrictive continuous models above cover many more municipalities across Germany. Both continuous measures produce the same relationship between electoral safety and permitting.

Treated episodes begin from a lower level of support than stable-support episodes: median baseline vote share is .81 and .90, respectively. Restricting the comparison group to the treated support range leaves the estimates nearly unchanged. A flexible spline adjustment attenuates the total-permit estimate to -9.5 percent and the multifamily estimate to -10.9 percent, with both confidence intervals including zero. ATT reweighting yields estimates between the flexible and baseline results. County clustering, two-way clustering by municipality and state-year, omission of the population offset, and retention of only the earliest qualifying episode preserve the direction and approximate magnitude (Appendix Tables C.5, C.6, and C.7). HonestDiD bounds for the average annual estimates become wide when post-election deviations may be half as large as the largest pre-election movement, so the dynamic estimates cannot rule out modest departures from parallel trends (Appendix Table C.8).

Finally, I replace the Poisson model with the doubly robust difference-in-differences estimator of Callaway and Sant'Anna (2021), applied to permits per 1,000 residents. Estimates are negative for all three outcomes across every aggregation window. Using every available year, they imply declines of 5.5 percent in total permits and 8.3 percent in multifamily permits. None of the fifteen estimates reaches the five-percent level, and joint tests reveal no differential pre-election pattern. This estimator corroborates the direction of the result while producing smaller and less precise magnitudes (Appendix Table C.9).

3.7 Assessing alternative explanations

The theory attributes the permit decline to mayors becoming less willing to incur immediate political costs after their electoral cushion narrows. I assess six rival accounts using characteristics measured before the focal re-election. Each adjusted estimate is paired with an unadjusted estimate from the same complete-case sample, so movement cannot be attributed to changing coverage (Appendix Tables C.10 and C.11). Appendix Table C.12 repeats the adjustments in the broader continuous panels.

Local development preferences, party control, and fiscal capacity could affect both the mayor's support and later permitting. I measure development preferences with pre-election homeownership and prior housing-related citizen initiatives (Einstein, Glick and Palmer, 2019; Institut für Demokratie- und Partizipationsforschung, 2025). I then adjust for mayoral party, the council left share, tax revenue, and debt (Bremer, Di Carlo and Wansleben, 2023; Chou and Dancygier, 2021; de Benedictis-Kessner, Jones and Warshaw, 2025). The opposition and party adjustments leave the estimates virtually unchanged. The fiscal adjustments move them by no more than 0.3 percentage points. Concurrent council elections are balanced across the groups, and the permit response does not vary systematically with concurrency. I therefore interpret the result as the response of local political leadership under the same mayor (Appendix C.12).

Adverse local conditions could cost the incumbent votes and reduce housing demand without changing her behavior (de Benedictis-Kessner and Warshaw, 2020; Hopkins and Pettingill, 2018). Adjusting for pre-election rents, five-year rent growth, and net migration attenuates the estimates by 2.5 to 3.3 percentage points on the same 616 episodes, but all remain negative. County-by-calendar-year fixed effects absorb shocks shared by nearby municipalities, and the dynamic estimates show no earlier divergence. The evidence does not fit a broad demand contraction that begins before re-election or extends across the surrounding county.

Developer capture offers a different account of the association between electoral safety and permitting. If safer mayors are more willing to accommodate organized development interests, the permit decline should be largest where real-estate and construction employment were initially strongest (Hankinson, Magazinnik and Weissman, 2024; Solé-Ollé and Viladecans-Marsal, 2012). The interaction estimates do not support this alternative. Evidence from US cities points the same way: developers connected to a mayor do sell more housing, but a decomposition attributes more than ninety percent of the citywide permitting effect to the mayor's own policy stance rather than to favors for donors (Yu, 2026). Routine administrative cycles and reversion also do not explain the result. Both groups contain re-elected incumbents, so a common new-term cycle is absorbed by the average post-election change. Treated episodes show no excess permitting increase into the baseline term, and neither the level nor the trend of permitting during that term predicts the subsequent loss of support (Appendix Figure C.10 and Table C.13).

The permit decline remains after accounting for development preferences, partisanship, fiscal conditions, local demand, industry strength, and administrative timing. Taken together, these tests tie deteriorating electoral support to reduced permitting and weigh against the main alternative explanations.

4 Why Housing Supply Is Politically Costly

Administrative records reveal how permitting changes with electoral risk, but not why politicians value additional housing or which project features make approval costly. I examine these premises with an original survey of German mayors and municipal councillors. Because the survey does not vary electoral risk, it tests these premises rather than whether electoral risk changes project support.

4.1 Survey and experimental design

I fielded an online survey among German mayors and municipal councillors in February and March 2026. I recruited respondents through two channels: 40,055 individualized invitations to elected officials in a stratified sample of municipalities and forwarding requests sent to 8,657 municipal administrations. The individual channel produced 4,390 substantive responses, an 11.0 percent response rate. Municipal links produced another 848 responses, for

5,238 in total. Of these respondents, 4,213 reached 100 percent progress. The primary analyses use current elected municipal officeholders who answered the relevant item, so sample sizes vary across modules. Respondents in the individual frame came from municipalities with somewhat narrower council margins, higher density and housing costs, and lower homeownership than nonrespondents. Women, SPD politicians, and Green politicians were also more likely to respond (Appendix Table D.3). The survey therefore describes participating officeholders rather than a population-weighted cross-section of all German local politicians. Appendix D reports recruitment, sample composition, selection, and the survey design.

The survey includes a conjoint experiment in which each respondent makes five paired choices between hypothetical housing projects. Four independently generated tasks vary project scale, completion time, price restrictions, access rules, planning pathway, fiscal effects, and expected reactions among affected residents and the wider municipality. Three feasibility rules exclude unrealistic combinations (Sichart and Wagner, 2025): projects adding at least five percent to the housing stock cannot finish within two years, ten-percent expansions require a zoning change, and local-resident priority appears only in new areas completed after at least three years. The fifth task repeats the first pair with the positions reversed to assess response consistency (Clayton et al., 2026).

I estimate the registered stacked profile-level OLS specification using the four independently generated tasks. Fiscal effects, price restrictions, and expected reactions remain independently assigned, so their coefficients identify average marginal component effects over the randomized profile distribution (Hainmueller, Hopkins and Yamamoto, 2014). For the attributes linked by feasibility rules, Appendix E.1 also reports nonparametric estimates within their exact common-support strata. These estimates lie within 0.7 percentage points of the registered coefficients and leave every sign and inferential conclusion unchanged. The conjoint experiment and moderator analyses were registered in a pre-analysis plan[10]. The salience, beliefs, and policy-instrument modules are descriptive. Appendix D reports the at-

[10] https://osf.io/nd7wx/overview

Figure 4: Opposition, fiscal costs, and delay reduce project support

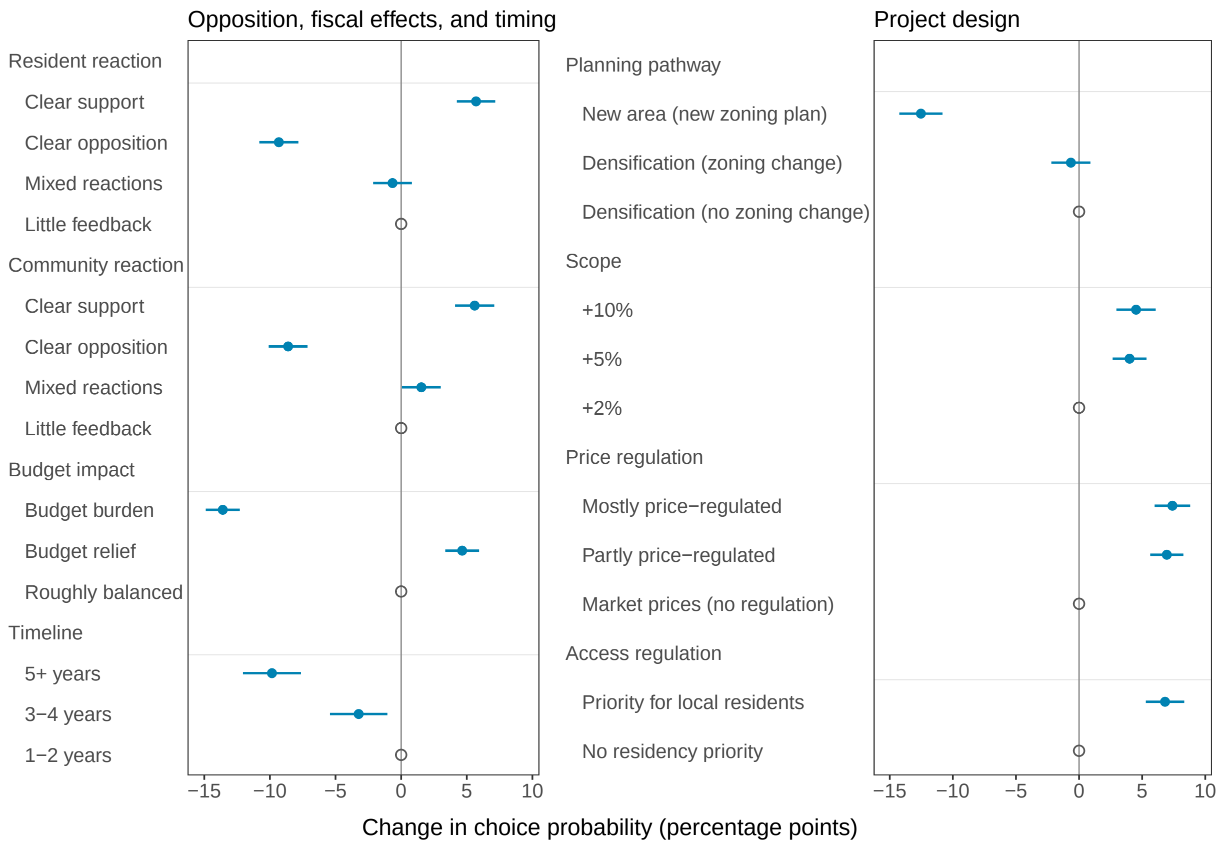

Notes: Pre-registered coefficients from four independently generated paired choices among current elected municipal officeholders. The outcome is whether each profile is selected. Open circles mark reference categories. Bars are 95 percent confidence intervals from stacked profile-level OLS estimates with standard errors clustered by respondent. Appendix Figure E.1 reports restriction-adjusted AMCEs for attributes linked by feasibility rules. $N = 4{,}251$ officeholders and 33,362 profiles with observed choices.

tributes and restrictions, and Appendix E reports the registered analyses and sample checks.

4.2 Immediate costs and delays reduce project support

The conjoint turns the theory's political tradeoff into a choice between concrete housing projects. If immediate costs and delayed benefits shape support, politicians should be less likely to select projects that provoke opposition, impose a fiscal burden, require a more demanding planning process, or take longer to complete. Figure 4 shows how each randomized conjoint attribute changes the probability that a project is chosen.

The largest deterrent is a net fiscal burden, which lowers the probability that a project is chosen by 13.6 percentage points relative to a fiscally balanced project. A new residential area requiring a new zoning plan reduces support by 12.5 points relative to densification that requires no zoning change. Completion after five or more years reduces support by 9.8 points relative to completion within two years. Clear opposition among directly affected residents and in the wider municipality lowers support by 9.3 and 8.6 points. Price restrictions increase support. Within non-fast projects in new residential areas, the only settings where access rules were randomized, priority for local residents also increases support. Larger projects receive slightly more support. This scale result does not support the pre-registered expectation that politicians prefer smaller projects.

Project timing matters for politicians' choices. A longer completion period postpones the point at which the municipality gains new housing. Creating a new residential area also makes the project visible and adds an approval stage before construction can begin. Across the randomized attributes, politicians are less likely to support projects whose political and administrative costs arise before the housing is complete. The penalties for delay, planning demands, and opposition persist when fiscal burden cannot distinguish the two profiles, while the fiscal and opposition penalties persist when a long delay cannot distinguish them (Appendix Figure E.2). The result is not driven by one categorical objection.[11]

Officeholders select the same underlying project in 79.3 percent of repeated choices. A response-noise correction increases the coefficient magnitudes without changing their ordering. The estimates are also stable when I use all respondents, exclude the fastest five percent of officeholders, adjust for respondent characteristics, or report marginal means (Appendix Figures E.7–E.9 and E.8).[12]

[11] Appendix Section E.3 reports the pre-registered moderator analysis by perceived electoral safety. The low-safety subgroup is too small for a precise comparison ($N = 180$), and perceived safety is measured rather than experimentally varied. The subgroup analysis therefore does not identify the causal effect of electoral safety on project choice.

[12] Separate exploratory analyses test whether the conjoint effects vary with perceived and observed housing-market conditions. No pattern recurs across attribute and moderator families (Appendix Figures E.4–E.6).

Figure 5: Politicians perceive a gap between effective policy and voter preferences

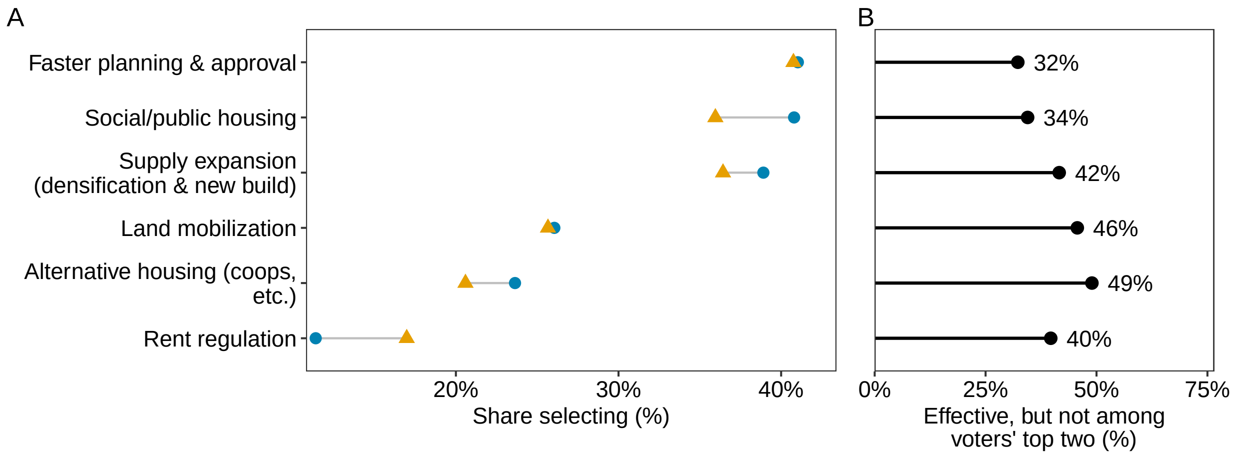

Notes: Panel A compares the share of current officeholders selecting each policy as one of the most effective measures (blue circles, $N = 4{,}341$) with the share selecting it as one of the two measures that receive the strongest support among their own voters (orange triangles, $N = 4{,}041$). Panel B conditions on officeholders who selected each policy as effective and answered the voter-preference question. Conditional sample sizes range from 451 to 1,662.

4.3 Effective policy and perceived voter preferences diverge

Politicians can respond to housing pressure through several policies, only some of which expand supply. The survey asks which policies respondents regard as most effective and which they believe have the greatest support among their own voters. Figure 5 compares these answers across respondents and within the same politician.

Faster planning, social or public housing, and supply expansion are the three policies most often selected as effective. Politicians perceive less voter support for each supply-side response. Rent regulation reverses the pattern: politicians select it as effective less often than they rank it among the two policies with strongest voter support. Among the 1,582 officeholders who select supply expansion as effective and answer the voter-preference question, 42 percent do not rank it among the two measures they believe receive the strongest support from their own voters. The pattern shows that politicians can see policy value in additional supply while expecting competing instruments to offer greater electoral rewards.

Figure 6: Politicians rarely expect additional housing to raise prices or rents

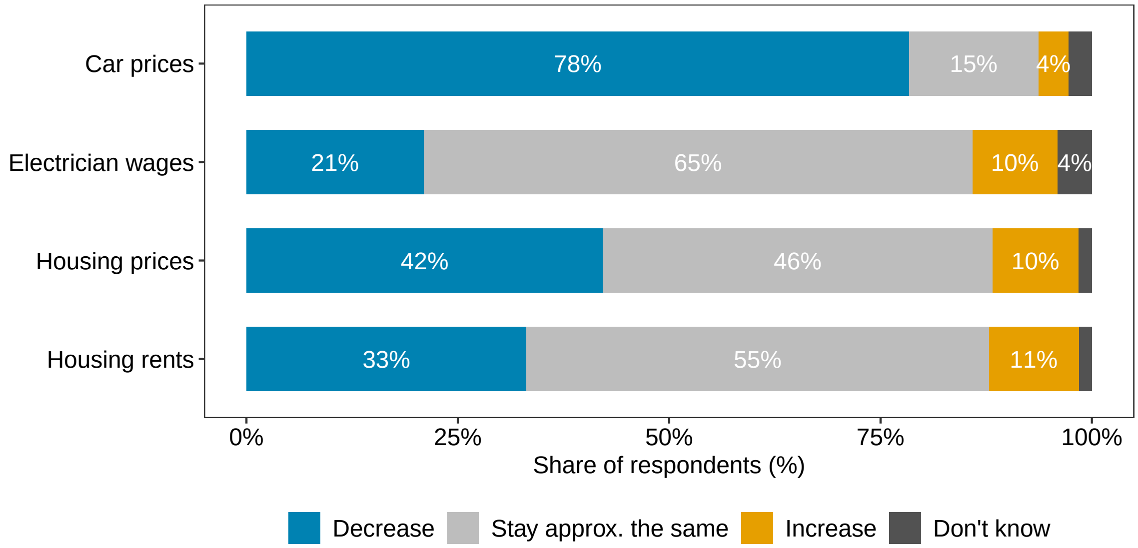

Notes: The figure compares current officeholders' beliefs about supply responses across product, labor, purchase-price, and rent domains (N = 4,402–4,407). The car-price responses are recoded from an equivalent question framed around a decrease in supply.

It does not establish that every respondent would favor a particular housing project after weighing its other costs.

4.4 Politicians have reasons to expand supply

Politicians nevertheless treat housing as important. Four in five officeholders rate housing as a rather or very important local issue. Appendix Figure E.10 places this assessment alongside five other municipal policy areas. I next examine whether politicians expect additional supply to improve affordability rather than make it worse.

Only about one in ten officeholders expect additional housing to raise purchase prices or rents. This belief is far less common than the supply skepticism documented among US voters by Elmendorf, Nall and Oklobdzija (2025). Forty-two percent expect purchase prices to fall and 33 percent expect rents to fall, while the modal response is that both remain approximately unchanged. Politicians therefore rarely view additional supply as counterproductive, although many expect a limited affordability effect.

Thirty-nine percent select supply expansion as one of the two most effective measures for addressing local housing problems. Faster planning and social or public housing rank slightly higher, and both can also expand housing availability. Politicians thus see housing as a problem and often regard supply-side action as useful. These responses and the conjoint estimates suggest that politicians weigh opposition and fiscal burdens that arise before a project's benefits materialize and perceive that voters prefer other responses.

The survey establishes two premises of the argument. Politicians often see benefits in expanding supply, and their project choices respond to opposition, fiscal effects, planning demands, and delay. Many also expect voters to prefer other instruments. The survey does not vary electoral risk, so it cannot identify mediation of the German permit decline.

5 Electoral Risk and Housing Supply in US Cities

Sections 3 and 4 show that housing supply falls as electoral risk rises and that the costs identified by the theory shape politicians' project choices. Section 2 predicts the reverse when electoral accountability weakens: once reelection is no longer possible, losing support no longer threatens another term in the same office. Term limits provide this test. In a political-agency model, Smart and Sturm (2013) show that reelection incentives can induce politicians to ignore their private information and choose a politically safe action. Term limits remove this distortion in their model. Existing studies find that term limits change legislative effort and fiscal policy (Alt, Bueno de Mesquita and Rose, 2011; Besley and Case, 1995; Fouirnaies and Hall, 2022).

I test this prediction in US cities, where mayoral term limits are common: 191 of the 754 cities in the data operated under one between 1990 and 2024.[13] The United States

[13] Author's calculation from the term-limit histories described below.

provides a demanding test because leading accounts attribute housing restraint to homeowner interests and unrepresentative participation. Roughly two-thirds of US households own their home, and municipal reliance on property taxes gives many voters a direct stake in local land use (Fischel, 2001).[14] Discretionary hearings amplify these interests because homeowners and nearby residents dominate public comment on housing proposals (Einstein, Glick and Palmer, 2019; Sahn, 2025). Off-cycle elections narrow participation further (Anzia, 2014). Opposition also extends beyond homeowners, as renters in expensive cities can resist nearby construction (Hankinson, 2018). My argument predicts that the officeholder's electoral position determines whether these pressures bind.

I combine mayoral elections from de Benedictis-Kessner, Jones and Warshaw (2025) and the American Local Elections Database (de Benedictis-Kessner et al., 2023) with annual permit counts from the Census Building Permits Survey for 754 cities from 1990 through 2024. I code final terms from city charters, ballot measures, official election returns, and mayoral service records. Appendix F describes the data construction, coding decisions, and sample restrictions.

I compare the same mayor's permitting in her final term with her permitting in earlier terms. I estimate a Poisson model of annual permit counts with mayor and calendar-year fixed effects, a log-population offset, and standard errors clustered by city:

$$\mathbb{E}[Y_{ict} \mid \mathbf{X}] = \exp\left\{\tau L_{ict} + \alpha_i + \lambda_t + \log P_{ct}\right\}, \tag{4}$$

where L_{ict} indicates that mayor i serves in a final term. The comparison is identified by 220 mayors in 112 cities who are observed both before and after reaching the term limit under a dated and sourced statutory ceiling. Eligibility changes according to the city's term-limit rule rather than the outcome of an election. The main concern is that final terms occur later in a mayor's career, so experience and ineligibility change together. A term limit also

[14]Homeownership rates are from the Census Bureau's Housing Vacancy Survey and Eurostat's income and living conditions statistics.

bars another term in the same office without necessarily ending the mayor's electoral career. Mayors who seek higher office continue to face voters, so the estimate may understate the response when a term limit also ends the mayor's electoral career (Fouirnaies and Hall, 2022). Appendix F reports the legal coding and alternative specifications.

Figure 7: Term limits and housing permits in US cities

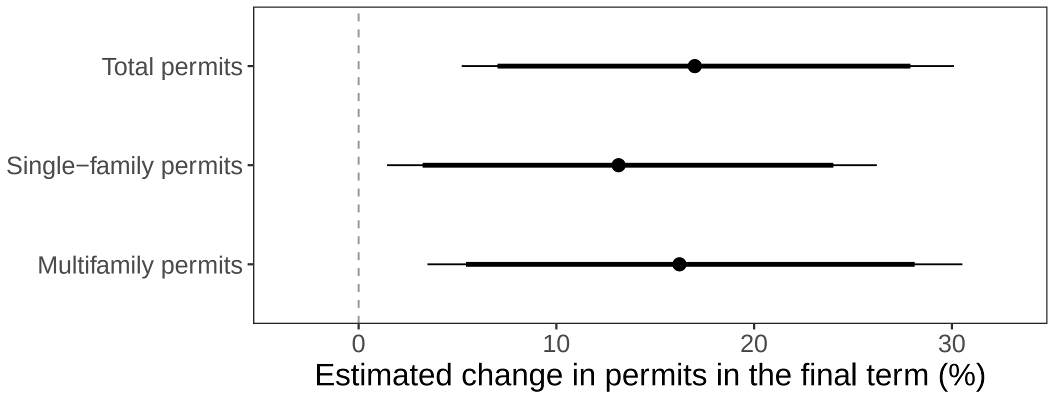

Notes: The figure reports $100 \times [\exp(\hat{\tau}) - 1]$ from Equation 4. The models compare the same mayor's permitting in her final term with earlier terms and include mayor and calendar-year fixed effects with a log-population offset. The sample contains 220 mayors in 112 cities from 1990 through 2024. Election years are omitted, and standard errors are clustered by city. Thin bars show 95 percent confidence intervals and thick bars show 90 percent confidence intervals.

Figure 7 reports the term-limit estimates. Cities issue 17.0 percent more total permits in a mayor's final term. Multifamily permits increase by 16.2 percent, while single-family permits increase by 13.1 percent. The German and US estimates move in opposite directions, as the argument predicts: permitting falls after a mayor loses electoral support and rises when a term limit ends eligibility for another term. The final-term estimates for total and multifamily permits remain positive when excluding mayors who later return to office, restricting the comparison to cities with two-term ceilings, focusing on cities above the median population, and allowing each state to follow a separate linear trend (Appendix Figure F.1).

6 Conclusion

Electoral risk determines whether the immediate political costs of construction constrain local politicians' supply decisions. I show this by following German mayors who remain in office after a sharp loss of electoral support and US mayors after a term limit ends their eligibility for another term. Permitting falls in the first setting and rises broadly in the second, while the German decline is largest for multifamily housing. An original survey of German local politicians identifies the project features that create this pressure: resident opposition, fiscal burdens, planning demands, and delay. The survey supports this account by showing that politicians value additional supply but penalize immediate costs and delay.

These findings show how a classic argument for political insulation applies to housing. Its advocates held that voter pressure favors current consumption over investment, while insulated governments can bear present costs for future growth (Przeworski and Limongi, 1993; Przeworski et al., 2000). Regime comparisons could not show whether electoral pressure changed the decisions behind aggregate growth. Housing makes the policy response observable. A project creates an immediate and attributable electoral penalty, while its benefits arrive only after planning and construction and remain difficult to credit to the official who approved it. Electoral accountability is therefore asymmetric in long-term policy: politicians can be punished for present costs before they can be rewarded for future benefits. Electoral risk determines how strongly this asymmetry constrains them. Once a term limit removes the value of another term in the same office, municipal permitting rises under the same mayor. Accountability can therefore deter investments that officeholders themselves value when their political costs arrive first.

Term limits, higher-level mandates, and fiscal incentives target this political constraint. Term limits most directly remove reelection pressure, and municipal permitting rises under the same US mayor once that pressure is gone. Yet term limits are blunt because they weaken accountability across every issue. Their welfare effect is ambiguous because removing the reelection constraint gives officeholders more freedom to act on their own information (Smart and Sturm, 2013). Higher-level mandates shift responsibility for a housing decision away from local politicians. If state or national law requires municipalities to allow apartment buildings in specified places, local politicians can tell opponents that their hands are tied. This approach will increase supply only if the government that gains authority favors construction. Fiscal incentives work differently. Payments for each home and reimbursements for the schools, roads, and services needed by new residents make approval more attractive and let politicians reassure voters that housing will not crowd out current services, a concern that makes public housing electorally costly (Hilbig and Wiedemann, 2026).

Housing is one instance of a broader political problem. Climate mitigation, infrastructure, and pension reform also ask elected officials to incur visible costs before benefits arrive and before beneficiaries can reward them (Finnegan, 2022; Gailmard and Patty, 2019; Jacobs, 2011). In each domain, long-term policy depends on whether officeholders can withstand opposition at the moment of choice. Electoral risk therefore helps explain why governments defer investments even when officeholders recognize their long-run value.

References

Abou-Chadi, Tarik, Denis Cohen and Thomas Kurer. 2025. "Rental Market Risk and Radical Right Support." *Comparative Political Studies* 58(13):2866–2901.
URL: *https://journals.sagepub.com/doi/10.1177/00104140241306963*

Ahlfeldt, Gabriel M., Stephan Heblich and Tobias Seidel. 2023. "Micro-geographic property price and rent indices." *Regional Science and Urban Economics* 98:103836.

Alt, James, Ethan Bueno de Mesquita and Shanna Rose. 2011. "Disentangling Accountability and Competence in Elections: Evidence from U.S. Term Limits." *The Journal of Politics*

73(1):171–186.

Anderson, Carter, Ella Briman, Aidan Ferrin, Charlotte Hampton, Emelia Malhotra, Spriha Pandey, James Robinson, Lila Sugerman, Jesse VanNewkirk, Marina Wang, Jessica Yu, Bill Zheng and Brendan Nyhan. 2025. "Debunking NIMBY Myths Increases Support for Affordable Housing, Especially Near Respondents' Homes." *Journal of Experimental Political Science* pp. 1–14.

Ansell, Ben, Frederik Hjorth, Jacob Nyrup and Martin Vinæs Larsen. 2022. "Sheltering Populists? House Prices and the Support for Populist Parties." *The Journal of Politics* 84(3):1420–1436.
URL: *https://www.journals.uchicago.edu/doi/10.1086/718354*

Anzia, Sarah F. 2014. *Timing and Turnout: How Off-Cycle Elections Favor Organized Groups*. Chicago: University of Chicago Press.

Arnold, Felix. 2018. "Turnout and Closeness: Evidence from 60 Years of Bavarian Mayoral Elections." *The Scandinavian Journal of Economics* 120(2):624–653.

Arnold, R. Douglas. 1990. *The Logic of Congressional Action*. New Haven: Yale University Press.

Baker, Andrew C., David F. Larcker and Charles C.Y. Wang. 2022. "How much should we trust staggered difference-in-differences estimates?" *Journal of Financial Economics* 144(2):370–395.

Besley, Timothy and Anne Case. 1995. "Does Electoral Accountability Affect Economic Policy Choices? Evidence from Gubernatorial Term Limits." *The Quarterly Journal of Economics* 110(3):769–798.

Bratu, Cristina, Oskari Harjunen and Tuukka Saarimaa. 2023. "JUE Insight: City-wide effects of new housing supply: Evidence from moving chains." *Journal of Urban Economics* 133:103528.

Bremer, Björn, Donato Di Carlo and Leon Wansleben. 2023. "The constrained politics of local public investment under cooperative federalism." *Socio-Economic Review* 21(2):1007–1034.

Broockman, David, Christopher S. Elmendorf and Joshua Kalla. 2024. "The Symbolic Politics of Housing." OSF preprint.
URL: *https://osf.io/preprints/osf/surv9*

Broockman, David E. and Christopher Skovron. 2018. "Bias in Perceptions of Public Opinion among Political Elites." *American Political Science Review* 112(3):542–563.

Brouwer, N.R. and Jessica Trounstine. 2024. "NIMBYs, YIMBYs, and the Politics of Land Use in American Cities." *Annual Review of Political Science* 27(1):165–184.
URL: *https://www.annualreviews.org/content/journals/10.1146/annurev-polisci-041322-041133*

Callaway, Brantly and Pedro H.C. Sant'Anna. 2021. "Difference-in-Differences with multiple time periods." *Journal of Econometrics* 225(2):200–230.
URL: *https://linkinghub.elsevier.com/retrieve/pii/S0304407620303948*

Canes-Wrone, Brandice, Michael C. Herron and Kenneth W. Shotts. 2001. "Leadership and Pandering: A Theory of Executive Policymaking." *American Journal of Political Science* 45(3):532–550.

Chen, Jiafeng and Jonathan Roth. 2024. "Logs with Zeros? Some Problems and Solutions." *The Quarterly Journal of Economics* 139(2):891–936.

Chou, Winston and Rafaela Dancygier. 2021. "Why Parties Displace Their Voters: Gentrification, Coalitional Change, and the Demise of Public Housing." *American Political Science Review* 115(2):429–449.

Clayton, Katherine, Yusaku Horiuchi, Aaron R. Kaufman, Gary King and Mayya Komisarchik. 2026. "Correcting Measurement Error Bias in Conjoint Survey Experiments."

American Journal of Political Science pp. 1–19.
URL: *https://onlinelibrary.wiley.com/doi/10.1111/ajps.70076*

Dancygier, Rafaela and Andreas Wiedemann. 2025. "The financialization of housing and its political consequences." *American Journal of Political Science* 69(4):1354–1373.
URL: *https://onlinelibrary.wiley.com/doi/10.1111/ajps.12928*

Dancygier, Rafaela M. 2010. *Immigration and Conflict in Europe*. Cambridge University Press.

de Benedictis-Kessner, Justin and Christopher Warshaw. 2020. "Accountability for the Local Economy at All Levels of Government in United States Elections." *American Political Science Review* 114(3):660–676.

de Benedictis-Kessner, Justin, Diana Da In Lee, Yamil R. Velez and Christopher Warshaw. 2023. "American local government elections database." *Scientific Data* 10(1):912.
URL: *https://www.nature.com/articles/s41597-023-02792-x*

de Benedictis-Kessner, Justin, Daniel Jones and Christopher Warshaw. 2025. "How partisanship in cities influences housing policy." *American Journal of Political Science* 69(1):64–77.
URL: *https://onlinelibrary.wiley.com/doi/10.1111/ajps.12856*

de Chaisemartin, Clément and Xavier D'Haultfœuille. 2020. "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects." *American Economic Review* 110(9):2964–2996.

Diamond, Rebecca, Tim McQuade and Franklin Qian. 2019. "The Effects of Rent Control Expansion on Tenants, Landlords, and Inequality: Evidence from San Francisco." *American Economic Review* 109(9):3365–3394.

Einstein, Katherine Levine, David M. Glick and Maxwell Palmer. 2019. *Neighborhood Defenders: Participatory Politics and America's Housing Crisis*. Cambridge University Press.
URL: *https://www.cambridge.org/core/product/identifier/9781108769495/type/book*

Elmendorf, Christopher S., Clayton Nall and Stan Oklobdzija. 2025. "The Folk Economics of Housing." *Journal of Economic Perspectives* 39(3):45–66.
URL: *https://pubs.aeaweb.org/doi/10.1257/jep.20241428*

Fernandez, Raquel and Dani Rodrik. 1991. "Resistance to Reform: Status Quo Bias in the Presence of Individual-Specific Uncertainty." *American Economic Review* 81(5):1146–1155.

Findeisen, Francesco. 2022. "The Housing Crisis in Superstar Cities: Labour Markets, Price Inflation, and Financialization." *European Journal of Sociology* 63(3):363–392.

Finnegan, Jared J. 2022. "Institutions, Climate Change, and the Foundations of Long-Term Policymaking." *Comparative Political Studies* 55(7):1198–1235.

Fischel, William A. 2001. *The Homevoter Hypothesis*. Cambridge: Harvard University Press.
URL: *http://www.jstor.org/stable/10.2307/j.ctv1p6hp64*

Fouirnaies, Alexander and Andrew B. Hall. 2022. "How Do Electoral Incentives Affect Legislator Behavior? Evidence from U.S. State Legislatures." *American Political Science Review* 116(2):662–676.

Fu, Jiawei and Tara Slough. 2026. "Heterogeneous Treatment Effects and Causal Mechanisms." *American Political Science Review* pp. 1–18.

Gailmard, Sean and John W. Patty. 2019. "Preventing Prevention." *American Journal of Political Science* 63(2):342–352.

Glaeser, Edward L and Erzo F. P Luttmer. 2003. "The Misallocation of Housing Under Rent Control." *American Economic Review* 93(4):1027–1046.

Granzier, Riako, Vincent Pons and Clémence Tricaud. 2023. "Coordination and Bandwagon Effects: How Past Rankings Shape the Behavior of Voters and Candidates." *American Economic Journal: Applied Economics* 15(4):177–217.

Gyourko, Joseph, Christopher Mayer and Todd Sinai. 2013. "Superstar Cities." *American*

Economic Journal: Economic Policy 5(4):167–199.

Gyourko, Joseph and Raven Molloy. 2015. Regulation and Housing Supply. In *Handbook of Regional and Urban Economics*, ed. Gilles Duranton, J. Vernon Henderson and William C. Strange. Vol. 5 Elsevier pp. 1289–1337.

Hainmueller, Jens, Daniel J. Hopkins and Teppei Yamamoto. 2014. "Causal Inference in Conjoint Analysis: Understanding Multidimensional Choices via Stated Preference Experiments." *Political Analysis* 22(1):1–30.

Hankinson, Michael. 2018. "When Do Renters Behave Like Homeowners? High Rent, Price Anxiety, and NIMBYism." *American Political Science Review* 112(3):473–493.

Hankinson, Michael, Asya Magazinnik and Anna Weissman. 2024. "When Do Local Interest Groups Participate in the Housing Entitlement Process?" *Journal of Political Institutions and Political Economy* 5(1):47–69.

Heddesheimer, Vincent, Hanno Hilbig, Florian Sichart and Andreas Wiedemann. 2025. "GERDA: The German Election Database." *Scientific Data* 12(1):618.
URL: *https://www.nature.com/articles/s41597-025-04811-5*

Held, Alexander and Pauliina Patana. 2023. "Rents, refugees, and the populist radical right." *Research & Politics* 10(2).

Hilbig, Hanno and Andreas Wiedemann. 2026. "How budget trade-offs undermine electoral incentives to build public housing." *American Journal of Political Science* 70(1):202–220.
URL: *https://onlinelibrary.wiley.com/doi/10.1111/ajps.12939*

Hopkins, Daniel J. and Lindsay M. Pettingill. 2018. "Retrospective Voting in Big-City US Mayoral Elections." *Political Science Research and Methods* 6(4):697–714.

Institut für Demokratie- und Partizipationsforschung. 2025. "Datenbank Bürgerbegehren." Wuppertal: Bergische Universität Wuppertal, in cooperation with Mehr Demokratie e.V.
URL: *https://www.datenbank-buergerbegehren.info/*

Ioannides, Yannis M. and L. Rachel Ngai. 2025. "Housing and Inequality." *Journal of Economic Literature* 63(3):916–963.

Jacobs, Alan M. 2011. *Governing for the Long Term.* Cambridge University Press.

Kayser, Mark Andreas and René Lindstädt. 2015. "A Cross-National Measure of Electoral Competitiveness." *Political Analysis* 23(2):242–253.

Kydland, Finn E. and Edward C. Prescott. 1977. "Rules Rather than Discretion: The Inconsistency of Optimal Plans." *Journal of Political Economy* 85(3):473–491.

Larsen, Martin Vinæs. 2021. "How Do Voters Hold Politicians Accountable for Personal Welfare? Evidence of a Self-Serving Bias." *The Journal of Politics* 83(2):740–752.

Larsen, Martin Vinæs, Frederik Hjorth, Peter Thisted Dinesen and Kim Mannemar Sønderskov. 2019. "When Do Citizens Respond Politically to the Local Economy? Evidence from Registry Data on Local Housing Markets." *American Political Science Review* 113(2):499–516.

Larsen, Martin Vinæs and Niels Nyholt. 2024. "Understanding Opposition to Apartment Buildings." *Journal of Political Institutions and Political Economy* 5(1):29–46.

Maskin, Eric and Jean Tirole. 2004. "The Politician and the Judge: Accountability in Government." *American Economic Review* 94(4):1034–1054.

Mast, Evan. 2023. "The effect of new market-rate housing construction on the low-income housing market." *Journal of Urban Economics* 133:103383.

Mense, Andreas. 2025. "The Impact of New Housing Supply on the Distribution of Rents." *Journal of Political Economy Macroeconomics* 3(1):1–42.

Mildenberger, Matto and Alexander Sahn. 2025. "The Effect of Policy Traceability on Legislative Incentives." *Legislative Studies Quarterly* 50(4):e70036.

Molloy, Raven, Charles G. Nathanson and Andrew Paciorek. 2022. "Housing supply and affordability: Evidence from rents, housing consumption and household location." *Journal* *of Urban Economics* 129:103427.

Müller, Wolfgang C. and Kaare Strøm. 1999. *Policy, Office, or Votes? How Political Parties in Western Europe Make Hard Decisions*. Cambridge: Cambridge University Press.

OECD. 2017. *Land-use Planning Systems in the OECD*. Paris: OECD Publishing.

Olson, Mancur. 1965. *The Logic of Collective Action*. Harvard University Press.

Peterson, Paul E. 1981. *City Limits*. Chicago: University of Chicago Press.

Przeworski, Adam and Fernando Limongi. 1993. "Political Regimes and Economic Growth." *Journal of Economic Perspectives* 7(3):51–69.

Przeworski, Adam, Michael E. Alvarez, José Antonio Cheibub and Fernando Limongi. 2000. *Democracy and Development: Political Institutions and Well-Being in the World, 1950–1990*. Cambridge: Cambridge University Press.

Rambachan, Ashesh and Jonathan Roth. 2023. "A More Credible Approach to Parallel Trends." *Review of Economic Studies* 90(5):2555–2591.
URL: *https://academic.oup.com/restud/article/90/5/2555/7039335*

Reisenbichler, Alexander and Pascal D. Koenig. 2025. "Cause for Celebration or Concern? Voter Reactions to Rising House Prices." *Comparative Political Studies* 58(7):1352–1388.
URL: *https://journals.sagepub.com/doi/10.1177/00104140241269896*

Roth, Jonathan, Pedro H.C. Sant'Anna, Alyssa Bilinski and John Poe. 2023. "What's trending in difference-in-differences? A synthesis of the recent econometrics literature." *Journal of Econometrics* 235(2):2218–2244.

Sahn, Alexander. 2025. "Public comment and public policy." *American Journal of Political Science* 69(2):685–700.
URL: *https://onlinelibrary.wiley.com/doi/10.1111/ajps.12900*

Sheffer, Lior, Peter John Loewen and Jack Lucas. 2024. "Long-term policymaking and politicians' beliefs about voters: Evidence from a 3-year panel study of politicians." *Governance*

37(2):395–410.

Sichart, Florian and Markus Wagner. 2025. "Beyond Isolated Effects: Rethinking Causal Questions in the Presence of Linked Attributes." OSF preprint.
URL: *https://osf.io/preprints/osf/7xskg*

Smart, Michael and Daniel M. Sturm. 2013. "Term Limits and Electoral Accountability." *Journal of Public Economics* 107:93–102.

Solé-Ollé, Albert and Elisabet Viladecans-Marsal. 2012. "Lobbying, political competition, and local land supply: Recent evidence from Spain." *Journal of Public Economics* 96(1-2):10–19.

Soroka, Stuart N. 2014. *Negativity in Democratic Politics: Causes and Consequences*. Cambridge: Cambridge University Press.

Strotz, R. H. 1955. "Myopia and Inconsistency in Dynamic Utility Maximization." *The Review of Economic Studies* 23(3):165–180.

Trounstine, Jessica. 2018. *Segregation by Design*. Cambridge University Press.

Trounstine, Jessica. 2023. "You Won't be My Neighbor: Opposition to High Density Development." *Urban Affairs Review* 59(1):294–308.

Weaver, R. Kent. 1986. "The Politics of Blame Avoidance." *Journal of Public Policy* 6(4):371–398.

Yu, Rui. 2026. "Returns to Political Contributions in Local Housing Markets." *Review of Economics and Statistics* pp. 1–44.
URL: *https://direct.mit.edu/rest/article-abstract/doi/10.1162/REST.a.1732*

Supplementary Information

A Interviews

A.1 Recruitment, ethics, and use

Princeton University's IRB approved the interview protocol (#16641). Before each interview, participants were told the study's purpose, the intended use of quotations, the voluntary nature of participation, and the available confidentiality protections. Participants gave oral informed consent and could end the interview or decline any question. No compensation or deception was used.

Recruitment was purposive rather than probability-based. Potential participants were contacted because their political office, administrative role, professional work, or research gave them direct knowledge of German housing policy. The invitation offered an interview in person, by telephone, or by Zoom. The corpus therefore covers deliberately varied positions in the housing process, not a sampling frame from which response rates or population frequencies can be inferred.

The corpus contains 28 interviews conducted between September 2024 and January 2026. Twenty-six were recorded and transcribed in German with the same automated speech-recognition system (Whisper large-v3-turbo). Two were documented in contemporaneous notes only. Quotations in the manuscript were checked against these transcripts and, where available, an independently generated Zoom transcript. They were translated from German by the author. Generic IDs and role descriptions protect confidentiality. The interviews motivate and illustrate assumptions in the argument and informed the survey instrument. They are not a representative sample or an independent estimate of how common any mechanism is.

A.2 Semi-structured guide

The guide was adapted to each participant's role and covered four common domains: the causes of local housing shortages, the role of electoral and party competition in planning decisions, the stakeholders who influence housing policy, and the relative effectiveness and political feasibility of supply, planning, subsidy, and rent policies. Follow-up questions asked about the timing of costs and benefits and about concrete development decisions.

A.3 Interview roster

Table A.1: Completed interviews

ID	Generic role	Date
Federal level		
A1	Member of the German Parliament	September 26, 2024
A2	Chief of staff to a member of the German Parliament	October 30, 2024
A3	Former state secretary	November 12, 2024
State level		

ID	Generic role	Date
B1	Member of a state parliament	November 11, 2024
B2	Member of a state parliament	November 15, 2024
B3	State-government official	December 5, 2024
B4	State-government official	December 11, 2024
B5	State legislator and building-policy spokesperson	February 7, 2025
Local level		
F1	Former state and federal legislator	November 6, 2024
F2	City councillor	December 4, 2024
F3	Head of planning in a major German city	December 16, 2024
F4	Mayor	November 11, 2025
F5	Planning official in a major German city	November 13, 2025
F6	Mayor	November 19, 2025
F7	Mayor	December 29, 2025
F8	Mayor	January 7, 2026
F9	Works in the administration of a major German city	January 7, 2026
F10	Mayor	January 8, 2026
Interest groups		
D1	Real-estate industry representative	November 15, 2024
D2	Tenant-association representative	November 20, 2024
D3	Landlord-association representative	December 4, 2024
D4	Regional economic-area representative	December 13, 2024
Private actors		
E1	Real-estate business owner	December 4, 2024
E2	Real-estate business owner	December 9, 2024
Experts		
G1	Research expert	September 20, 2024
G2	Research expert	October 1, 2024
G3	Research expert	November 12, 2024
G4	Research expert	November 5, 2024

Notes: Names, organizations, and identifying position details are withheld. Generic role descriptions preserve the distinctions relevant to the research design.

B German Observational Data

B.1 Building permits and completions

The housing outcomes come from the Building Permits Statistics (*Baugenehmigungsstatistik*) and Building Completions Statistics (*Baufertigstellungsstatistik*), accessed through restricted on-site use at the Research Data Centers of the German statistical offices. The underlying registries record the universe of permitted and completed construction. The extracted files contain approximately 315,700 permit records and 309,200 completion records at the municipality-year-source level. Across the two sources, nearly 370 columns describe dwelling units, buildings, living space, costs, building height, tenure, and builder category.

I aggregate counts to the municipality-year and retain raw permit counts for estimation. Total permitted dwelling units are divided into single-family and multifamily construction. Additional outcomes distinguish low-rise from multi-story buildings and separate private households, housing companies and real-estate funds, and public or non-profit developers. The last category combines public developers, including municipalities, counties, states, and social-insurance bodies, with non-profit organizations. It records who built the unit and carries no information about subsidy status. Subsidized units built by private housing companies remain in the corporate category, while unsubsidized units built by a public or non-profit body enter the public/non-profit category. Completion outcomes use the corresponding registry and therefore test whether the permitting response becomes realized construction rather than merely a paper authorization.

B.2 Boundary harmonization

German municipal boundaries change frequently through mergers, incorporations, and identifier revisions. I map every source record to a consistent municipal geography before joining elections and local covariates. Aggregations are additive for permit and completion counts. Population and area are harmonized to the same unit before constructing density and per-capita rates. Municipality-year keys are unique after the build. The final theory-test panel contains 373,255 rows and 127 columns. The 1995-onward housing analysis contains 318,365 municipality-years in 10,981 municipalities.

B.3 Local elections, mayor identity linkage, and privacy

The harmonized GERDA files contain 73,535 municipal council election records and 49,717 non-superseded mayoral contests. Council turnout averages 66 percent across the recorded elections. Across mayoral contests with turnout data, the aggregate ratio of voters to eligible voters is 59 percent. The council files also record party vote and seat shares. The mayoral candidate file supplies first-round vote shares, candidate counts, and winner status.

GERDA's mayoral election data link consecutive terms using public candidate records, official person-level term counters, and a conservative cross-state alias map. The Hessen workbook supplies term and re-election counters even where names are redacted, expanding longitudinal coverage without reconstructing protected identities. Every linked row records an identity method and, where available, the official term and re-election counters used in the match.

The build converts name and municipality combinations to pseudonymous identifiers before creating downstream files. Candidate names and restricted source rows are absent from all analysis exports. The main treatment requires only public first-round vote shares and candidate counts. Restricted challenger identities are neither necessary nor used.

Table B.1: Coverage of the German housing and mayor data

Quantity	Value
Municipality-years, 1995 onward	318,365
Municipalities, 1995 onward	10,981
Permitted units	9,518,869
Single-family permitted units	4,441,702
Multifamily permitted units	5,077,167
Municipality-years linked to a mayor	115,974
Distinct linked mayors	13,137
Mayors observed in a second or later term	6,633
Vote-loss transitions	262
Stable-support transitions	1,209
Municipalities in stacked design	1,183

Notes: Permit totals and mayor counts refer to the harmonized 1995-onward analysis panel. A mayor is counted as term two or later only when `mayor_term_number` is at least two. Repeated annual observations within a single term do not qualify.

C Additional German Observational Results

C.1 Stacked episodes and treatment timing

Each episode joins two consecutive terms of one re-elected mayor in a single municipality. The mayor's party label and number of first-round candidates remain unchanged across the two elections, which must be three to nine years apart. Ninety-eight percent of eligible episodes are five to seven years apart. Equation 2 classifies the resulting episodes as treated or stable-support episodes. In the latter group, the mayor's vote share changes by less than five percentage points in either direction. Election years are excluded, and the post-election period begins one year after re-election. The 262 treated and 1,209 stable-support episodes contribute 14,185 episode-year observations before outcome-specific Poisson fixed-effect separation.

Figure C.1 shows the calendar timing of every episode, sorted by focal re-election year. Table C.1 reports state composition separately because Equation 3 absorbs shocks at the state-by-calendar-year level. Separating composition from timing prevents differences in state sample size from appearing as differences in temporal coverage.

Figure C.1: Treatment timing in the stacked design

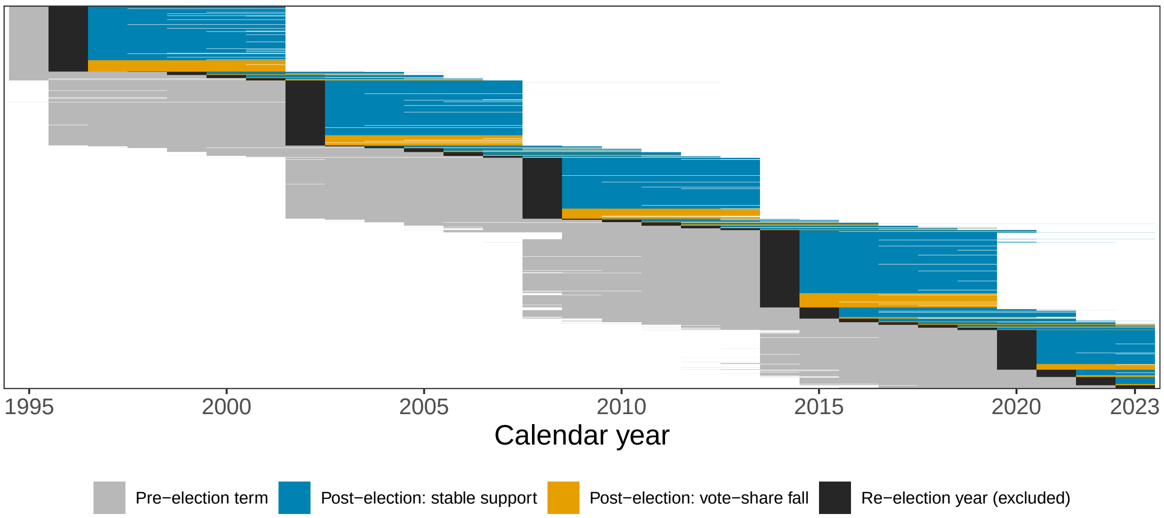

Notes: Each row is one episode spanning consecutive terms of a re-elected mayor, sorted by the focal re-election year. Gray cells mark the pre-election term. Blue and orange cells identify the post-election periods for the stable-support and vote-share-fall episodes defined in Equation 2. Black cells mark the excluded re-election year. Blank cells lie outside the two terms included in that episode.

Table C.1: State composition of the stacked design

State	Focal years	Stable support	Vote-share fall	Total
Lower Saxony	2014–2021	18	2	20
North Rhine-Westphalia	2014–2020	46	8	54
Hesse	1999–2023	125	46	171
Rhineland-Palatinate	2002–2022	7	1	8
Bavaria	1996–2023	937	192	1,129
Brandenburg	2019–2021	1	2	3
Saxony	2015–2023	63	11	74
Saxony-Anhalt	2001–2023	12	0	12

Notes: Focal years are the first and last re-election years among qualifying episodes in each state; they do not describe the underlying GERDA data coverage. Stable support and vote-share fall refer to the groups defined in Equation 2. Each episode contributes once.

C.2 Own vote share and mayors' perceived electoral safety

Treatment in the stacked design is a fall of more than ten points in a returning mayor's first-round share. I use the politician survey to test whether own vote share maps onto mayors' perceived electoral safety. The survey asked current officeholders how close their last election was (1 = very close to 4 = very clear) and how likely their re-election is (1 = very unlikely to 5 = very likely). I sum the two items into a perceived-safety index (2–9) and link mayors who identified their municipality to the most recent mayoral election before the survey. This yields 298 linked mayors with the index.[15]

Own first-round share strongly predicts perceived safety. With state fixed effects, the index rises by 0.43 points per ten points of own share (SE 0.04), a third of a standard deviation, and own share explains a quarter of the within-state variation. Both components move: the last election is judged clearer (0.28 points per ten points) and re-election likelier (0.16 points). Among mayors who won with less than 70 percent, 36 percent call re-election "very likely"; among those who won with at least 90 percent, 68 percent do. Figure C.2 shows that the relationship is approximately linear.

Figure C.2: Mayors' perceived electoral safety by own first-round vote share

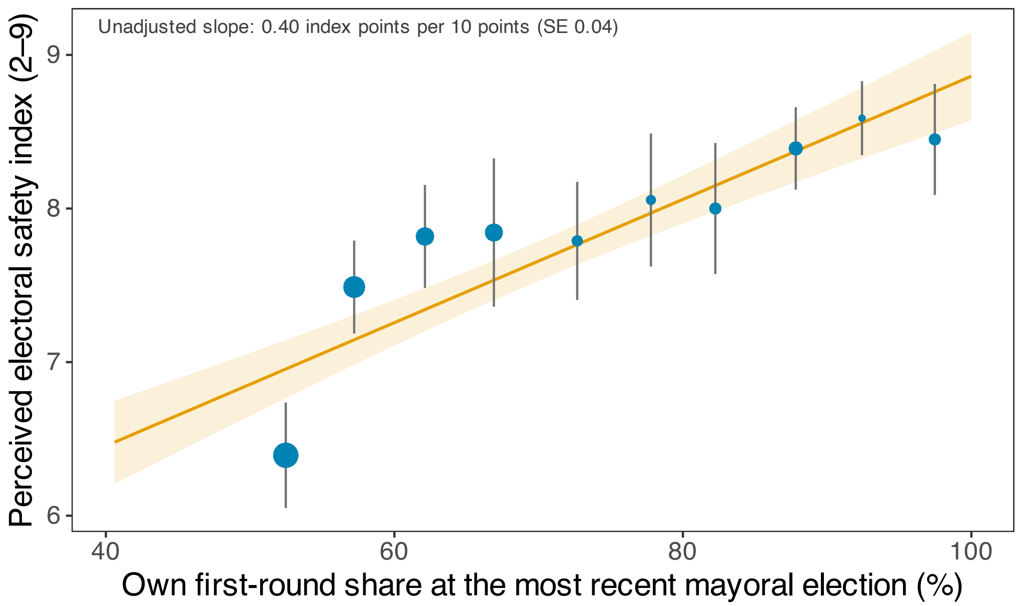

Notes: Bin means of the perceived-safety index (2–9) within 5-point bins of the mayor's own first-round share at the most recent mayoral election before the survey, for bins with at least five respondents. Vertical lines are 95 percent confidence intervals and point size is proportional to the number of mayors in the bin. The line is an unadjusted linear fit on the underlying 298 mayors, shown with its 95 percent confidence band and slope. With state fixed effects, the index rises by 0.43 points per ten points of own share (SE 0.04).

To separate current standing from the path to it, I restrict the sample to 91 mayors whose preceding election was won by the same person. Conditional on current share, the coefficient on the change since that election is -0.05 (SE 0.11), while the current-share coefficient is 0.45 (SE 0.11). The estimate detects no separate relationship between perceived safety and how the mayor reached her current position. The survey observes each mayor once and includes only 15 respondents who meet the treatment definition, so it cannot reproduce the stacked comparison directly. In the stacked design, treated mayors fall on average from 81 to 65 percent, while stable-support mayors remain at 83 percent. Applying the level slope to that 16-point gap gives 0.7 index points, or 0.55 standard deviations of perceived safety.

[15] Elections held during fielding in 2026 are excluded because respondents did not yet know their results. The most recent election precedes the survey by up to six years.

C.3 Model sample and unadjusted trajectories

Table C.2 describes the exact episode sample used in the stacked design. Each episode receives equal weight. Permit rates and municipality characteristics are measured before the focal re-election, so they cannot reflect the post-election response. The final three columns report the difference between vote-share-fall and stable-support episodes, its 95 percent confidence interval, and a direct test of equality. The difference in vote-share changes is mechanical because that change defines the groups. Beyond this difference, vote-share-fall episodes begin with 1.7 percentage points less support and 0.40 fewer single-family permits per 1,000 residents. The remaining differences are not statistically distinguishable from zero at the 5 percent level. The baseline-support imbalance is examined further in Appendix Table C.5.

Table C.2: Descriptive statistics and group differences in the stacked design

Variable	Vote-share fall	Stable support	Difference	95% CI	p-value
Baseline mayor vote share (percent)	81.5 (10.2)	83.1 (14.8)	-1.7	[-3.1, -0.2]	0.029
Change in mayor vote share (percentage points)	-16.1 (5.4)	-0.5 (2.5)	-15.6	[-16.3, -15.0]	< 0.001
Pre-election population (thousands)	5.2 (6.5)	6.2 (25.6)	-0.9	[-2.6, 0.7]	0.261
Pre-election population density	162.3 (178.5)	159.5 (219.7)	2.8	[-22.4, 28.0]	0.826
Pre-election total permits per 1,000 residents	4.88 (3.36)	5.27 (3.64)	-0.39	[-0.84, 0.06]	0.089
Pre-election single-family permits per 1,000 residents	3.29 (2.35)	3.69 (2.58)	-0.40	[-0.71, -0.09]	0.012
Pre-election multifamily permits per 1,000 residents	1.59 (2.02)	1.58 (1.97)	0.01	[-0.26, 0.28]	0.940
Observed pre-election years	5.2 (1.9)	5.1 (2.0)	0.0	[-0.2, 0.3]	0.794
Observed post-election years	4.5 (1.2)	4.5 (1.2)	-0.0	[-0.2, 0.1]	0.912
Episodes	262	1,209			

Notes: The first two columns report episode-level means with standard deviations in parentheses. Differences are vote-share-fall minus stable-support means. Confidence intervals and p-values use standard errors clustered by municipality. The vote-share-change difference is mechanical because the treatment is defined by that change. Pre-election permit rates average the available years before the focal re-election. Population and density are measured in the final observed pre-election year. The table includes all 1,471 episodes before outcome-specific Poisson fixed-effect separation.

Figure C.3 plots unadjusted permit rates for the vote-loss and stable-support groups. It is a descriptive view of the outcome paths, not an alternative estimator. The adjusted comparison in Figure 2 accounts for episode fixed effects, state-by-calendar-year shocks, baseline support, and pre-election density.

C.4 Dynamic specification

I estimate the event-time coefficients in Figure 2 with the following dynamic extension of Equation 3:

$$\mathbb{E}\big[Y_{et}^h \mid \mathbf{X}\big] = \exp\Big\{\alpha_e + \lambda_{s(e)t} + \sum_{k\in\mathcal{K}} \rho_{kh}\mathbf{1}\{\widetilde{r}_{et} = k\} + \sum_{k\in\mathcal{K}} \tau_{kh} D_e \mathbf{1}\{\widetilde{r}_{et} = k\} + \gamma_h\big(Z_e \times \text{Post}_{et}\big) + \delta_h\big(V_{e0} \times \text{Post}_{et}\big) + \log(\text{Population}_{et})\Big\}. \tag{5}$$

Figure C.3: Unadjusted permit trajectories in the stacked-design sample

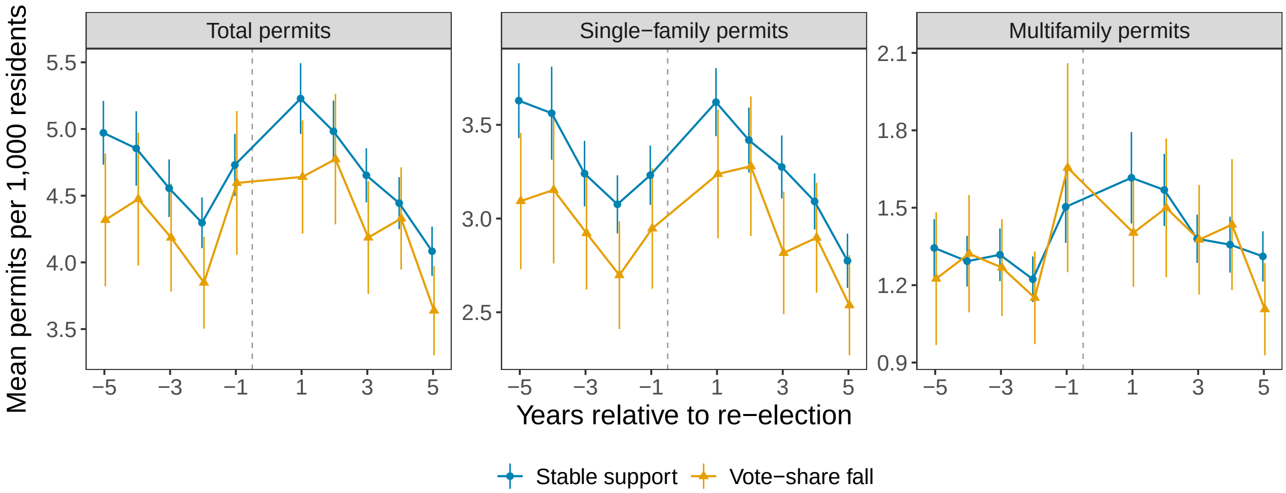

Notes: Points are unadjusted group means for the exact event stack. Permit rates use the population measured in thousands, so a count divided by population is permits per 1,000 residents. Bars are 95 percent intervals around each mean with standard errors clustered by municipality. The focal re-election year is excluded. These trajectories are descriptive and are not used for hypothesis tests.

The relative-time measure $\widetilde{r}_{et}$ is binned at -5 and 5, $\mathcal{K} = \{-5, \ldots, -2, 1, \ldots, 5\}$, and the year before re-election ($k = -1$) is the reference period. The election year remains excluded. The coefficients τ_{kh} compare treated and stable-support episode paths in each relative year. All other terms retain the definitions in Equation 3.

C.5 Direct comparisons across housing outcomes

Figure 3 shows whether the point estimates follow the ordering predicted by the theory. Comparing whether one estimate is statistically distinguishable from zero while another is not would not establish that the two effects differ. I therefore estimate each pair of outcomes jointly. The model gives each outcome its own episode and state-by-calendar-year fixed effects, retains the adjustments in Equation 3, and clusters across both outcomes within municipality. The reported contrast is the difference between the effect on the more politically costly outcome and the effect on its less costly counterpart.

Table C.3 shows the predicted ordering for all four pairwise comparisons. The estimated decline is larger for multifamily permits and completions, multi-story houses, and corporate-developer units. The direct differences range from 8 to 22 percentage points, but every 95 percent confidence interval includes zero. The composition evidence is therefore suggestive rather than a statistically established difference across housing types.

Table C.3: Direct comparisons across housing outcomes

Comparison	Less costly	More costly	Difference	95% CI	p
Single- vs multifamily permits	-7.3%	-16.8%	-10.2%	[-23.8, 5.8]	0.198
Single- vs multifamily completions	-7.7%	-15.4%	-8.4%	[-21.6, 7.0]	0.268
Low-rise vs multi-story houses	-7.1%	-27.8%	-22.3%	[-43.0, 5.8]	0.109
Household vs corporate-developer units	-5.9%	-20.1%	-15.0%	[-33.7, 8.9]	0.199

Notes: Effects are percentage changes. The difference is the more-costly coefficient minus the less-costly coefficient, transformed to a percentage. Its 95 percent confidence interval and p-value test equality of the two effects directly. Each row estimates both outcomes jointly with outcome-specific episode and state-by-calendar-year fixed effects. Standard errors are clustered by municipality.

Figure 3 also reports permits for units developed by public bodies and non-profit organizations. These units account for 1.9 percent of all permitted housing, and 95.5 percent of municipality-years record none. The fixed-effects Poisson model therefore retains 3,773 episode-years across the 396 episodes that ever record a public or non-profit permit. The point estimate implies a 59.9 percent decline after a sharp loss of support, with a 95 percent confidence interval from -84.1 to 1.0 percent. The direction is consistent with budget trade-offs creating an expected support penalty for public housing (Hilbig and Wiedemann, 2026). The developer category does not identify subsidized housing, so the estimate cannot establish an effect on German *Sozialwohnungen*.

C.6 Direction and dose

Figure C.4 re-estimates Equation 3 after defining the treated group as either a vote-share fall or increase greater than ten points. The decrease estimates reproduce Table 1. All increase estimates are positive but imprecise.

The disjoint dose bands in Figure C.5 show that the ten-to-fifteen-point drop band produces the largest negative multifamily estimate. The single-family estimate is positive in the five-to-ten-point loss band. In the opposite direction, increases greater than fifteen points produce the largest positive total and multifamily point estimates. The largest-change bands are noisy because relatively few episodes remain and because very large changes combine heterogeneous electoral shocks. The figure is a diagnostic of treatment magnitude, not a claim of linear dose response.

Figure C.6 varies the threshold used to define a vote-share fall from five to twenty percentage points while retaining the stable-support comparison group. Total and multifamily point estimates are negative at every threshold. The single-family estimate is slightly positive at five points and negative thereafter. The main ten-point threshold is the only definition at which the total and multifamily estimates reach the five-percent level. At higher thresholds, the treated sample falls from 262 episodes at ten points to 52 at twenty points.

C.7 Continuous fixed-effects evidence

For each permit outcome h, I estimate

$$\mathbb{E}[Y_{iht} \mid S_{it}, \alpha_i, \lambda_t] = P_{it} \exp(\beta_h S_{it} + \alpha_i + \lambda_t)\,, \tag{6}$$

Figure C.4: Vote-share falls and increases

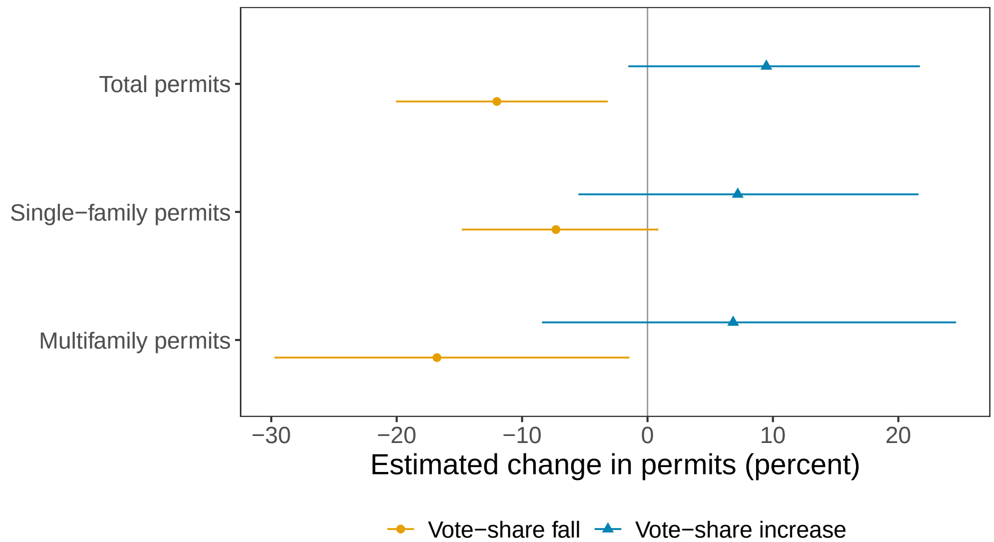

Notes: Estimates of τ_h from Equation 3, replacing the treated category in Equation 2 with a greater-than-ten-point fall or increase. Each is compared with the same stable group. Bars are 95 percent confidence intervals.

Figure C.5: Direction and magnitude of the electoral change

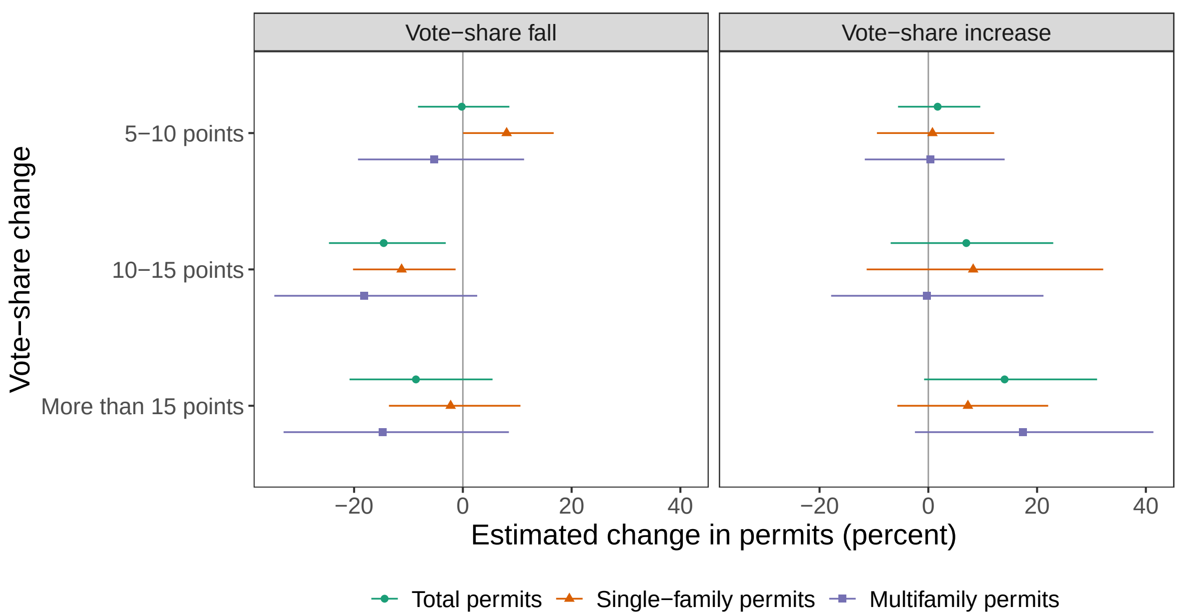

Notes: Estimates from Equation 3, replacing D_e in Equation 2 with mutually exclusive vote-share-change bands. Each band is compared with stable-support episodes, in which vote share changes by less than five percentage points in either direction. Bars are 95 percent confidence intervals.

Figure C.6: Sensitivity to the vote-share-loss threshold

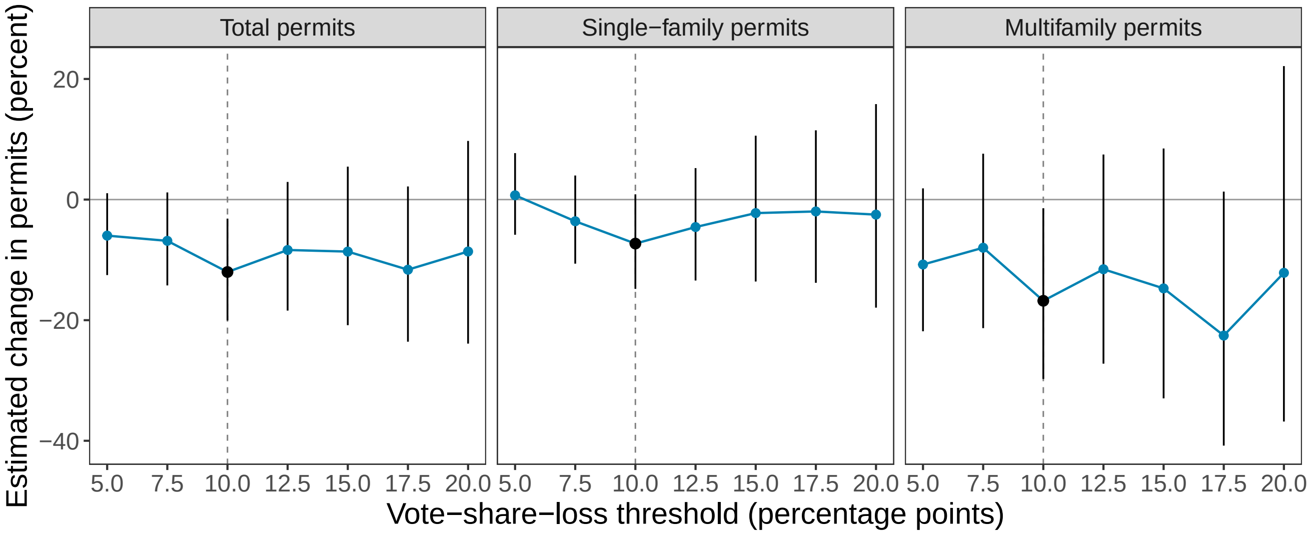

Notes: Estimates from Equation 3 as the minimum vote-share fall defining treatment varies from five to twenty percentage points. Stable-support episodes continue to require a change smaller than five points in absolute value. Bars are 95 percent confidence intervals. The dashed line and black point identify the ten-point definition in Equation 2.

where Y_{iht} is the permit count in municipality i and year t, P_{it} is the population exposure, α_i denotes municipality fixed effects, and λ_t denotes year fixed effects. The electoral-safety measure S_{it} is either the incumbent mayor's first-round vote share or the vote-share gap between the two largest named parties in the most recent council election. Each measure is carried forward until the next election. The coefficient β_h therefore uses changes within municipalities after absorbing common year shocks.

Table C.4 reports effects for a ten-percentage-point increase in S_{it}. The council winning margin is the main continuous measure because councils formally adopt local land-use plans and the margin captures changes in the leading party's institutional position. A wider council margin predicts more permits for all three outcomes, with the largest estimate for multifamily housing. The equivalent mayoral vote-share specification produces the same ordering. Because Equation 6 allows electoral safety to move in either direction and uses all available changes, it provides supporting fixed-effects evidence rather than the stacked difference-in-differences comparison in Equation 3.

Two limitations reinforce that supporting role. The carried-forward electoral measure is not strictly exogenous if permitting during one electoral spell affects the next election result. In addition, a continuous two-way fixed-effects coefficient can combine heterogeneous responses with implicit weights that are difficult to interpret as one common effect (de Chaisemartin and D'Haultfœuille, 2020). I therefore use these models to assess whether the broader within-municipality association points in the same direction, not as a substitute for the discrete mayoral design.

Table C.4: Supporting fixed-effects evidence

Design	Outcome	Change	95% CI	p	N
Council winning margin	Total permits	4.9%	[1.1, 8.8]	0.011	206,977
Council winning margin	Single-family permits	2.4%	[1.2, 3.7]	0.000	207,036
Council winning margin	Multifamily permits	7.1%	[1.9, 12.5]	0.007	205,898
Mayor vote share	Total permits	2.1%	[0.9, 3.2]	0.001	111,571
Mayor vote share	Single-family permits	1.2%	[0.7, 1.8]	0.000	111,657
Mayor vote share	Multifamily permits	3.4%	[0.9, 5.9]	0.007	110,909

Notes: Poisson estimates of Equation 6. Mayor safety is the incumbent's first-round vote share. Council safety is the vote-share gap between the two largest named parties. Electoral values from the most recent election are carried forward until the next election. Models include municipality and year fixed effects and a log-population offset. Standard errors are clustered by municipality. Effects are reported for a ten-percentage-point increase.

C.8 Pretreatment heterogeneity

The mayor's starting position merits a flexible test because a linear interaction could miss a threshold or nonmonotonic pattern. I divide the 262 vote-loss episodes into quintiles of first-round vote share at the preceding election, with 52 or 53 treated episodes in each group. I apply the same cut points to stable-support episodes and allow both their post-election path and the treatment effect to vary across quintiles. Figure C.7 shows no systematic gradient. Joint tests of equal treatment effects across quintiles produce $p = .856$ for total, $p = .710$ for single-family, and $p = .618$ for multifamily permits. The average decline is therefore not concentrated among mayors who began just above one particular support threshold.

I next examine whether the effects of vote-share falls and increases vary with a broader set of political and municipal conditions already in place before the focal re-election. The moderators capture population size and density, council partisanship, prior citizen initiatives, the mayor's party family, homeownership, vacancy, rents, rent growth, migration, and municipal finances.

For each moderator M_e, I augment Equation 3 with $M_e \times \text{Post}_{et}$ and $D_e \times \text{Post}_{et} \times M_e$. The first term allows stable-support episodes with different pretreatment characteristics to experience different post-election changes. The triple interaction tests whether the treatment effect varies with the moderator. Figure C.8 reports conditional effects at one standard deviation below and above the mean for continuous moderators. Figure C.9 separately compares municipalities by prior citizen initiatives and the mayor's party family. The party specification includes separate family indicators and tests their treatment interactions jointly. Every measure is fixed before the focal election. Following Fu and Slough (2026), I treat these estimates as descriptive. Interpreting the tax-revenue interaction as evidence for a fiscal mechanism would require revenue not to alter other pathways from support losses to permits, even though revenue may also affect administrative capacity, public investment, and housing demand.

Several interactions cross conventional thresholds before accounting for the number of comparisons. Only one survives a Benjamini–Hochberg correction across all 78 tests. After a vote-share fall, multifamily permitting declines by 38.9 percent at one standard deviation below mean tax revenue. At one standard deviation above mean tax revenue, the point estimate is a statisti-

Figure C.7: The vote-share loss across starting positions

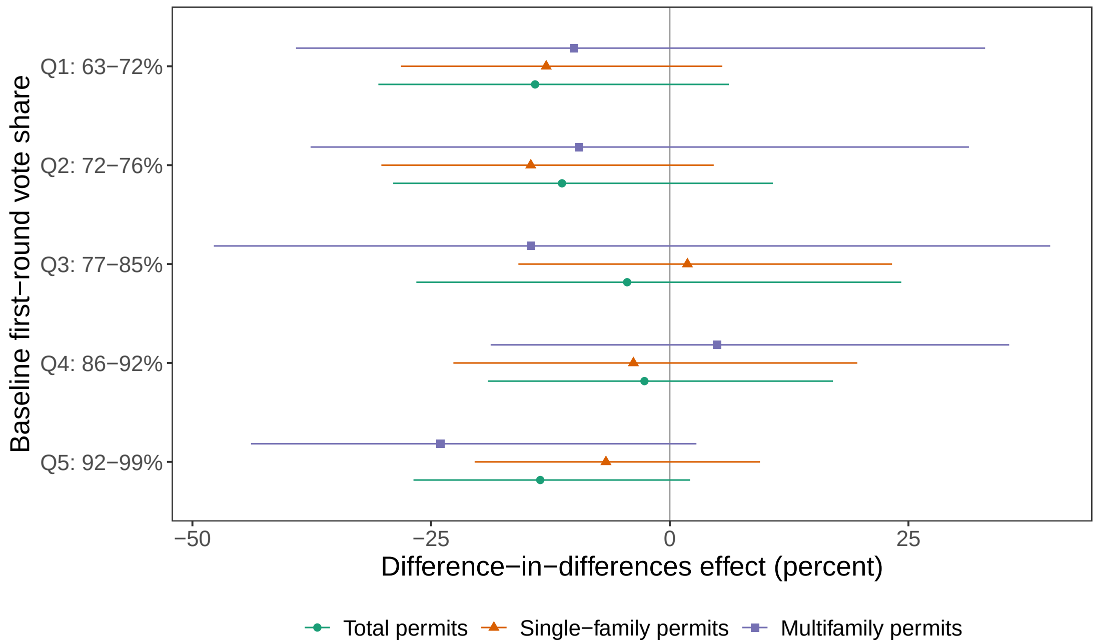

Notes: Conditional estimates from Equation 3 by quintile of the treated episodes' vote share at the preceding election. The cut points are applied to stable-support episodes. The models include quintile-by-post and treatment-by-post-by-quintile interactions while retaining the baseline-support and pre-election-density adjustments. Each quintile contains 52 or 53 treated vote-loss episodes. Bars are 95 percent confidence intervals.

cally imprecise 15.1 percent increase. The interaction has $p < .001$ and $q = .014$. The result implies that the negative multifamily response is concentrated among municipalities with fewer fiscal resources and is consistent with fiscal capacity buffering the permitting response to electoral risk. It does not establish that fiscal capacity is the operative mechanism. The tax sample contains only 56 treated and 209 stable-support episodes, and revenue could moderate several pathways. I therefore treat this pattern as suggestive rather than as a separate main result. Neither the citizen-initiative tests nor the joint party-family tests survive the correction. None of the interactions for vote-share increases does so either.

C.9 Common support, inference, and episode construction

The stacked design is concentrated in Bavaria, which contributes 192 of 262 treated episodes and 937 of 1,209 stable-support episodes. The main text therefore also reports less restrictive fixed-effects estimates from broader mayoral and council panels.

Stable-support episodes begin from higher vote shares than treated episodes. Their median baseline share is 0.90, compared with 0.81 among treated episodes. About 80 percent of stable-support episodes nevertheless lie within the first-to-ninety-ninth-percentile range of treated baseline support. Table C.5 addresses the remaining imbalance in four ways. Restricting the comparison group to the treated range and trimming both groups to the treated central range leave the estimates nearly unchanged. A three-degree-of-freedom spline in baseline support attenuates the effects and widens the intervals. ATT reweighting produces estimates between the trimmed

Figure C.8: Pretreatment heterogeneity across municipal conditions

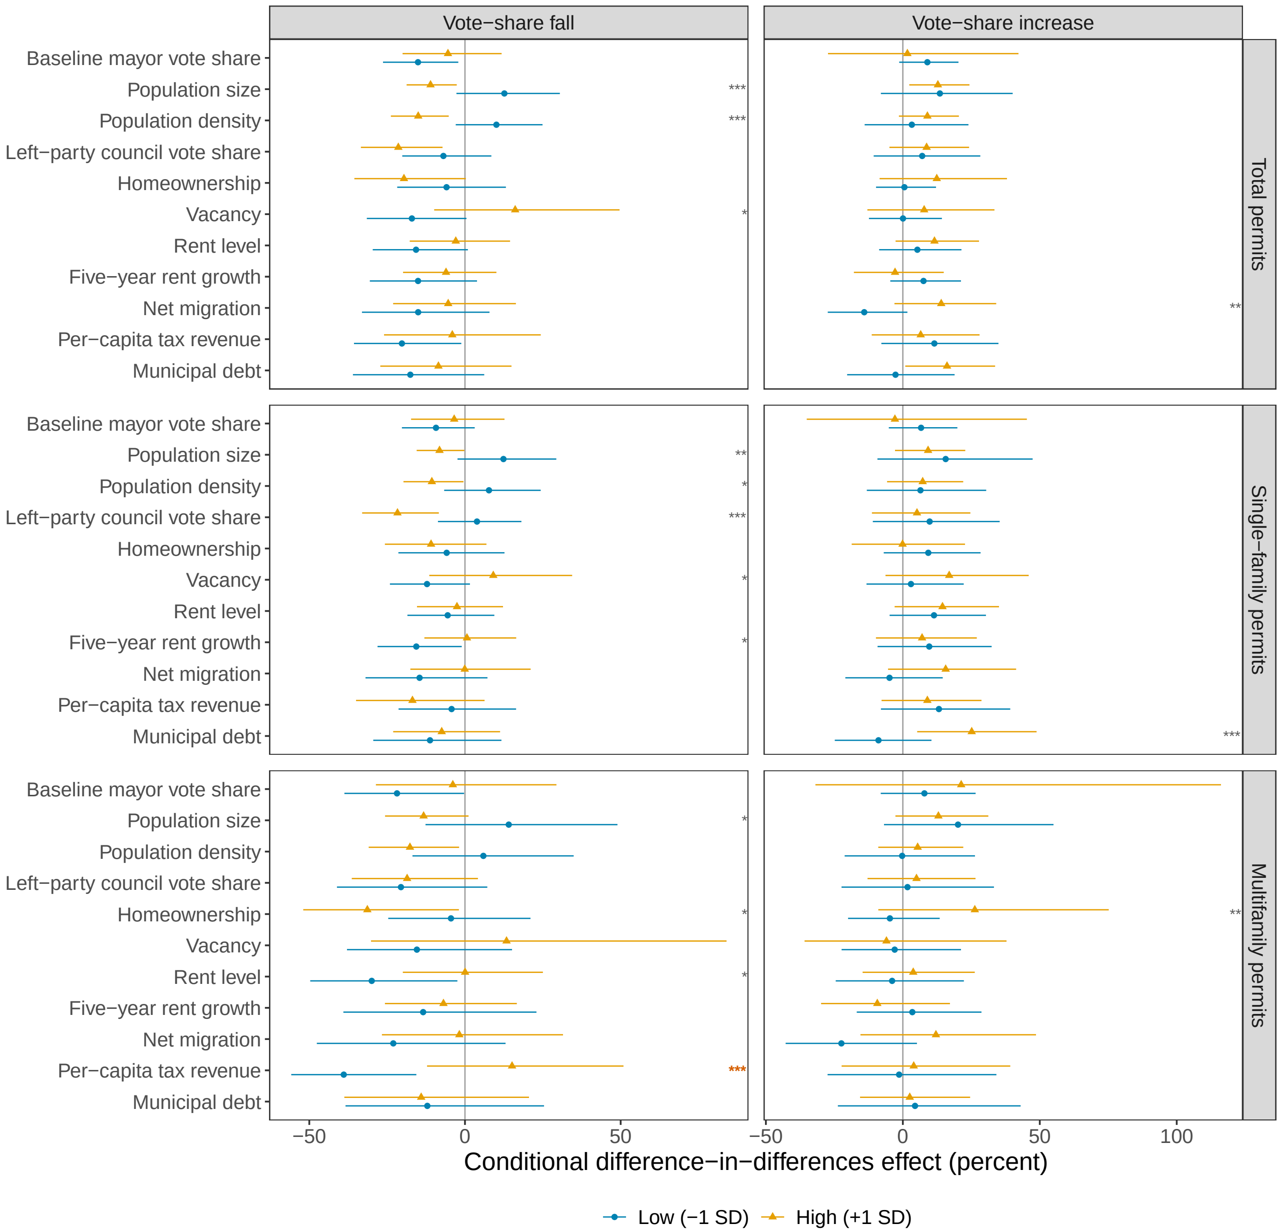

Notes: Conditional estimates from Equation 3 augmented with moderator-by-post and treatment-by-post-by-moderator interactions. Points compare moderator values one standard deviation below and above the mean. Bars are 95 percent confidence intervals. Stars at the right report the nominal triple-interaction p-value, with $^{*}p < .1$, $^{**}p < .05$, and $^{***}p < .01$. The orange star identifies the interaction that survives the Benjamini–Hochberg correction across all 78 tests shown in Figures C.8 and C.9.

Figure C.9: Pretreatment heterogeneity by citizen initiatives and mayoral party

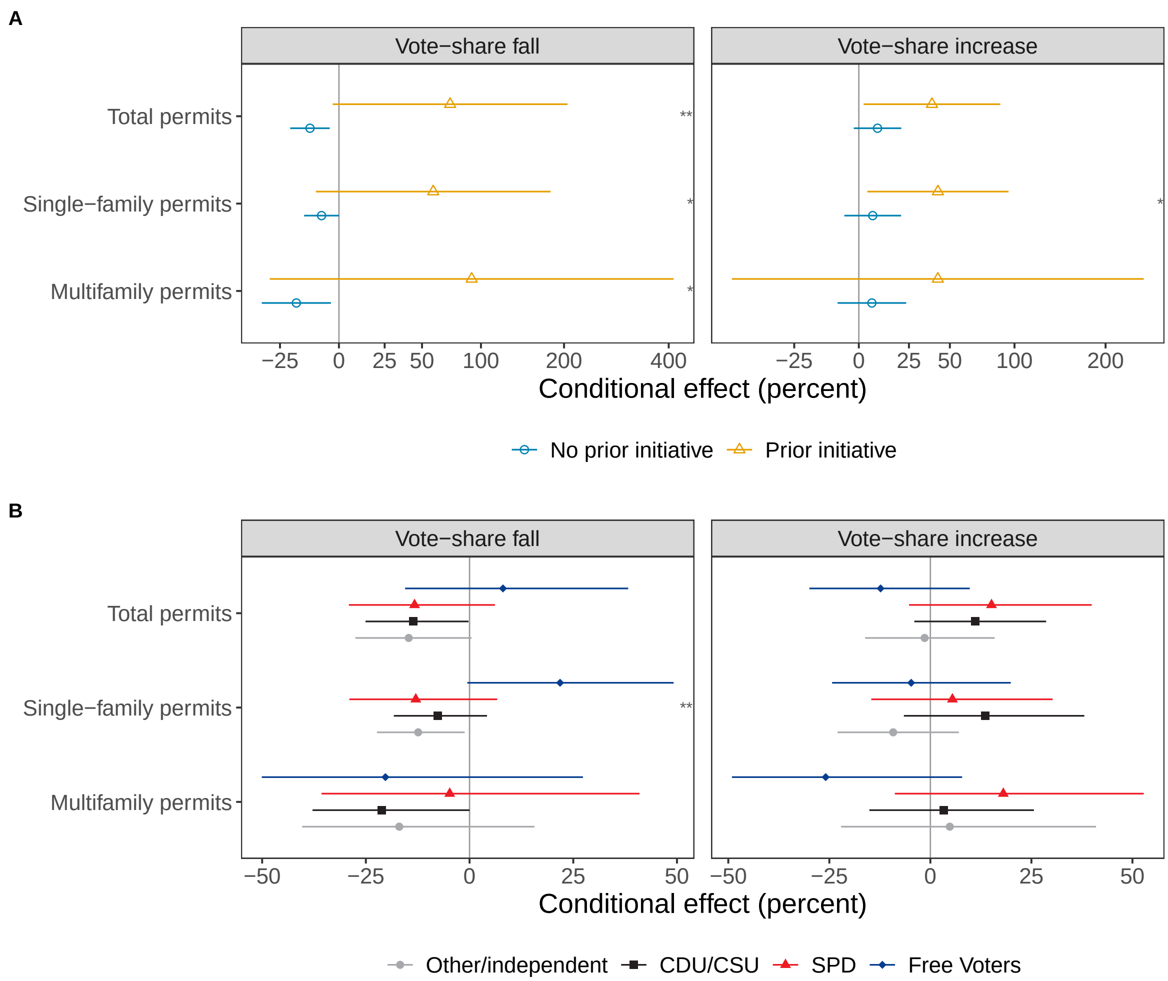

Notes: Panel A compares conditional estimates by the prior absence or presence of a housing-related citizen initiative. Only three treated vote-loss episodes, five treated vote-increase episodes, and ten stable-support episodes have a prior initiative. Panel B compares other or independent mayors with CDU/CSU, SPD, and Free Voter mayors. Bars are 95 percent confidence intervals. Panel A uses the underlying Poisson log scale, labeled as percentage changes, to display its wide intervals without truncation. Stars at the right report nominal moderator-test p-values, with $^{*}p < .1$, $^{**}p < .05$, and $^{***}p < .01$. The party-family test is joint across categories. None of these results survives the Benjamini–Hochberg correction across all 78 tests.

and flexible-adjustment results. These checks show that extreme stable-support episodes do not drive the result, while the weaker flexible estimates counsel against treating the exact magnitude as fixed.

Table C.5: Sensitivity to baseline-support overlap

Specification	Total	Single-family	Multifamily
Baseline	-12.0^{***} [-20.1, -3.2]	-7.3^{*} [-14.8, 0.9]	-16.8^{**} [-29.7, -1.4]
Stable episodes within treated range	-12.5^{***} [-19.9, -4.4]	-7.5^{*} [-14.9, 0.6]	-15.3^{**} [-27.4, -1.2]
Both groups within treated p1-p99	-12.3^{***} [-19.8, -4.2]	-7.3^{*} [-14.8, 0.8]	-15.3^{**} [-27.5, -1.0]
Flexible baseline-support adjustment	-9.5^{*} [-18.3, 0.1]	-6.7 [-14.7, 2.1]	-10.9 [-25.6, 6.8]
ATT reweighting	-10.8^{**} [-18.6, -2.2]	-7.0^{*} [-14.7, 1.3]	-14.4^{*} [-27.3, 0.9]

Notes: Cells report percentage effects with 95 percent confidence intervals. The second row retains stable-support episodes within the full treated range. The third retains both groups within the first and ninety-ninth percentiles of treated baseline support. The fourth replaces the linear baseline-support interaction with a natural cubic spline. The fifth weights stable-support episodes by the odds of treatment after estimating treatment from baseline support, pre-election density, state, and focal year. Weights are capped at the comparison-group ninety-ninth percentile. Standard errors are clustered by municipality. $^{*}p < .1$, $^{**}p < .05$, $^{***}p < .01$.

Table C.6 varies the dependence structure used for inference. Clustering by county permits arbitrary dependence across municipalities in the same county. Two-way clustering by municipality and state-year additionally allows common disturbances within a state-year beyond the fixed-effect adjustment. Both produce standard errors close to the municipality-clustered baseline.

Table C.6: Alternative clustering of standard errors

Inference	Total	Single-family	Multifamily
Municipality clusters	-12.0^{***} [-20.1, -3.2]	-7.3^{*} [-14.8, 0.9]	-16.8^{**} [-29.7, -1.4]
County clusters	-12.0^{***} [-19.7, -3.6]	-7.3^{*} [-14.9, 1.0]	-16.8^{**} [-29.8, -1.4]
Municipality and state-year clusters	-12.0^{***} [-20.2, -3.0]	-7.3^{*} [-15.2, 1.3]	-16.8^{**} [-30.2, -0.8]

Notes: Cells report percentage effects with 95 percent confidence intervals from Equation 3. Only the variance estimator changes across rows. $^{*}p < .1$, $^{**}p < .05$, $^{***}p < .01$.

Table C.7 addresses three further specification and sample concerns. First, county-by-calendar-year fixed effects replace the state-by-calendar-year fixed effects in Equation 3. They absorb shocks shared by municipalities in the same county and year, including more localized changes in housing demand. The multifamily result is essentially unchanged, the total-permit result remains negative but attenuates, and the single-family estimate moves close to zero.

Second, the main specification enters current population as an exposure offset. This choice preserves zero permit counts and gives the Poisson coefficient a rate interpretation, but population could itself change after the election. Omitting the offset leaves the sign, magnitude, and precision of all three estimates virtually unchanged. A frozen pre-election population offset produces exactly the same slope estimates as the no-offset model because it is constant within an episode and is absorbed by the episode fixed effect.

Third, I retain only the earliest eligible episode in each municipality. This restriction prevents municipalities with multiple qualifying re-elections from contributing more than once. It leaves 217 treated and 966 stable-support episodes before outcome-specific separation. All three estimates retain their direction and remain close to the baseline magnitudes. Precision declines in the smaller sample, but repeated episodes do not drive the pattern.

Every specification adjusts for density measured before the focal re-election and interacts it with the post-election period. This timing prevents post-election population or permitting changes from entering density as a contemporaneous control. Density is observed in the year immediately before re-election for 1,465 of the 1,471 episodes. For the remaining six, I use the closest observed pre-election year.

Table C.7: Specification and sample robustness

Specification	Total	Single-family	Multifamily
Current-population offset	-12.0^{***} [-20.1, -3.2]	-7.3^{*} [-14.8, 0.9]	-16.8^{**} [-29.7, -1.4]
County-by-year fixed effects	-7.6^{*} [-15.2, 0.7]	-0.5 [-8.2, 7.8]	-17.1^{**} [-29.0, -3.2]
No population offset	-12.1^{***} [-20.0, -3.3]	-7.5^{*} [-14.9, 0.6]	-16.7^{**} [-29.6, -1.5]
First episode per municipality	-11.8^{**} [-20.7, -1.9]	-7.4 [-15.8, 1.8]	-15.1^{*} [-29.3, 2.0]

Notes: Cells report $100[\exp(\tau_h) - 1]$ with 95 percent confidence intervals in brackets. The first row reproduces Equation 3. The second replaces state-by-calendar-year with county-by-calendar-year fixed effects. The third omits the population offset. The fourth retains the earliest eligible episode in each municipality. All models include episode fixed effects and the post-election interactions in Equation 3. Standard errors are clustered by municipality. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

C.10 Sensitivity to violations of parallel trends

I also use the relative-magnitude sensitivity analysis of Rambachan and Roth (2023) to ask how the dynamic estimates change when post-election violations of parallel trends may be larger than the observed pre-election movements. Table C.8 applies the procedure to the equal-weight average of event years one through five. This estimand differs from the pooled coefficient in Table 1. Even its conventional 95 percent intervals narrowly include zero for total and multifamily permits. Allowing a post-election violation half as large as the largest pre-election movement makes the bounds much wider. The annual event-study estimates are therefore too noisy to deliver informative robust bounds against modest violations of parallel trends.

Table C.8: Sensitivity of the average dynamic effect to trend violations

Outcome	Restriction	Estimate	95% CI
Total permits	Conventional	-11.4%	[-22.0, 0.5]
Total permits	$\bar{M} = 0.5$	-11.4%	[-40.0, 29.4]
Single-family permits	Conventional	-5.1%	[-14.3, 5.2]
Single-family permits	$\bar{M} = 0.5$	-5.1%	[-33.7, 35.5]
Multifamily permits	Conventional	-18.4%	[-35.5, 3.2]
Multifamily permits	$\bar{M} = 0.5$	-18.4%	[-62.7, 66.0]

Notes: The estimate is the equal-weight average of the treatment-by-event-time coefficients for years $t = 1$ through $t = 5$ in Equation 5. The conventional row imposes exact parallel trends. Relative-magnitude rows allow the largest post-election violation to be $\bar{M}$ times the largest pre-election movement. Bounds use four pre-election coefficients and 95 percent coverage.

C.11 Doubly robust difference-in-differences estimates

The Callaway–Sant'Anna estimator provides a deliberately different translation of the main comparison (Callaway and Sant'Anna, 2021). The unit is an eligible episode aligned around the focal re-election. Vote-loss episodes become treated in year $t = +1$, while stable-support episodes form the never-treated comparison group. Because the estimator models an additive change in a continuous outcome, the outcome is the number of permits per 1,000 residents rather than the raw count used by the Poisson model (Chen and Roth, 2024). The estimator combines an outcome regression with inverse-probability weighting, so the group-time average treatment effects are consistent if either adjustment model is correctly specified. The propensity and outcome models use baseline vote share, focal election year, pre-election log population and density, and the pre-election outcome mean. Standard errors use a municipality-clustered multiplier bootstrap, which allows episodes from the same municipality to remain dependent.

Annual municipal permit counts are highly variable, so Table C.9 reports several pooled summaries. The all-years estimate averages every available year in the adjacent terms. The remaining estimates compare symmetric means over two, three, four, or five years on either side of the focal re-election. Each fixed-window estimate requires complete observations throughout that window. The two-year sample also requires observations in years $t = -4$ and $t = -3$ so that its pre-election trajectory can be tested over four years. Longer balanced windows therefore use fewer episodes. Reporting the complete set avoids making the conclusion depend on one aggregation horizon.

Every aggregation produces negative estimates for all three outcomes. The two- and three-year estimates are largest, especially for multifamily permits. Four- and five-year estimates are smaller as the balanced sample contracts. The all-years estimates are roughly half the corresponding Poisson magnitudes. None of the fifteen estimates reaches the five-percent level, with the smallest p-value equal to .072. None of the row-specific joint tests rejects parallel pre-election trends, although the shrinking treated samples also limit their power.

Table C.9: Callaway–Sant'Anna estimates across aggregation horizons

Aggregation	Outcome	ATT	95% CI	% of mean	p	Pretrend p	Treated
All available years	Total permits	-0.270	[-0.572, 0.032]	-5.5%	0.080	0.623	261
All available years	Single-family permits	-0.137	[-0.360, 0.086]	-4.2%	0.227	0.547	261
All available years	Multifamily permits	-0.132	[-0.313, 0.049]	-8.3%	0.152	0.875	261
2-year means	Total permits	-0.373	[-0.786, 0.040]	-9.6%	0.077	0.904	207
2-year means	Single-family permits	-0.126	[-0.434, 0.183]	-4.8%	0.424	0.974	207
2-year means	Multifamily permits	-0.267	[-0.558, 0.024]	-20.8%	0.072	0.884	207
3-year means	Total permits	-0.304	[-0.675, 0.066]	-7.5%	0.107	0.725	193
3-year means	Single-family permits	-0.132	[-0.426, 0.161]	-4.8%	0.377	0.932	193
3-year means	Multifamily permits	-0.192	[-0.419, 0.035]	-14.8%	0.098	0.765	193
4-year means	Total permits	-0.178	[-0.530, 0.174]	-4.3%	0.321	0.778	169
4-year means	Single-family permits	-0.094	[-0.359, 0.171]	-3.3%	0.488	0.996	169
4-year means	Multifamily permits	-0.085	[-0.280, 0.111]	-6.8%	0.395	0.424	169
5-year means	Total permits	-0.169	[-0.493, 0.154]	-4.0%	0.306	0.440	161
5-year means	Single-family permits	-0.062	[-0.317, 0.192]	-2.1%	0.631	0.819	161
5-year means	Multifamily permits	-0.114	[-0.289, 0.062]	-9.2%	0.205	0.624	161

Notes: Doubly robust Callaway–Sant'Anna estimates in permits per 1,000 residents. The percentage column divides the ATT by the treated pre-period mean. Fixed k-year rows compare means over $t = -k, \ldots, -1$ and $t = +1, \ldots, +k$. All-years rows use every available year in the adjacent terms. Pretrend p-values jointly test the pre-election treatment-by-event-time coefficients in a linear event-time regression of permit rates on the same episodes. Standard errors for the ATT come from a municipality-clustered multiplier bootstrap. Pretrend tests cluster by municipality.

C.12 Assessing alternative explanations

Table C.10 applies four sets of adjustments to the main stacked difference-in-differences. All covariates are fixed before the focal re-election and interacted with the post-election period. Each adjusted row follows an unadjusted estimate on exactly the same episodes. The comparison therefore distinguishes changes caused by covariate adjustment from changes caused by data coverage.

The anti-development preference checks use an indicator for a housing-related citizen initiative during the preceding five years and, in a separate census sample, the latest homeownership rate observed before the focal election. The citizen-initiative model is estimated in the council-data sample and also includes the partisan adjustments described next. The estimates scarcely move after either preference adjustment.

The partisanship check allows the post-election change to vary with the mayor's party family and the combined SPD, Green, and Left vote share in the most recent council election no more than six years before the focal election. These adjustments also leave the estimates stable across all three outcomes.

The fiscal check uses per-capita tax revenue and municipal debt in the year before re-election. I remove administrative sentinel values, winsorize both measures at the first and ninety-ninth percentiles, and standardize them across eligible episodes. The demand check uses the pre-election rent level, rent growth during the preceding five years, and net migration. Fiscal adjustment leaves the estimates unchanged. Within the same 616-episode demand sample, the demand adjustment attenuates the total, single-family, and multifamily estimates by 3.3, 3.3, and 2.5 percentage points. Every estimate remains negative.

Table C.10: Pre-election adjustments in the main mayor design

Specification	Total	Single-family	Multifamily	Episodes	Treated
Main specification	-12.0^{***} [-20.1, -3.2]	-7.3^{*} [-14.8, 0.9]	-16.8^{**} [-29.7, -1.4]	1,471	262
Census benchmark	-10.3 [-23.2, 4.7]	-7.5 [-18.5, 4.9]	-9.9 [-29.3, 14.7]	626	111
+ Homeownership	-9.4 [-21.6, 4.6]	-7.5 [-18.4, 4.8]	-9.2 [-28.0, 14.5]	626	111
Council-data benchmark	-15.0^{***} [-24.3, -4.5]	-10.7^{**} [-19.4, -1.2]	-18.4^{**} [-33.0, -0.6]	807	149
+ Partisanship	-15.0^{***} [-24.3, -4.6]	-10.7^{**} [-19.2, -1.3]	-19.5^{**} [-34.1, -1.8]	807	149
+ Partisanship and opposition	-14.9^{***} [-24.2, -4.4]	-10.7^{**} [-19.2, -1.2]	-19.3^{**} [-33.8, -1.5]	807	149
Fiscal-sample benchmark	-13.3 [-27.2, 3.4]	-10.0 [-22.8, 4.8]	-15.1 [-34.9, 10.7]	264	56
+ Tax revenue and debt	-13.0 [-27.0, 3.6]	-9.8 [-22.6, 5.0]	-15.1 [-34.9, 10.6]	264	56
Demand-sample benchmark	-11.0 [-23.8, 4.0]	-8.7 [-19.4, 3.5]	-9.8 [-29.3, 15.1]	616	110
+ Rents, rent growth, and migration	-7.7 [-20.3, 6.8]	-5.4 [-16.1, 6.6]	-7.3 [-26.6, 17.0]	616	110

Notes: Entries are percentage effects, $100 \times [\exp(\widehat{\tau}_h) - 1]$, with 95 percent confidence intervals in brackets. Every row estimates Equation 3 using the treatment in Equation 2. Covariates are measured before re-election and interacted with the post-election period. Benchmark and adjusted rows within each block use identical episodes. Comparisons within blocks, rather than significance across samples, identify the contribution of the adjustment. Standard errors are clustered by municipality. $^{*}p < .1$, $^{**}p < .05$, $^{***}p < .01$.

Concurrent council elections. Council elections coincide with 81 percent of focal mayoral elections, reflecting the electoral calendars of the German states. Concurrency is almost identical in vote-share-fall and stable-support episodes (80.2 and 81.1 percent, $p = .79$), and the permit response does not vary systematically with concurrency ($p = .34$ for total and $p = .44$ for multifamily permits). Among the 52 treated episodes with nonconcurrent elections, total permits fall by 4.3 percent (95 percent CI: $-17.2, 10.5$). This subsample has a minimum detectable effect of 18.6 percent at 80 percent power, and its confidence interval includes the pooled estimate. I therefore interpret the design as identifying the response of local political leadership under the same mayor, not the mayor acting independently of her council.

C.13 Predetermined development-industry strength

The developer-capture concern implies that the permitting response should be larger where organized development interests were initially stronger. I proxy their local economic presence with municipality-level employment shares in real estate and construction from the Federal Employment Agency. The primary measures use the latest observation before the earlier election in each two-term episode, so they precede both the focal vote-share change and the post-election permitting period. I standardize each share and interact it with treatment and the post-election indicator in Equation 3. A negative triple interaction would indicate a larger permit decline where the industry was stronger.

Table C.11 provides no such pattern. The interaction estimates are not systematically negative, and none of the six primary comparisons survives a Benjamini–Hochberg correction. A secondary version measured before the focal re-election produces the same general conclusion. Employment shares are an imperfect proxy for political organization and cannot capture firm-specific lobbying or informal relationships. The test therefore shows that the main result is not concentrated in municipalities with a larger predetermined development industry, rather than ruling out capture in every form.

Table C.11: Effects by predetermined development-industry strength

Timing	Moderator	Total		Single-family		Multifamily		Episodes
		Coef.	(SE)	Coef.	(SE)	Coef.	(SE)	
Before baseline election	Real-estate employment	0.051	(0.060)	−0.041	(0.059)	0.166*	(0.099)	599
Before baseline election	Construction employment	0.124	(0.078)	0.043	(0.067)	0.171	(0.135)	599
Before focal re-election	Real-estate employment	−0.018	(0.055)	−0.091*	(0.049)	0.098	(0.088)	888
Before focal re-election	Construction employment	0.085	(0.068)	0.055	(0.058)	0.075	(0.113)	888

Notes: Entries are triple-interaction coefficients, with clustered standard errors in parentheses. The moderators are municipal employment shares standardized across eligible episodes. A negative coefficient indicates that the permit decline in the main within-mayor specification becomes larger as industry employment increases. Every model follows Equation 3 and also allows the moderator to predict a different post-election change. Standard errors are clustered by municipality. Stars report unadjusted tests. No interaction survives a Benjamini–Hochberg correction. $^{*}p < .1$, $^{**}p < .05$, $^{***}p < .01$.

Because treated and stable-support episodes both follow re-elected incumbents, an administrative cycle common to all new terms is absorbed by the post-election coefficient. Figure C.10 extends the permitting history to address reversion from an unusually active baseline term. It finds no statistically detectable differential increase, although the intervals remain wide. The absence of a treatment gradient across baseline vote-share quintiles provides a separate check against reversion from one unusually strong electoral starting point.

Table C.12 repeats the same adjustment logic in the larger panels used for the continuous mayor vote-share and council-margin models. Covariates are measured before the election that sets each electoral spell and then carried through that spell. Council-margin coefficients move little within the paired samples. The weaker precision for total and multifamily permits in some restricted samples is already present in the corresponding benchmark rows. These models remain corroborative because carried electoral measures can respond to prior permitting and continuous fixed-effects coefficients may combine heterogeneous responses.

C.14 Permitting before the focal re-election

A simple backlash-and-reversion account requires treated municipalities to permit unusually heavily before the focal re-election. I first ask whether permitting during the baseline term predicts which mayors subsequently experience a sharp loss of support. For each episode, I measure the

Table C.12: Pre-election adjustments in continuous fixed-effects models

Specification	Total	Single-family	Multifamily	N	Municipalities
Mayor vote share					
Main specification	2.1*** [0.9, 3.2]	1.2*** [0.7, 1.8]	3.4*** [0.9, 5.9]	110,909	4,538
Council-data benchmark	2.1*** [0.9, 3.3]	1.4*** [0.7, 2.1]	3.2*** [0.9, 5.6]	63,785	3,204
+ Partisanship	1.9*** [0.8, 3.0]	1.4*** [0.7, 2.1]	3.1** [0.6, 5.6]	63,785	3,204
+ Partisanship and opposition	1.9*** [0.8, 3.1]	1.4*** [0.7, 2.1]	3.1** [0.6, 5.7]	63,785	3,204
Census benchmark	1.9** [0.2, 3.7]	2.0*** [1.0, 3.0]	2.3 [-0.7, 5.4]	37,221	4,314
+ Homeownership	2.0** [0.2, 3.7]	2.0*** [1.0, 3.0]	2.3 [-0.7, 5.5]	37,221	4,314
Fiscal-sample benchmark	1.6** [0.2, 3.0]	2.3*** [1.2, 3.4]	2.3 [-0.5, 5.2]	25,717	1,992
+ Tax revenue and debt	1.4** [0.0, 2.7]	2.3*** [1.2, 3.4]	1.5 [-0.9, 4.0]	25,717	1,992
Demand-sample benchmark	2.4*** [0.7, 4.1]	2.0*** [1.0, 3.0]	3.4** [0.7, 6.1]	35,580	4,298
+ Rents, rent growth, and migration	2.0** [0.4, 3.6]	1.9*** [0.9, 2.9]	2.6** [0.1, 5.2]	35,580	4,298
Council winning margin					
Main specification	4.9** [1.1, 8.8]	2.4*** [1.2, 3.7]	7.1*** [1.9, 12.5]	205,898	7,388
Council-data benchmark	3.2 [-1.3, 8.0]	3.0*** [1.6, 4.4]	3.6 [-1.7, 9.2]	166,911	7,182
+ Partisanship	3.3 [-0.8, 7.6]	2.8*** [1.5, 4.1]	3.9 [-1.0, 9.1]	166,911	7,182
+ Partisanship and opposition	3.2* [-0.4, 7.0]	2.8*** [1.5, 4.1]	3.8* [-0.6, 8.4]	166,911	7,182
Census benchmark	0.9 [-1.7, 3.6]	2.5*** [0.7, 4.4]	−0.4 [-4.5, 4.0]	65,349	6,801
+ Homeownership	1.0 [-1.6, 3.7]	2.6*** [0.7, 4.5]	−0.3 [-4.5, 4.0]	65,349	6,801
Fiscal-sample benchmark	1.4 [-1.4, 4.3]	3.5*** [1.3, 5.8]	1.3 [-3.2, 6.0]	37,721	2,709
+ Tax revenue and debt	1.2 [-1.7, 4.2]	3.6*** [1.3, 5.9]	0.6 [-3.6, 5.1]	37,721	2,709
Demand-sample benchmark	2.3 [-1.1, 5.9]	3.0*** [1.1, 5.0]	1.7 [-3.9, 7.6]	70,416	7,349
+ Rents, rent growth, and migration	3.0** [0.1, 5.9]	2.8*** [1.0, 4.7]	3.1 [-1.4, 7.7]	70,416	7,349

Notes: Entries are percentage changes associated with a ten-percentage-point increase in mayor vote share or the council winning margin, with 95 percent confidence intervals in brackets. All models include municipality and year fixed effects and a log-population offset. Covariates are fixed before the election that sets the carried treatment. Benchmark and adjusted rows within each block use identical municipality-years. Standard errors are clustered by municipality. $^{*}p < .1$, $^{**}p < .05$, $^{***}p < .01$.

mean and annual trend in $\log(1 + \text{permits per capita})$ between the baseline and focal elections. Table C.13 regresses the sharp-loss indicator on these standardized measures.

Table C.13: Baseline-term permitting and a subsequent loss of support

Permit outcome	Baseline-term mean (percentage points)	Annual trend (percentage points)
Total permits	−1.5 [−4.9, 1.8]	1.4 [−1.1, 3.9]
Single-family permits	−1.8 [−5.0, 1.5]	−0.1 [−2.6, 2.5]
Multifamily permits	0.0 [−3.2, 3.3]	1.0 [−1.5, 3.5]
Episodes	1,204	

Notes: Cells report the percentage-point change in the probability of a sharp support loss associated with a one-standard-deviation increase in baseline-term permitting. Each row comes from a model that enters the mean and annual trend in $\log(1 + \text{permits per capita})$ jointly. Sharp-loss episodes are those in which the mayor's first-round vote share fell by more than ten percentage points; stable-support episodes remain within five points. All models adjust for baseline support and pre-election log population density and include state and focal-year fixed effects. Brackets contain 95 percent confidence intervals based on standard errors clustered by municipality. The sample contains 1,204 episodes, including 217 sharp-loss episodes. Estimating the trend requires at least four observations during the baseline term.

Treated episodes permitted slightly less during the baseline term than stable-support episodes: the normalized differences are −0.114 for total, −0.142 for single-family, and −0.035 for multifamily permits. In the adjusted models, no level or trend coefficient exceeds 1.8 percentage points per standard deviation in absolute value. Effects of 3.6–4.8 percentage points would have been detectable at 80 percent power. The test weighs against moderate selection on prior permitting but cannot exclude smaller relationships.

The event-study leads in Figure 2 compare years within the baseline term. They cannot detect whether treated municipalities entered that term at an unusually high permitting rate and then reverted toward their earlier rate. Figure C.10 therefore adds an earlier comparison. For each episode with sufficient history, I compare the focal baseline term with permitting before the election that began that term. The equal-length specification matches the preceding window to the observed baseline-term length. The three-year specification uses three years on each side of the baseline election. Both use episode and state-by-calendar-year fixed effects, a population offset, baseline support, and density measured before the baseline election.

The estimated differential increase into the baseline term ranges from 0.6 percent for single-family permits to 9.7 percent for multifamily permits. Every interval includes zero. The diagnostic does not detect the unusually elevated baseline permitting required by a simple backlash-and-reversion account, but neither does it rule out a moderate increase. Only 160 treated episodes have the complete history required by the equal-length comparison, and the multifamily interval is especially wide.

Figure C.10: Permitting before and during the focal baseline term

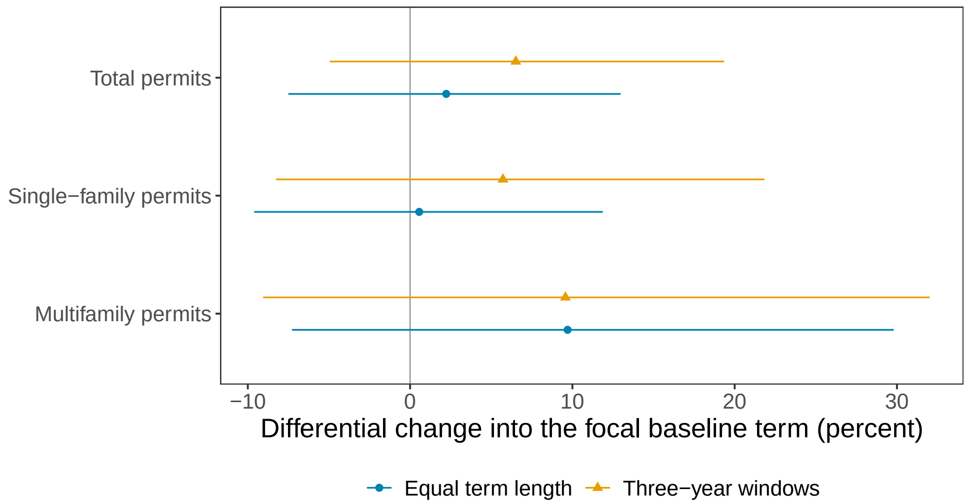

Notes: Points report the treated-minus-stable difference in the change from earlier permitting history into the focal baseline term. Equal-length windows match the preceding history to the observed baseline-term length. Three-year windows use the three years before the baseline election and the final three years of the baseline term. Bars are 95 percent confidence intervals. Standard errors are clustered by municipality.

D Politician Survey: Sample and Design

D.1 Sampling, recruitment, and completion

The survey was administered online in German in February and March 2026. I sampled municipalities by state, population, and local CDU/CSU vote share, including every municipality with at least 100,000 residents. The sampling frame combined 40,055 individualized addresses collected from official municipal and party websites with 8,657 municipal administration addresses. Administrations were asked to forward the invitation where individual contacts were sparse.

Table D.1 reports the final survey dataset used in the paper. It contains 5,238 substantive responses, including 4,390 from individualized links and 848 from municipal links. Among them, 4,213 reached 100 percent progress. All analyses retain every respondent who answered the relevant item.

Table D.1: Survey response pipeline

Stage	Individual	Municipal	Total
Invitations sent	40,055	8,657	48,712
Substantive responses	4,390	848	5,238
Substantive response rate (%)	11.0	–	–
Reached 100% progress	3,626	587	4,213
Reached 100% progress (%)	83	69	80

Notes: Substantive responses exclude Qualtrics metadata rows and records without survey content. The individual-channel response rate divides 4,390 substantive responses by 40,055 invitations. No response rate is defined for administration links because each link could be forwarded to several politicians. Progress is the Qualtrics completion measure. Item-level analyses include partial responses when the relevant question was answered.

D.2 Sample composition and selection

The respondent sample includes all major parties, all 16 states, and municipalities across the population distribution. Most respondents are current elected officials, and mayors constitute 11.2 percent of respondents who report their current or most recent office (Table D.2).

Table D.2: Survey respondent composition

	N	Percent
Panel A: Officeholding status		
Current elected official	4,749	93.6
Former elected official	210	4.1
Other local-government role	114	2.2
Panel B: Current or most recent office		
Council or other office	4,404	88.8
Mayor	556	11.2
Panel C: Party		
Other or independent	1,137	25.5
CDU/CSU	1,060	23.8
SPD	824	18.5
Greens	684	15.4
Free Voters	351	7.9
FDP	140	3.1
Left	134	3.0
AfD	122	2.7
Panel D: Gender		
Male	3,039	69.3
Female	1,348	30.7

Notes: Percentages are calculated among respondents with nonmissing information for the relevant panel. Party categories combine multiple selections and free-text responses into a single harmonized classification.

Table D.3 compares respondents with nonrespondents in the individual invitation frame. Respondents come from municipalities with somewhat narrower council winning margins, higher density and housing costs, and lower homeownership. They are also more likely to be women, SPD politicians, or Green politicians. These differences matter for population descriptions. I therefore interpret the survey as evidence about participating officials rather than as a population-weighted estimate for all German local politicians.

Table D.3: Respondents and nonrespondents in the individual invitation frame

	Respondents	Non-respondents	SMD	*p*-value
	N = 3,797 resp., 31,479 non-resp.			
Council winning margin	0.108	0.126	-0.179	$< .001$
Population (thousands)	221.995	170.387	0.084	$< .001$
Population density	897.221	751.188	0.154	$< .001$
Rent (EUR/sqm)	10.809	10.432	0.122	$< .001$
Purchase price (EUR/sqm)	4060.507	3859.962	0.105	$< .001$
Vacancy rate	0.042	0.043	-0.053	0.002
Homeownership rate	0.484	0.518	-0.204	$< .001$
East Germany	0.132	0.114	0.056	$< .001$
Female	0.305	0.275	0.066	$< .001$
Mayor	0.074	0.068	0.024	0.144
CDU/CSU	0.173	0.202	-0.072	$< .001$
SPD	0.163	0.121	0.122	$< .001$
Green	0.132	0.072	0.199	$< .001$

Notes: Municipality characteristics are measured in the latest available year before the survey. SMD is the standardized mean difference, coded as respondents minus nonrespondents. The sample contains individual invitations that can be linked to the relevant measure.

D.3 Instrument and conjoint design

The instrument measures perceived local problem severity, beliefs about the price consequences of supply, evaluations of housing-policy instruments, a paired conjoint experiment, perceived electoral safety, and respondent characteristics. Belief questions precede the conjoint. Electoral questions and demographics appear at the end.

Each respondent receives five paired housing-development choices. Four tasks are independently generated. The fifth repeats the first with profile positions reversed to assess response reliability following Clayton et al. (2026). Attribute levels are randomized subject to three feasibility rules. Projects that expand the housing stock by five or ten percent cannot finish within one to two years, densification without a zoning change cannot produce a ten-percent stock increase, and priority for local residents appears only in new residential areas that take at least three years.

Table D.4: Conjoint attributes and levels

Attribute	Levels
Housing-stock increase	2%, 5%, 10%
Time to completion	1–2 years, 3–4 years, at least 5 years
Price binding	Market prices, partly price-bound, predominantly price-bound
Access rules	Priority for local residents, no residency priority
Planning pathway	New development area, densification with zoning change, densification without zoning change
Fiscal impact	Net burden, approximately balanced, net fiscal relief
Reaction of affected residents	Strong opposition, divided, few reactions, strong support
Reaction in the municipality	Strong opposition, divided, few reactions, strong support

Notes: Percentage increases were displayed relative to the actual housing stock when the municipality could be linked. Reference levels are 2 percent, one to two years, market prices, no residency priority, densification without a zoning change, a balanced fiscal effect, and little feedback from residents or the municipality.

D.4 Data quality and representativeness

The survey dataset contains 5,238 substantive respondents, including partial responses, and 4,213 respondents who reached 100 percent progress. The conjoint source file contains the 4,594 respondents who reached the module and 36,752 profiles from the four independently generated tasks. Removing profiles without an observed paired choice leaves 4,505 respondents and 35,364 profiles in the full analytic conjoint sample. The registered primary analysis then restricts this sample to 4,251 current elected municipal officeholders and 33,362 profiles. These counts refer to successive stages of one response pipeline rather than different survey waves. Among officeholders who complete the repeated task, intra-respondent reliability is 79.3 percent. Appendix E compares the main estimates with a correction for response noise, the full analytic sample, and a sample that excludes the fastest five percent of officeholders. The ordering and magnitudes change little. Municipality linkage occurs only after the survey response is stored, and analysis files replace names and email addresses with anonymous respondent identifiers.

E Additional Survey Evidence

The registered estimands and restricted randomization appear first, followed by tests of independent project costs, moderator analyses, response consistency, and sample choices. The section then reports marginal means, descriptive salience evidence, and departures from the pre-analysis plan.

E.1 Pre-analysis-plan estimands and restricted randomization

The pre-analysis plan registers the conjoint experiment, not the survey's salience, belief, policy-instrument, or sample-selection modules. Its directional predictions receive mixed support. Net fiscal relief raises project support and fiscal burdens reduce it, supporting the fiscal prediction. Opposition among affected residents and in the wider municipality reduces support, supporting both opposition predictions. Price regulation and priority for local residents raise support. The registered expectation that larger projects would receive less support is not supported. Five- and ten-percent projects are both modestly more likely to be chosen than two-percent projects.

The registered design excludes combinations that are infeasible in practice. Projects that expand the housing stock by five or ten percent cannot finish within one to two years, densification without a zoning change cannot increase the housing stock by ten percent, and local-resident priority appears only for non-fast projects in new residential areas. These rules leave fiscal effects, price restrictions, and expected reactions independently assigned. Their additive OLS coefficients therefore identify AMCEs directly. Scale, completion time, access, and planning pathway are linked by the restrictions. For these attributes, a nonparametric AMCE must compare the focal levels within every exact stratum where both are assignable and then average the stratum-specific contrasts using the known assignment probabilities (Hainmueller, Hopkins and Yamamoto, 2014).

Figure E.1 applies this procedure to all seven linked-attribute contrasts. Every adjusted estimate lies within 0.7 percentage points of its pre-registered additive OLS coefficient, and every sign and inferential conclusion is unchanged. Completion after five or more years reduces support by 9.2 points, while a new residential area reduces it by 12.2 points. Among non-fast new-area projects, priority for local residents increases support by 6.6 points. The feasibility restrictions therefore do not account for the timing, scale, or planning results. They narrow the access estimate to the project settings in which both access rules were legally plausible.

E.2 Do the political costs operate independently?

The main conjoint estimates average over tasks in which several project attributes can distinguish the two profiles. A single categorical objection, especially fiscal burden or a long delay, could therefore account for the apparent effects of the other attributes. I assess this possibility with a post hoc diagnostic. I classify each paired task according to whether fiscal burden appears for one profile, both profiles, or neither profile. I then estimate the pre-registered profile-level OLS while allowing every attribute coefficient to vary across these three randomized task strata. I repeat the procedure after classifying tasks by whether completion after five or more years appears for one profile, both profiles, or neither profile.

Figure E.2 shows that the relevant penalties remain negative in every stratum. When fiscal burden cannot decide the choice because both or neither profile imposes it, long completion, a new residential area, and opposition still reduce support. When long completion cannot decide the choice, fiscal burden and both measures of opposition continue to reduce support. The penalty for opposition among affected residents varies across fiscal-burden strata, and the budget-burden penalty varies across completion-time strata. These differences do not alter the central diagnostic: every cost remains negative when the other cost is held constant across profiles. The conjoint choices therefore reflect several independent tradeoffs rather than a single fiscal or temporal veto.

Figure E.1: Restriction-adjusted estimates reproduce every linked-attribute result

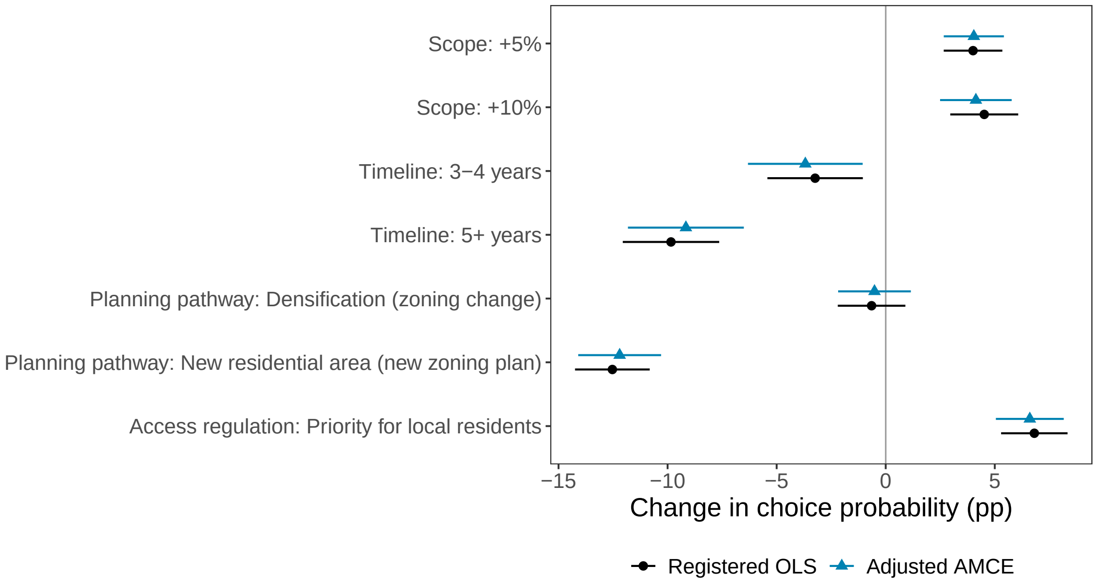

Notes: Black circles are coefficients from the pre-registered additive stacked OLS specification. Blue triangles are nonparametric AMCEs estimated among current officeholders. Each adjusted estimate compares the focal levels within every exact common-support stratum and averages the stratum-specific contrasts using the declared assignment probabilities. Bars are 95 percent confidence intervals with standard errors clustered by respondent. Sample sizes vary across contrasts because their common-support sets differ.

Figure E.2: Project costs continue to matter when another cost cannot decide the choice

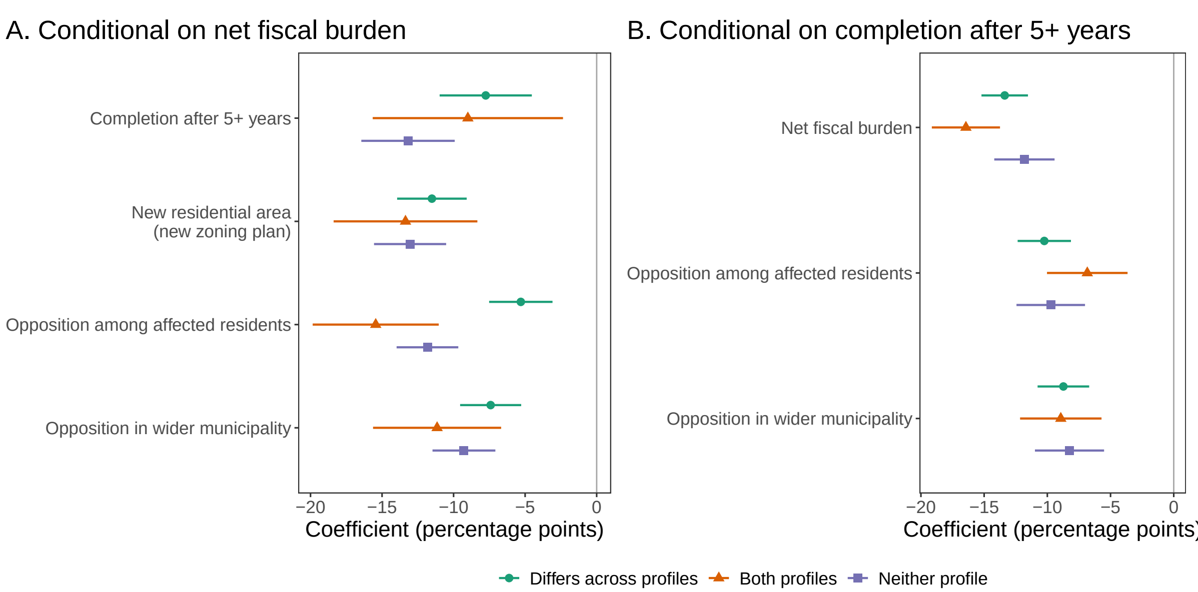

Notes: Points are coefficients from stacked profile-level OLS estimates among current officeholders, with all conjoint coefficients allowed to vary across the displayed randomized task strata. Panel A conditions on whether net fiscal burden differs across profiles, appears for both, or appears for neither. Panel B uses the same classification for completion after five or more years. Bars are 95 percent confidence intervals with standard errors clustered by respondent. The fiscal strata contain 7,304, 1,837, and 7,540 tasks. The completion-time strata contain 8,289, 3,502, and 4,890 tasks.

E.3 Pre-registered heterogeneity

The perceived-safety moderator asks whether randomized project attributes have different effects among politicians who consider themselves electorally safe and vulnerable. Because perceived safety is measured rather than assigned, this analysis compares experimental attribute effects across observed respondent groups. It does not identify the causal effect of electoral safety on project choice. Heterogeneous effects also support claims about mechanisms only under additional assumptions, while null differences generally do not establish that a mechanism is absent (Fu and Slough, 2026).

Figure E.3 implements the exact registered split between current officeholders who view re-election as rather or very likely and those who view it as rather or very unlikely. One of the seventeen attribute-level differences reaches the unadjusted five-percent threshold: the budget-burden penalty is smaller among low-safety respondents. It does not survive a Benjamini–Hochberg correction, and the remaining estimates show no consistent ordering across opposition, fiscal costs, planning requirements, or delay. The low-safety group contains only 180 respondents with observed conjoint choices and is too small for a precise comparison. The result therefore provides no systematic evidence that the randomized project attributes have different effects across perceived-safety groups, but it also does not tightly rule out such differences.

I also examine supply beliefs and housing-market pressure, along with exploratory measures of perceived market conditions, population change, migration, recent permitting composition, and prior anti-development petitions. The subgroup estimates reveal no single pattern that recurs across all attributes or moderator families. The clearest regularity is that politicians who perceive high demand or high prices, and those in municipalities with high current rents or purchase prices, are more favorable to larger and price-regulated projects. Growth-based measures are less consistent. Figures E.4–E.6 report these differences. Filled points denote nominally significant differences and should be interpreted as patterns across related attributes rather than as isolated tests.

Figure E.3: Conjoint effects by perceived electoral safety

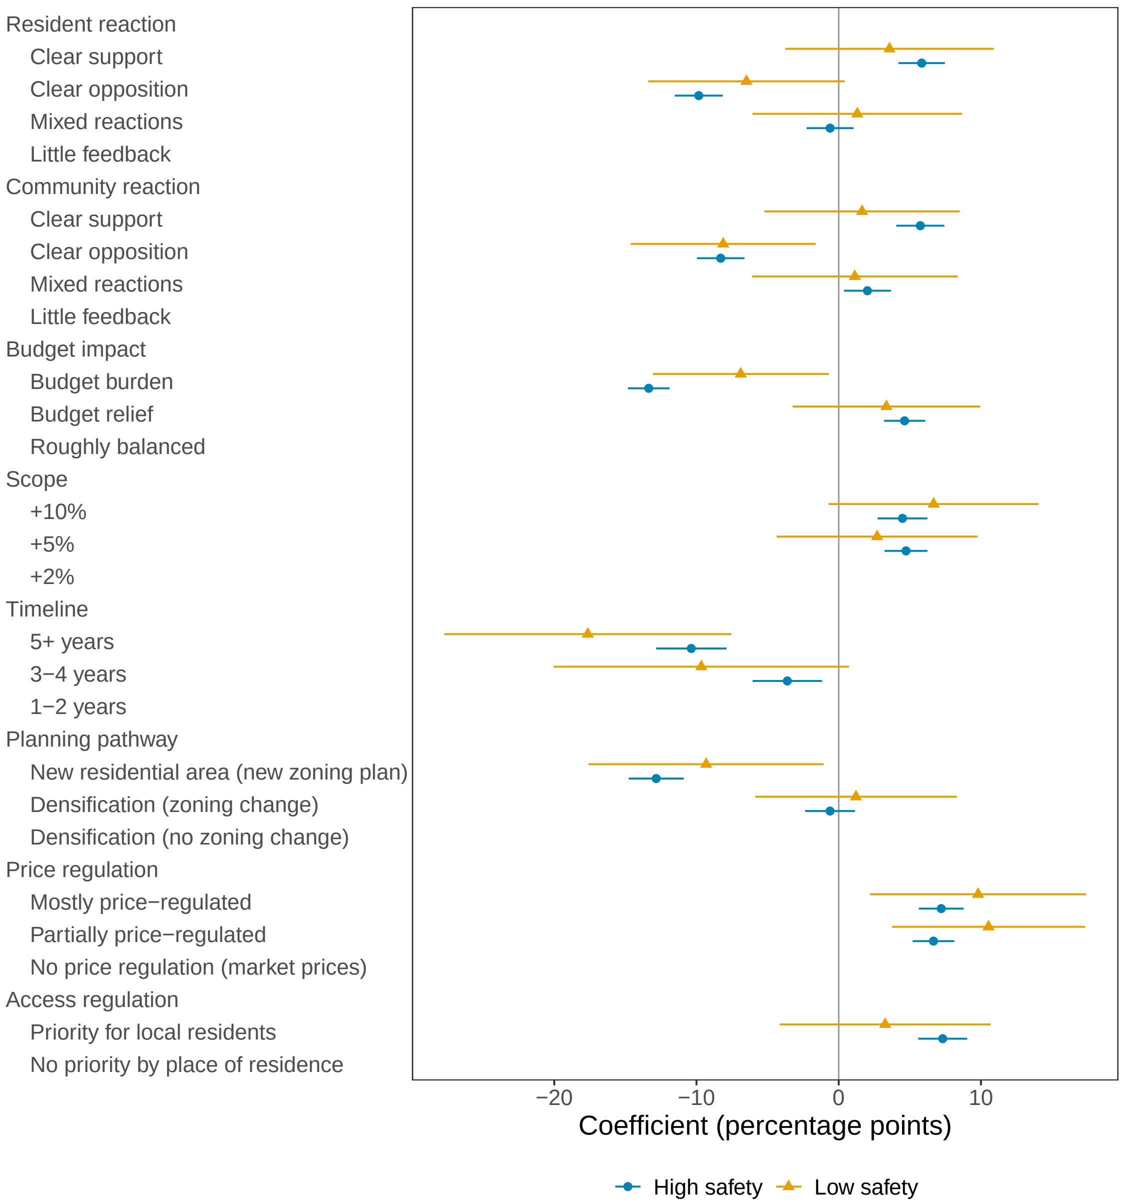

Notes: Points are coefficients from the pre-registered stacked profile-level OLS specification, estimated separately for current officeholders who view reelection as rather or very likely and those who view it as rather or very unlikely. Midpoint and don't-know responses are omitted. Bars are 95 percent confidence intervals with standard errors clustered by respondent. $N = 3{,}372$ high-safety and 180 low-safety respondents with observed choices.

Figure E.4: Conjoint heterogeneity by perceived market conditions

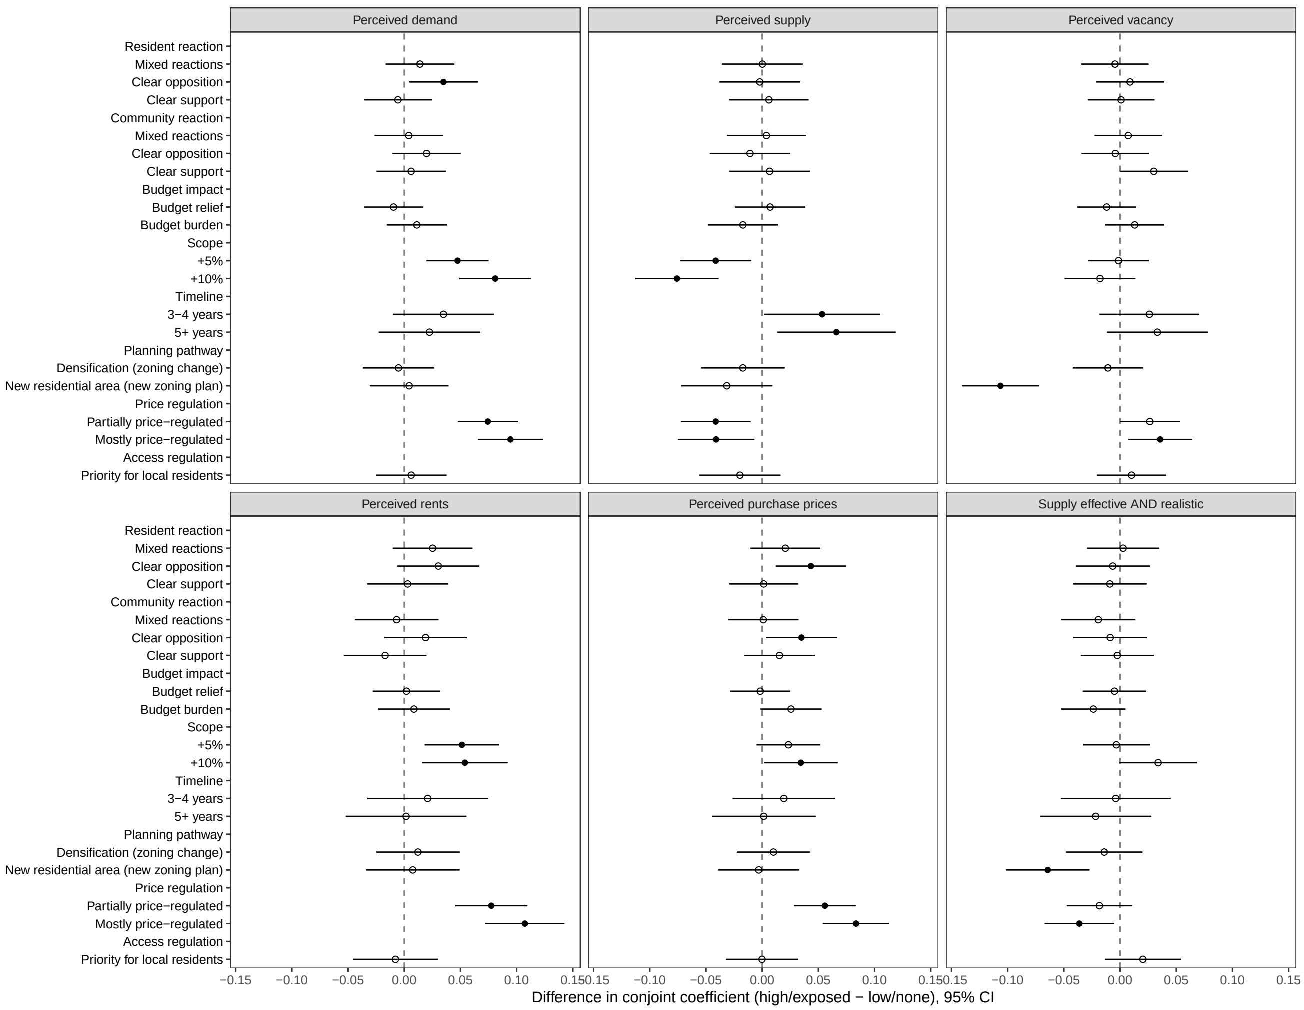

Notes: Points are differences in coefficients from the pre-registered stacked profile-level OLS specification between the high or exposed group and the low or unexposed group. Bars are 95 percent confidence intervals with standard errors clustered by respondent. Filled points have an unadjusted $p < .05$.

Figure E.5: Conjoint heterogeneity by municipal conditions

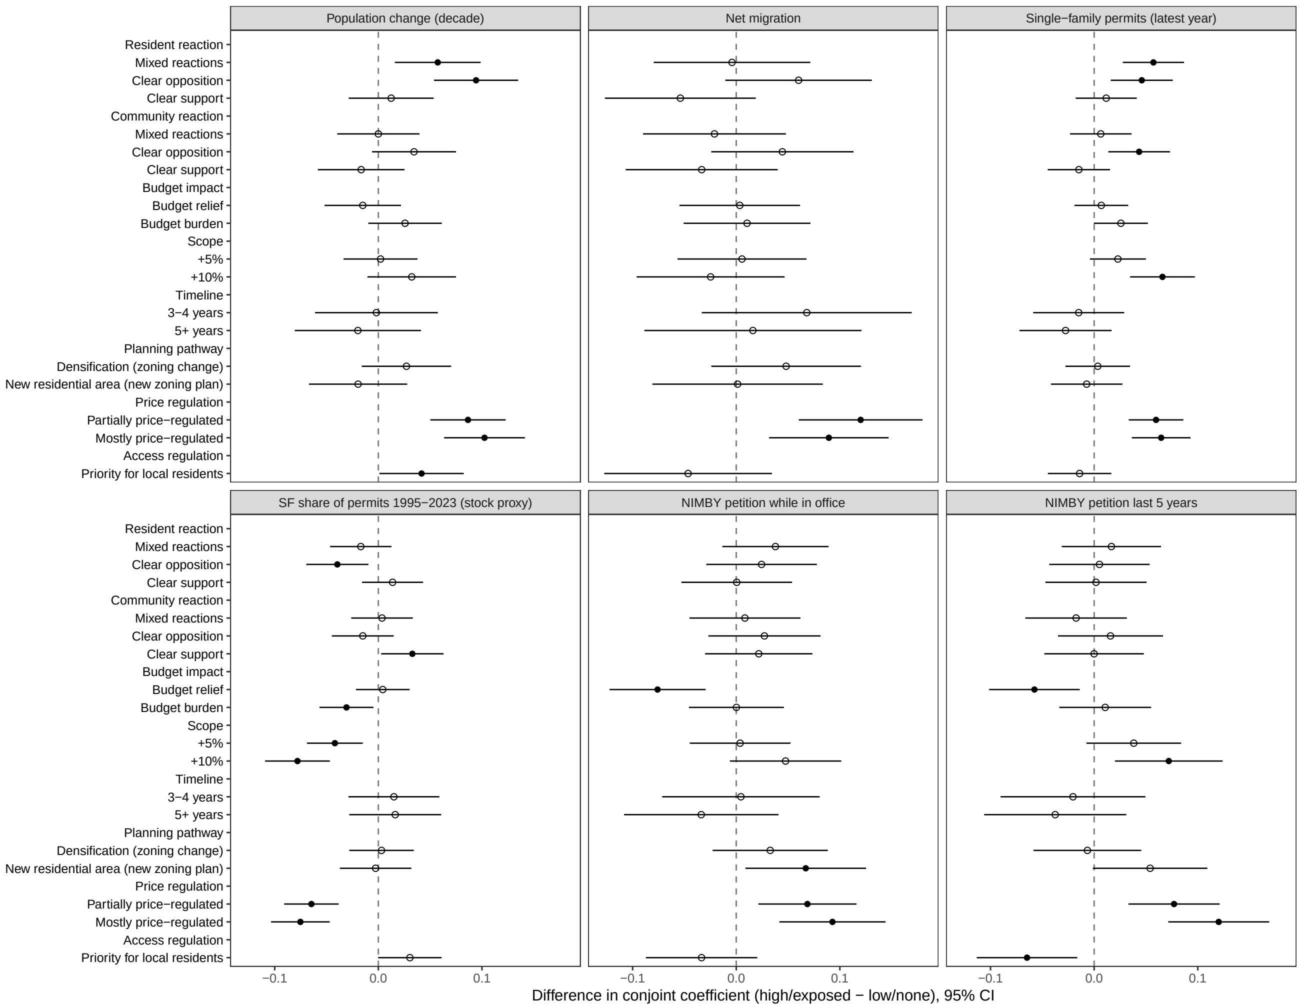

Notes: Points are differences in coefficients from the pre-registered stacked profile-level OLS specification between the high or exposed group and the low or unexposed group. Bars are 95 percent confidence intervals with standard errors clustered by respondent. The single-family stock measure is the cumulative share of permitted units in single-family buildings from 1995 through 2023. Petition data end in 2021. Filled points have an unadjusted $p < .05$.

Figure E.6: Conjoint heterogeneity by observed housing-market pressure

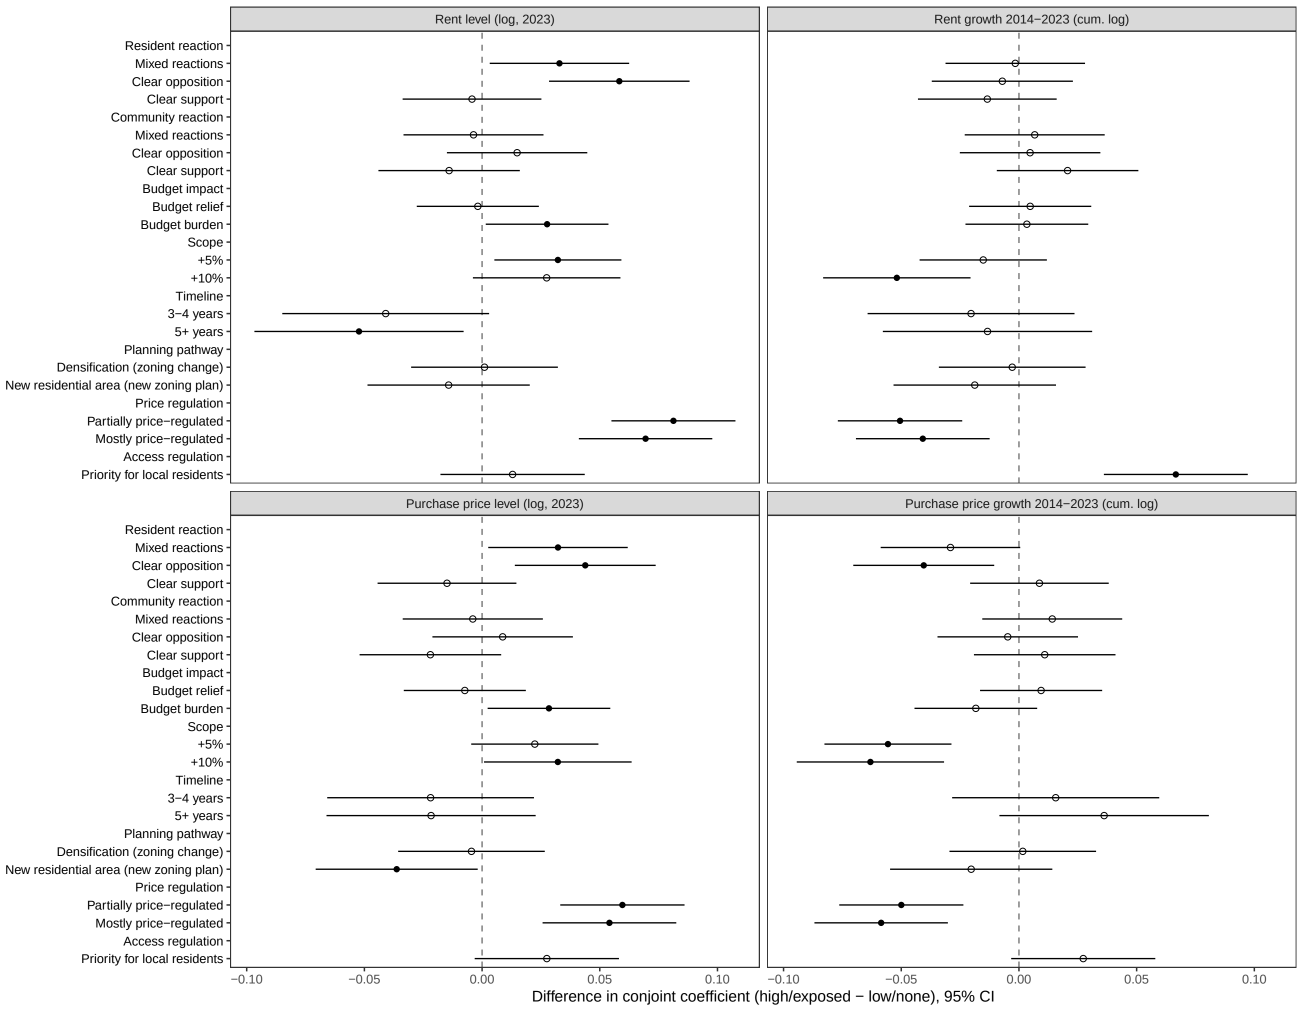

Notes: Points are differences in coefficients from the pre-registered stacked profile-level OLS specification between municipalities above and below the median of each measure. Price levels use 2023, and growth is cumulative from 2014 through 2023. Bars are 95 percent confidence intervals with standard errors clustered by respondent. Filled points have an unadjusted $p < .05$.

E.4 Response consistency and measurement noise

The fifth conjoint task repeats the first pair with its left and right positions reversed. A consistent respondent selects the same underlying project both times. Figure E.7 reports an intra-respondent reliability rate of 79.3 percent among current officeholders, well above the 50 percent expected from random choice.

Figure E.7: Consistency across the repeated conjoint task

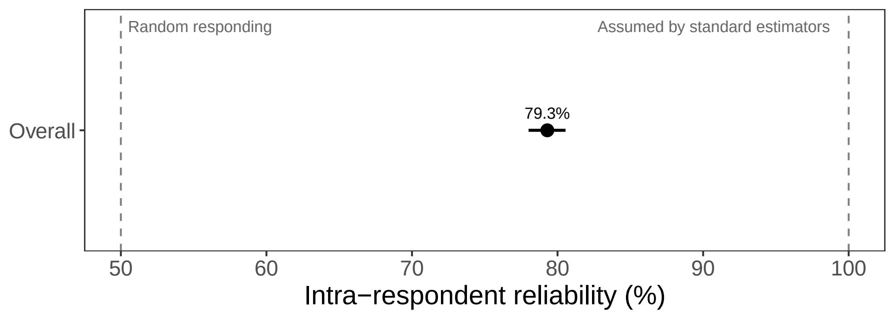

Notes: Intra-respondent reliability is the share of current officeholders choosing the same underlying project in Task 1 and its position-reversed repetition in Task 5. Bars are 95 percent Wilson confidence intervals. The reference lines mark random choice and perfect consistency. $N = 4{,}033$.

Response inconsistency attenuates conjoint effects toward zero. Figure E.8 applies the correction proposed by Clayton et al. (2026). Panel A reports the registered conjoint coefficients. The pre-analysis plan also registers marginal means, which appear in Panel B. Correcting for swapping error moves both sets of estimates away from their null values but leaves their ordering intact. The main text reports the uncorrected coefficients.

Figure E.8: Conjoint estimates with and without a response-noise correction

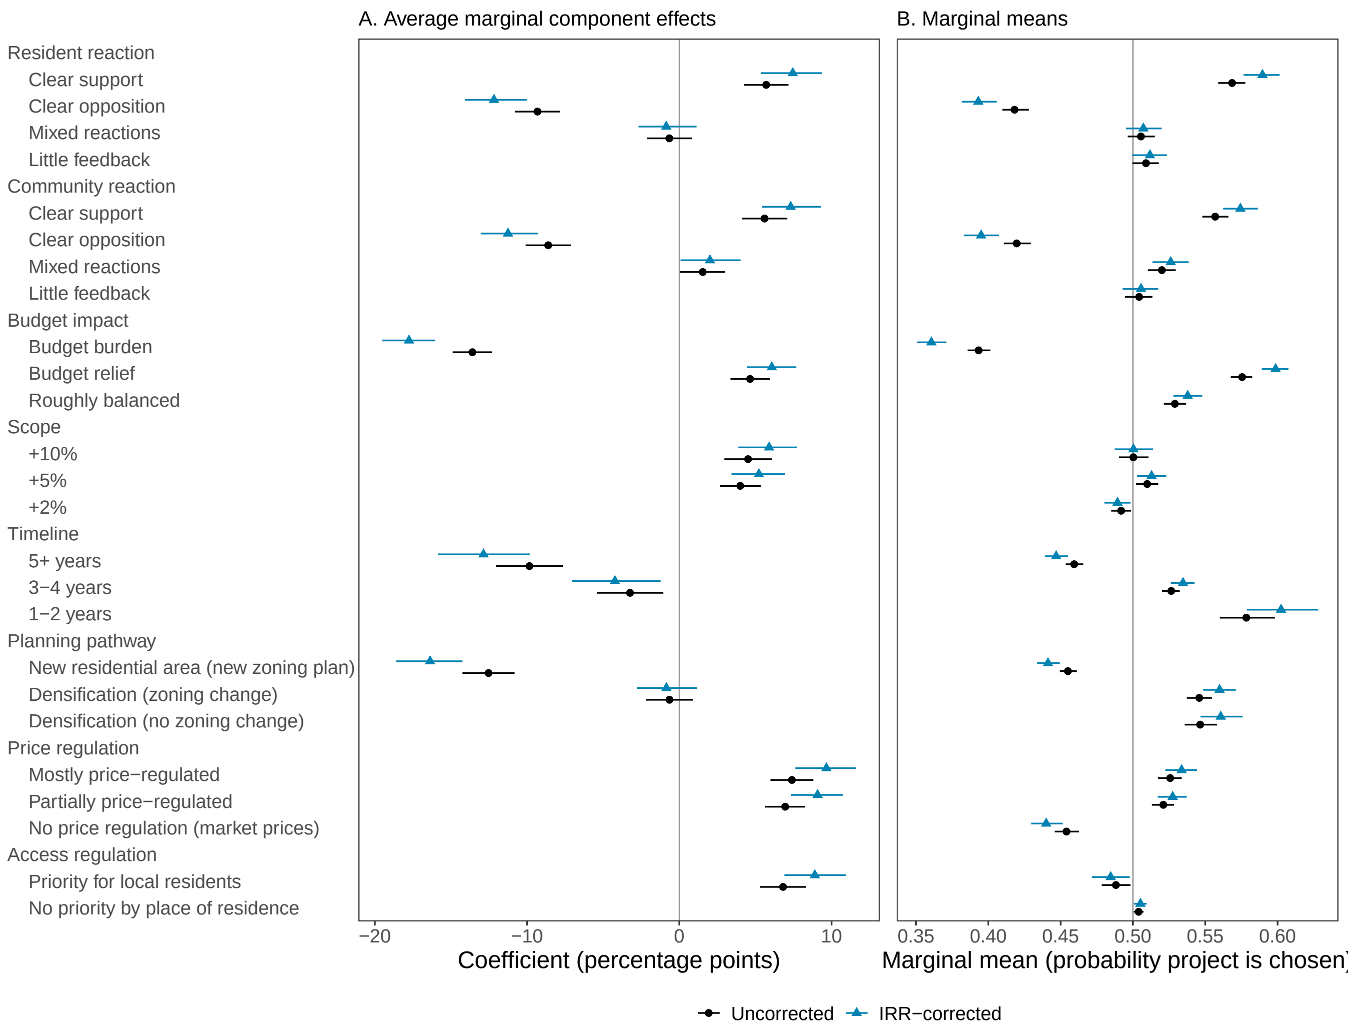

Notes: Panel A reports coefficients from the pre-registered stacked profile-level OLS specification; reference levels are omitted because their coefficients equal zero by definition. Panel B reports marginal means for every attribute level. Points are uncorrected and reliability-corrected estimates. The correction uses the Task 1–Task 5 agreement rate and a respondent-level bootstrap with 1,000 draws that recomputes reliability, conjoint coefficients, and marginal means. Bars are 95 percent confidence intervals.

Reliability is similar across the registered moderator groups. It equals 0.789 among respondents with high perceived safety and 0.749 among the smaller low-safety group, with overlapping intervals. Across party families it ranges from 0.752 to 0.821. A diagnostic model using completion time, housing familiarity, office, party, ideology, and state explains no meaningful share of inconsistent responses ($R^2 \leq .001$).

E.5 Sample and specification checks

Figure E.9 compares the primary officeholder estimates with the full respondent sample, officeholders excluding the fastest five percent, and the registered secondary covariate adjustment. These analyses contain 4,251, 4,505, 4,036, and 4,088 respondents, respectively. The covariate adjustment includes party, office, tenure, ideology, and state fixed effects. On the same complete-case sample, it changes no attribute coefficient by more than 0.02 percentage points. The results are therefore insensitive to the population and specification choices registered in the pre-analysis plan.

E.6 Housing as a local policy problem

The electoral mechanism is relevant only if local politicians perceive housing as a problem worth addressing. I therefore asked officeholders to rate the importance of six municipal policy areas. This module is descriptive and was not part of the conjoint pre-analysis plan. Figure E.10 reports the distribution of responses for every issue.

Housing is rated as rather or very important by 80 percent of officeholders. It ranks behind municipal finance, childcare and schools, and transportation, but clearly above economic development and environmental protection. The result establishes that housing pressure is salient to the surveyed officials without treating the sample as a population-weighted estimate of all German local politicians.

E.7 Departures from the pre-analysis plan

Several implementation details differ from the plan. North Rhine–Westphalia entered the sampling frame after fieldwork moved beyond the 2025 local election. The conjoint estimates use the four independently generated tasks, while the fifth task is reserved for reliability. The exact safety split omits the midpoint as registered and combines equivalent reelection items when the primary item is unavailable. The latest complete rent and purchase-price year is 2023 rather than 2024. The single-family composition measure uses the cumulative share of permits because a current municipal housing-stock share is unavailable. The citizen-initiative series ends in 2021, so exposure while in office is observed only through that year. These deviations affect coverage or measurement, not the randomized assignment of conjoint attributes.

Figure E.9: Conjoint effects across respondent samples

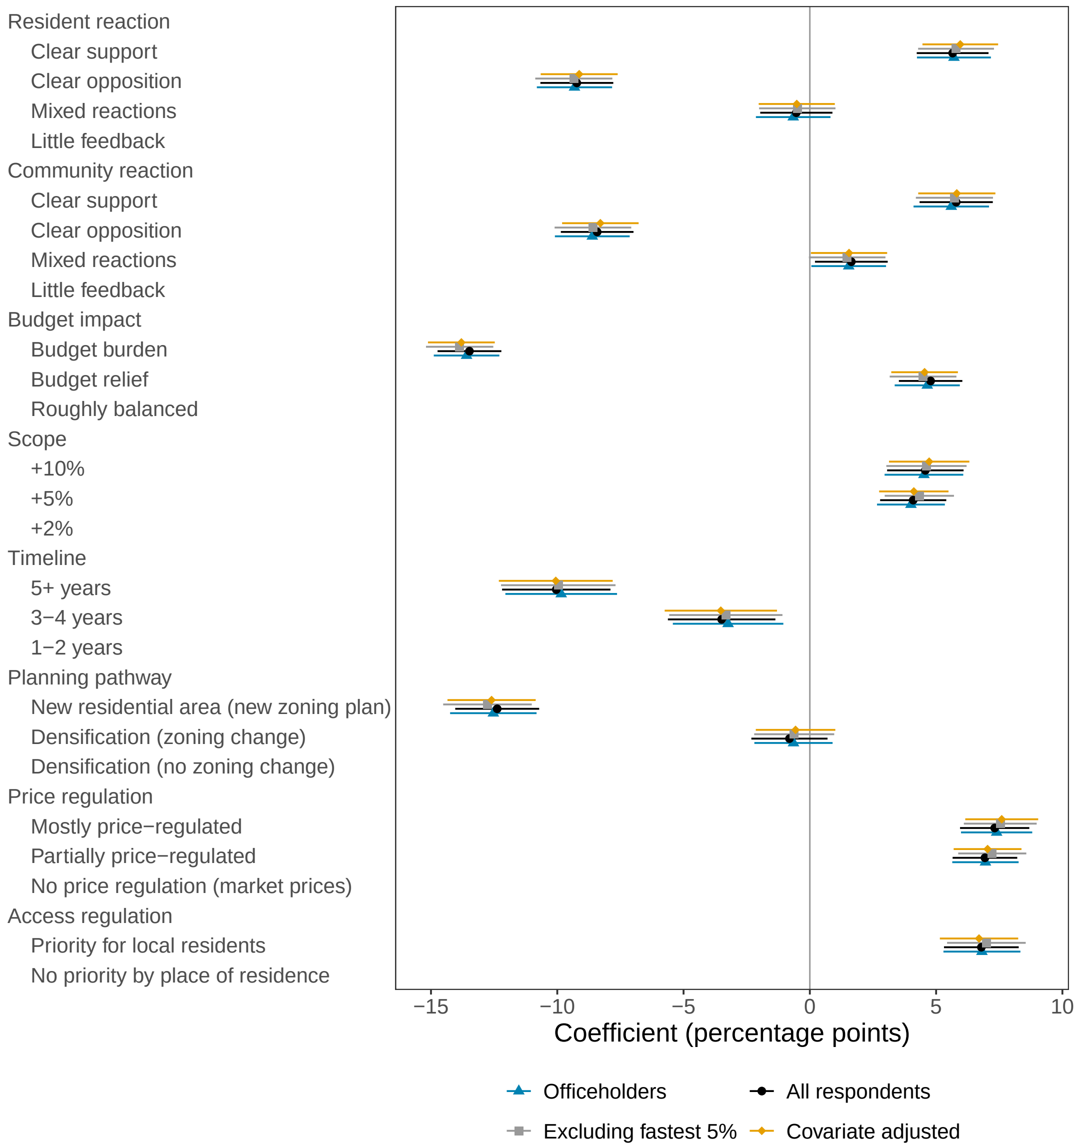

Notes: Points are coefficients from the pre-registered stacked profile-level OLS specification among current officeholders, all respondents, officeholders excluding the fastest five percent, and covariate-adjusted officeholders. The adjusted model includes party, office type, tenure, political and economic ideology, and state fixed effects. Bars are 95 percent confidence intervals with standard errors clustered by respondent.

Figure E.10: Perceived importance of local policy problems

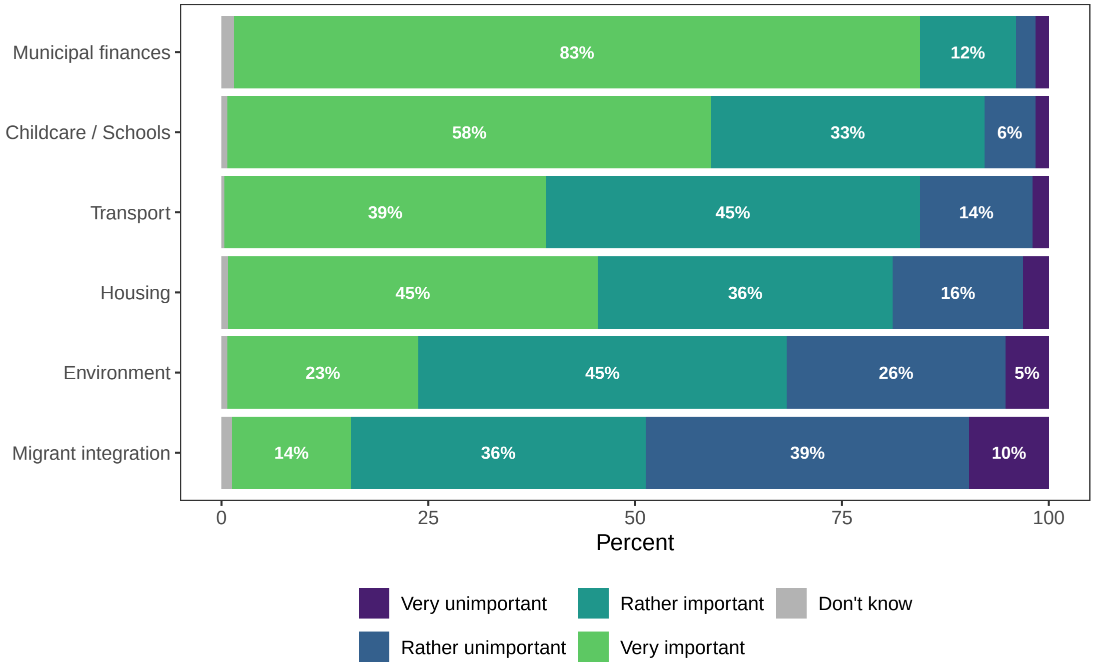

Notes: Shares report current officeholders' ratings across six municipal policy areas. Item-specific samples range from 4,483 to 4,486 officeholders.

F US Data, Design, and Additional Results

F.1 Data

The US analysis combines four sources. Mayoral election spells come from de Benedictis-Kessner, Jones and Warshaw (2025). Candidate-level vote shares and incumbency come from the American Local Elections Database (de Benedictis-Kessner et al., 2023). Annual permit counts come from the Census Building Permits Survey: the 1990–2021 counts from the replication archive of de Benedictis-Kessner, Jones and Warshaw (2025), and the 2022–2024 counts from the Census annual place files, which reproduce the archive's counts exactly on 99.8 percent of the 12,239 overlapping city-years. Population for 2022–2024 comes from the annual Census place estimates, chain-linked to each city's last panel population so that the offset has no break in vintage. Term-limit rules originate in the ICMA Municipal Form of Government Survey but enter the analysis only through the legal histories described below. Election spells are joined to city-years, all identifiers are standardized to Census place FIPS codes, and election years are dropped because annual permit counts mix months governed under the old and new term. The appendix vote-share extension uses the winning candidate's share of all votes cast. Council selections, uncontested placeholders, and other procedures that record no public vote do not enter that analysis. Implausible vote records were corrected against official canvasses.

ICMA reports whether a city limits mayoral terms and, in most waves, the statutory maximum. Two failure modes make this reporting insufficient for a within-mayor design. First, inferring the maximum from the longest observed tenure is circular: Chandler, Arizona, for example, appears in

ICMA waves with both a two-term and a four-term maximum, and a tenure-based rule resolves the contradiction toward whichever ceiling the observed careers happen to reach. Second, a mayor's early service, service under a superseded statute, or a return after time out of office all break the term count. I therefore code final-term status from city charters, ballot measures, canvasses, and official service records. For about three in five mayor-terms in the estimation sample the full legal history was hand-verified. The remaining mayor-terms rest on a dated statutory ceiling with a citation, applied to the mayor's consecutive-term count. Every observation in the main analysis therefore has a dated and sourced ceiling.

The permit panel contains total, single-family, and multifamily dwelling units from 1990 through 2024. Models estimate raw counts with a log-population offset. The main term-limit sample contains 1,200 city-years, 220 mayors, and 112 cities before outcome-specific Poisson separation. The total and single-family models retain 1,199 observations. The multifamily model retains 1,181.

F.2 Term limits

Equation 4 in the main text gives the within-mayor model, where L_{ict} indicates that mayor i of city c serves in a final term under the city's statutory ceiling and α_i are mayor fixed effects. The estimate is identified by the 220 mayors observed in both eligible and final terms. Total permits are 17.0 percent higher in the final term (95 percent confidence interval: 5.2 to 30.1, $p = .004$), multifamily permits are 16.2 percent higher (3.5 to 30.5, $p = .011$), and single-family permits are 13.1 percent higher (1.4 to 26.2, $p = .027$).

Panel A of Figure F.1 reports four alternatives. Adding state-specific linear trends reduces the unrestricted estimates to 6.3 percent for total permits and 11.2 percent for multifamily permits. Excluding return careers yields corresponding estimates of 12.6 and 16.6 percent. Restricting the sample to two-term ceilings produces increases of 12.8 and 16.5 percent, while the above-median-population sample produces increases of 19.5 and 20.6 percent. The total and multifamily estimates remain positive throughout. The single-family estimate ranges from approximately zero with state trends to 12.5 percent among larger cities.

A treatment-label permutation test provides design-based inference for the main comparison. I randomly reassign final-term status across each mayor's observed terms 1,000 times, preserving the number of eligible and final terms for every mayor. The one-sided permutation p-values are .001 for total permits and .004 for multifamily permits. The max-t p-value is .012 across the two focal year-fixed-effects outcomes and .037 across all nine combinations of three time controls and three permit outcomes.

F.3 Same-mayor vote-share extension

I also estimate a continuous extension that parallels the German fixed-effects analysis. Let V_{ict} denote the vote share with which mayor i won the election governing city c in year t. I estimate

$$\mathbb{E}[Y_{ict} \mid \mathbf{X}] = \exp\left\{\beta V_{ict} + \alpha_i + \lambda_{s(c)t} + \log P_{ct}\right\}. \tag{7}$$

Mayor fixed effects compare the same mayor across election cycles. State-by-year fixed effects absorb shocks common to cities in the same state and year. The sample contains 925 mayors in 506 cities.

Panel B of Figure F.1 reports the estimates for a ten-point higher winner vote share. Multifamily permits increase by 3.9 percent (95 percent confidence interval: -0.5 to 8.6, $p = .084$). Total permits increase by 1.7 percent (-1.5 to 4.9, $p = .301$), while single-family permits change by -0.3 percent (-3.5 to 3.0, $p = .845$). The multifamily estimate is 5.8 percent ($p = .024$) when the panel ends in 2021, but the model with mayor and calendar-year fixed effects is null. I therefore treat this specification as supporting evidence rather than as a second US headline result.

Figure F.1: Additional US specifications

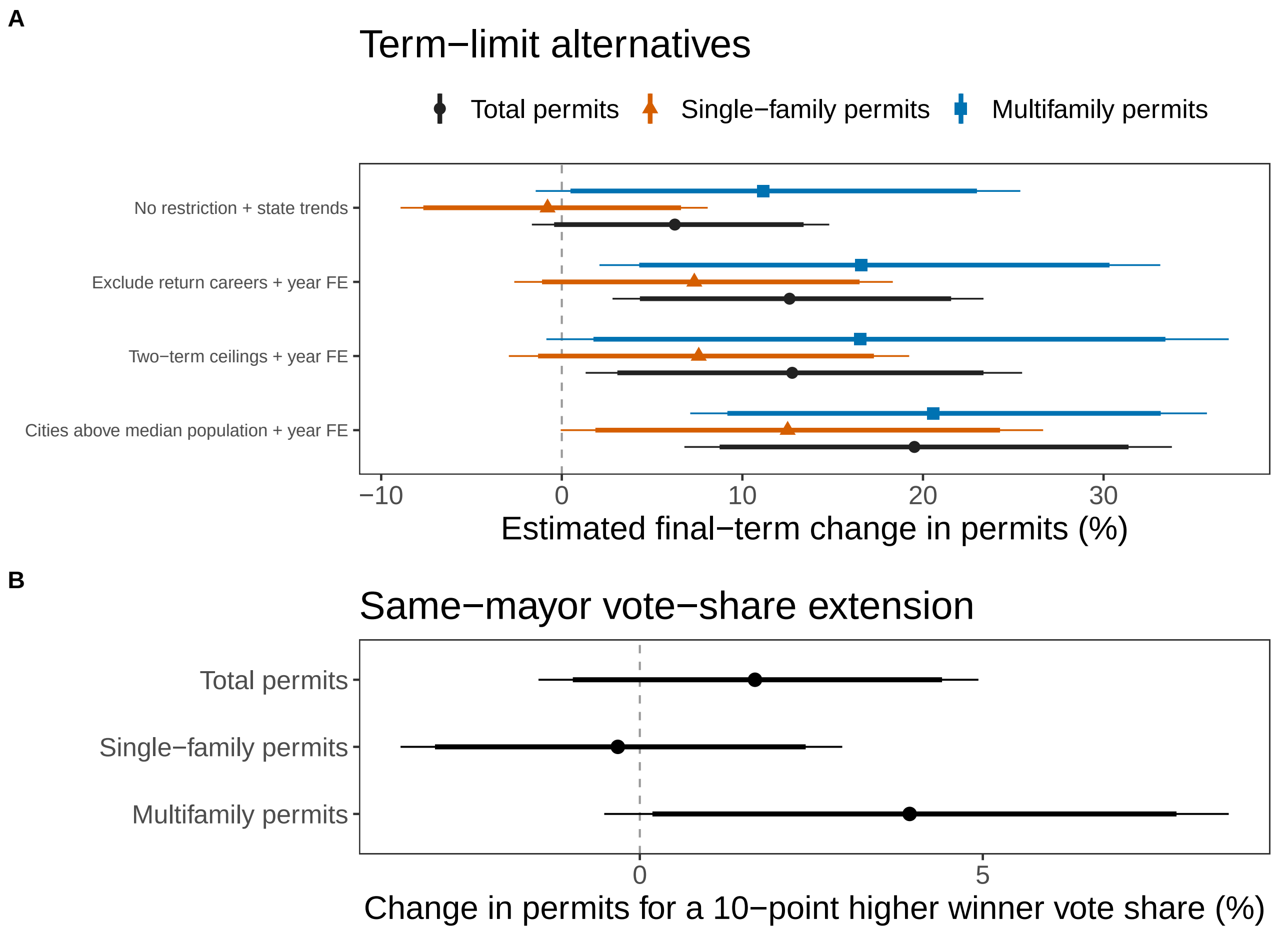

Notes: Panel A reports four alternatives to Equation 4. All use the dated-and-sourced legal tier, mayor fixed effects, Poisson permit counts with a log-population offset, and the 1990–2024 panel. "State trends" combines calendar-year fixed effects with state-specific linear trends. The remaining rows use calendar-year fixed effects and apply the sample restriction named on the axis. Panel B reports Equation 7 with mayor and state-by-year fixed effects. It shows the change associated with a ten-point higher winner vote share. Election years are omitted throughout, and standard errors are clustered by city. Thin bars show 95 percent confidence intervals and thick bars show 90 percent confidence intervals.