How Costly Is Permitting in Housing Development?*

Evan Soltas
Jonathan Gruber

February 2026

Abstract

Permitting costs are widely cited, but little analyzed, as a key burden on housing development in leading U.S. cities. We measure them using an implicit market for "ready-to-issue" permits in Los Angeles, where landowners can prepay permitting costs and sell preapproved land to developers at a premium. Using a repeat-listing difference-in-differences estimator, we find developers pay 50 percent more ($48 per square foot) for preapproved land. Comparing similar proposed developments, preapproval raises the probability of completing construction within four years of site acquisition by 10 percentage points (30 percent). Permitting can explain one third of the gap in Los Angeles between home prices and construction costs.

Keywords: building permit, land-use regulation, hedonic, zoning, capitalization
JEL Codes: K25, R38, R52

*Gruber: MIT Economics (gruberj@mit.edu). Soltas: Princeton Economics and SPIA (esoltas@princeton.edu).
Allan Lee and Audrey Wang provided excellent research assistance, and Charlotte Young, Christine Sinn, and Elizabeth Wright provided valuable additional support. We thank seminar audiences at Johns Hopkins, Maryland, Princeton, and Texas A&M for helpful comments. Soltas gratefully acknowledges funding from Arnold Ventures and the Griswold Center for Economic Policy Studies.

1 Introduction

The cost of housing is one of the top economic issues facing large, highly productive U.S. cities. In places like Boston and Los Angeles, rents and home prices are substantially higher than the national average.[1] Higher costs of living offset these cities' higher nominal incomes, leaving their residents not consistently better off than in lower-income, lower-cost cities (Diamond and Moretti, 2021). Moreover, scholars have estimated that the lack of housing in such cities meaningfully reduces aggregate U.S. output (Hsieh and Moretti, 2019; Duranton and Puga, 2023) and most greatly burdens the urban poor and middle class (Ganong and Shoag, 2017; Autor, 2019).

Among all large U.S. cities with current monthly rents above $2,000 on a typical apartment, none issued permits for more than five new homes per 100 residents cumulatively from 2015 to 2025. Why haven't these cities built their way out of such high rents? An academic literature has argued for the importance of restrictive land-use regulations, rather than high construction costs, in limiting a supply response. The most prominent evidence for this view is a series of articles, starting with Glaeser and Gyourko (2003), that finds large gaps between house prices and estimates of their construction costs, consistent with binding regulatory constraints.

A challenge facing both scholars and policy analysts is that the regulatory environment in real estate development is complex, involving not only explicit taxes, limits, and mandates but also a set of diffuse procedural burdens (Baum-Snow and Duranton, 2025). One of the most widely-mentioned sources of burden is permitting, or obtaining legal approval to build. Policy analysts have argued that permitting exposes projects to significant risk and delay (O'Neill et al., 2019a,b). Such arguments have fueled debate about an "abundance agenda" (Klein and Thompson, 2025) and a wave of reforms intended to simplify permitting.[2] Yet it remains unclear whether, in high-cost U.S. cities, permitting is a minor nuisance or a major impediment to housing supply. In particular, many of these cities also have restrictive zoning rules and high construction costs, forces that may be a far larger burden than permitting on its own.

Measuring permitting costs is challenging for several reasons. First, the costs are unobserved and likely have both time and resource components. Developers' compliance expenditures are proprietary information, nor is it clear what cost of capital should be used to value permitting-related delays and risks of project failure. Second, permitting costs are equilibrium objects, not policy parameters. Developers propose structures with permitting costs in mind, creating a selection problem that complicates the interpretation of permit data. Third, and most fundamentally, rights to build are tied to specific properties, such that permits are never traded separately from land

[1] In the latest data available as of 2024, rents are 35 to 54 percent above the national median, and home prices are 95 to 172 percent higher, depending on the city and data source (see Appendix Table A1).

[2] The Housing Affordability Institute, Mercatus Center, and Othering & Belonging Institute all regularly report on reforms. Our concept of "permitting" includes related processes that some jurisdictions call entitlement or planning.

and structures. Unlike in "cap-and-trade" pollution markets, therefore, economists lack an explicit permit price that would reveal compliance costs of marginal projects.

This paper measures permitting costs by studying an implicit market for development approval. In Los Angeles County, landowners may secure all necessary permits for a proposed project before sale, transferring the bundle of land and "ready-to-issue" (RTI) permits to developers. In theory, the premium paid for this bundle over raw, unpermitted land reveals developers' marginal willingness to pay to skip the permitting process. The RTI market is economically significant and mature: In some neighborhoods, it accounts for one in four land sales and one in ten properties suited for redevelopment ("likely teardowns"), with $353 million in transaction volume in 2024.[3] Exploiting this market, we estimate the approval premium, the impacts of preapproval on time-to-build and project completion, and permitting's share of the total regulatory tax on new housing.

To study the RTI market, we assemble a dataset of properties for sale advertised on the Multiple Listing Service (MLS) for Los Angeles County from 1995 to 2024. Our data contain 95,724 unique land and likely-teardown parcels, of which we observe 5.3 percent (5,092) are listed with and without permits. We use large language models (LLMs) to classify permit status from the text descriptions that real estate agents provide to the MLS, and we validate the LLM output with thorough human review.[4] Finally, we link listings to permits, property tax assessments, and zoning rules, allowing us to estimate time-to-build effects and perform supplementary analyses.

We begin by characterizing the permitting process in Los Angeles and the city's land market. Relative to other U.S. cities, permitting in Los Angeles appears burdensome: We estimate that building a midsize apartment complex takes around twice as long there as in Raleigh, NC or Fort Worth, TX. On average in Los Angeles, time-to-permit accounts for 40 percent of total time-to-build. Analyzing preapproval, we find it is more common on smaller lots and in denser neighborhoods, and it appears driven by specialized investors who acquire raw land specifically to permit and resell. While these facts lend support to our interpretation the approval premium as an equilibrium price of permits, they also raise the possibility of non-random selection into preapproval.

Our main empirical strategy is a repeat-listing difference-in-differences design. Among properties that are repeatedly listed, we compare price changes that coincide with changes in permit status to price changes on properties that remain raw land. By removing time-invariant characteristics, this design addresses a common challenge in cross-sectional hedonic regressions (Greenstone, 2017): unobserved fixed attributes that affect both prices and selection into preapproval. The capitalization of anticipated approval poses a second challenge (Bishop and Murphy, 2011; Bajari et al., 2012). If properties pay permitting costs before approval, land prices should rise in advance,

[3]See Appendix Figure A1 for a time series.

[4]We adopt emerging best practices for using LLMs in economics (Ludwig et al., 2025), and our work relates closely to earlier research that leveraged listing text data to measure housing quality (Nowak and Smith, 2020).

potentially contaminating the repeat-listing design. We adopt an event-study design to account for this anticipation effect.

We find a substantial approval premium. In the repeat-listing design, permit approval raises the price of vacant land by 50 percent on average. The cross-sectional design yields a similar estimate. In dollar terms, the premium amounts to \$48 per square foot of land for the average property, or approximately 36 percent of construction cost. For parcels with pre-existing structures, approval premia are smaller than those on vacant land in percentage terms but comparable in dollars, particularly if we reweight the data to be more similar on observables. Sensitivity analyses show our estimates are robust to time-varying controls for neighborhood characteristics and other listing information. Finally, we show permit arrival is partly anticipated, with a rise in listing mentions of pending permits two years before approval, coincident with the time path of land-price capitalization. However, anticipation bias in our approval premia appears negligible, as most land sales occur several years before permit news arrives.

Linking listings to permits, we also study the effects of preapproval on remaining time-to-build using a semiparametric hazard model. As we do not have a repeat-listing design for this outcome, we rely on comparisons of observably-similar projects with and without approved permits at the time of site acquisition. We find that preapproved projects are 8 to 12 percentage points more likely to be completed within four years of site acquisition, relative to a counterfactual four-year completion rate of 35 percent. To probe selection bias, we vary the set of control variables. Sensitivity analyses suggest that preapproval robustly increases the probability of completing quickly (i.e., within three years), whereas the impacts on ever completing are more sensitive to controls. Benchmarking our estimates to the average time-to-permit of preapproved properties, we find preapproval is effectively a one-to-one transfer of waiting time from developers to landowners.

To interpret these results, we propose a simple equilibrium model of housing development, wherein developers choose how much housing to supply and how much effort to invest in permitting speed. We obtain three results from the model. First, the approval premium can be expressed as the sum of two objects we call "pure wait" and "capitalized hassle," reflecting respectively the time and resource costs of permitting. Second, the effect of preapproval on remaining time-to-build allows us decompose the approval premium into these objects. Third, we relate permitting costs to the overall housing cost wedge, allowing us to contextualize permitting's importance within land-use regulation more broadly. The third result shows the permit share of the wedge can be measured either exclusively with land-price data or by the Glaeser and Gyourko (2003) house-price approach, which allows us to relate the approval premium to the city's overall regulatory burden.

Leveraging the model, we reach two conclusions about permitting in Los Angeles. First, at standard discount rates, pure wait accounts for much of the approval premium, suggesting that time-based measures of permitting can provide useful insight into its cost (Gabriel and Kung, 2025; Lyons and Sweeney, 2025; Manville et al., 2023). Second, on citywide average, permitting explains around one third of the gap between housing prices and construction costs, making it a key regulatory cost on development in Los Angeles. For such analyses, we exploit the size and richness of our listings data to compute detailed estimates of housing-price premia over construction costs, adapting methods from Glaeser and Gyourko (2003, 2018) and Gyourko and Saiz (2006). Overall, our analyses reveal rich and important impacts of permitting on housing and land markets.

In doing so, this paper contributes to literatures on regulation in public finance, urban economics, and real estate economics. An outstanding issue in this area has been to connect two distinct empirical approaches to land-use regulation. The first, beginning with Glaeser and Gyourko (2003), computes regulatory taxes as "top-down" wedges between housing prices and physical construction costs.[5] While this approach documents the presence of such wedges across diverse urban contexts, it typically treats the internal composition of the wedge as a "black box."

A second tradition has sought to open this black box by compiling indices of land-use regulation across jurisdictions (Quigley and Rosenthal, 2005; Gyourko et al., 2021; Bartik et al., 2024). These indices typically measure de-jure rules and questionably capture the de-facto reality of bureaucratic hassle. Also in this "bottom-up" tradition are studies of land-market capitalization of specific land-use regulations (e.g., Fu and Somerville, 2001; Brueckner and Sridhar, 2012; Turner et al., 2014). In related work, researchers have also examined the effects of regulations on development (e.g., Mayer and Somerville, 2000; Kahn et al., 2010; Jackson, 2016; Anagol et al., forthcoming; Büchler and Lutz, 2024; Kulka et al., 2024). Perhaps most closely related to our paper is Diamond et al. (2025), which infers regulatory burdens from developer behavior within a structural model. Our Los Angeles "ready-to-issue" setting provides complementary market-based evidence about de-facto compliance costs that are plausibly important but challenging to measure.

Methodologically, our analysis builds on a tradition of hedonic methods for valuing non-marketed attributes (Rosen, 1974; Greenstone, 2017). Our repeat-listing design fits within a branch of that literature which has used parcel-level panels, mostly for the valuation of environmental disamenities (Palmquist, 1982; Kohlhase, 1991; Davis, 2004; Buck et al., 2014; Moretti and Wheeler, 2025).[6] Our study also exhibits two applications of machine learning and LLMs to hedonic estimation: isolating attributes that existing approaches (e.g., keywords) capture with considerable noise, and soaking up once "unobservable" potential confounds (Nowak and Smith, 2020).

Finally, we relate to recent research on permitting processes for real estate and infrastructure. Manville et al. (2023) and Kestelman (2025) study permitting reform in Los Angeles, measuring its net impacts rather than total permitting costs. Our time-to-build analysis builds on efforts to

[5] Subsequent contributions include Glaeser et al. (2005a,b), Gyourko and Molloy (2015), Glaeser and Gyourko (2018), and Gyourko and Krimmel (2021).

[6] Our approach engages with recent advances that have stressed the advance capitalization of anticipated events and changes over time in the hedonic pricing function (Bishop and Murphy, 2011; Bajari et al., 2012; Banzhaf, 2021).

measure the speed of approval and development (Lyons and Sweeney, 2025; Gabriel and Kung, 2025) and connects to studies on the causes of rising infrastructure costs (Brooks and Liscow, 2023; Liscow, 2025). Where appropriate, we discuss how the richness of our data and context can inform the measurement of permitting costs in other real estate markets.

The paper proceeds as follows. Section 2 describes our data and setting in Los Angeles. Section 3 provides a model of permitting as a foundation for the empirical analysis. Sections 4 and 5 respectively estimate the approval premium and the effect of preapproval on remaining time-to-build. Section 6 relates our findings to the overall housing cost wedge. Section 7 concludes.

2 Setting and Data

This section introduces the setting and describes the data. We first explain the regulatory environment for real estate development in Los Angeles and the key features of the market for property with preapproved, or "ready-to-issue" (RTI), permits. We then review the data sources and the measurement of permit status. Finally, we present summary statistics, providing initial evidence of the approval premium and the role of investor intermediaries in preapproval.

2.1 Regulating Development in Los Angeles and Other U.S. Cities

Across U.S. Cities. While Los Angeles is not representative of the U.S. as a whole, the city does typify a class of places where housing affordability has become a defining policy challenge: places that have sustained high rents without growth. Across the 50 most-populous U.S. metropolitan areas as of 2025, none have both high rents and substantial recent development, defined as more than five new units per 1,000 residents per year on average over the preceding decade. By contrast, several metro areas have sustained faster growth at lower rents (see Panel A of Figure 1).

To investigate why housing production remains weak despite high prices, we examine the duration of the development process itself. Housing's "time-to-build" plausibly depends upon several factors: physical conditions like terrain, economic conditions like the construction labor market, and the regulatory environment. We explore these sources of time cost using permit microdata from 12 cities that we collected and harmonized, allowing us to measure the time from permit submission to project completion for observably-similar developments across cities.

In particular, we estimated the time-to-build for a standardized mid-size apartment building (see Panel B of Figure 1). This project would take on average 4.2 years in Los Angeles County, roughly twice as long as in Raleigh, NC and Fort Worth, TX. Time-to-build also appears high in other high-rent U.S. cities, such as New York and San Francisco.

Decomposing time-to-build into permitting and construction components, we find permitting directly accounts for approximately one-fifth of the total across-city variance in time-to-build (see

Figure 1: Housing Prices, Production, and Time-to-Build Across U.S. Cities

(a) Rents Versus Permits Across Metro Areas (b) Standardized Time-to-Build

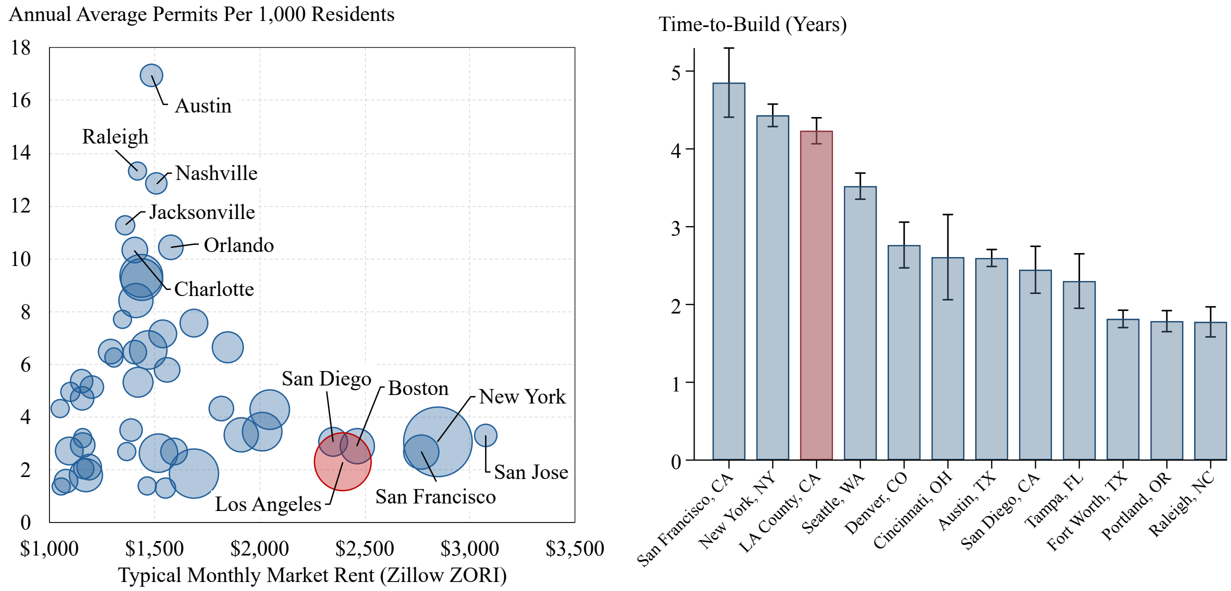

Notes: Panel A presents a scatterplot of average annual residential building permits versus monthly market rents across the 50 largest U.S. metropolitan areas by 2015 population. Building permits are averaged over 2015 through 2025 using the Census Building Permits Survey, then divided by 2015 population. Rents are measured using the Zillow Observed Rent Index (ZORI), which captures rents for "typical" units (quality-adjusted, between the 35th and 65th rent percentiles). MSA-average ZORIs over November 2015–2025 are used. Point sizes reflect 2015 population. Panel B presents estimated mean time-to-build (permit submission to the issuance of a certificate of occupancy) for a standardized apartment building (30 units, 2019 permit submission, 6,000 people per square mile tract density). Appendix C reviews the sample coverage and procedures used to clean and standardize the permit data.

Appendix Figure A2). The covariance of time-to-permit and remaining time-to-build (i.e., from permit issuance to completion) accounts for another third of the variance. The remaining variation reflects differences in construction time after issuance.

This variation in time-to-permit, as well as its covariance with remaining time-to-build, is highly suggestive of a role for regulation. To be sure, these patterns could also reflect physical or economic sources of delay, such as challenging terrain or tight labor markets that slow approval and construction. Indeed, the plausibility of these distinct interpretations makes evident the fundamental limitations of across-city comparisons.

Regulation in Los Angeles. Proposed developments in the City of Los Angeles face two main stages of review: a planning (or "entitlement") stage and a permit stage. The first is ostensibly focused on the proposal's compatibility with nearby land use but is ultimately discretionary. The second is ministerial and strictly technical (i.e., "by right"). We refer to both stages as "permitting."

Planning review is required whenever a project proposes 50 or more residential units or requests zoning variances. Because they frequently request variances, roughly three quarters of projects with 20 to 49 units also face a planning review (O'Neill et al., 2019a). While the Department of City Planning leads this process, their decisions can be appealed to one of seven regional Area Planning Commissions (APCs). These commissions are composed of political appointees who hold public hearings, acting as a venue for community voice and opposition.

Furthermore, around three in four developments with five or more units face some form of review under California Environmental Quality Act (CEQA), the state's environmental statute. One in five of these reviews are themselves appealed (O'Neill et al., 2019b). Proposals can also require approval from the California Coastal Commission (Kahn et al., 2010) or local utility boards (Gabriel and Kung, 2025). Surveys of California planners suggest that proposals are often abandoned, not merely delayed, at one of these review stages (Mawhorter and Reid, 2018).

Within-County Variation in Los Angeles. We estimate that, for the median project, the time from site acquisition to permit issuance accounts for 40 percent of total time-to-build. Panel A of Figure 2 presents this fact using two binned scatter plots of project-level data. In red diamonds, we compare the project's time from site acquisition to permit submission versus its total time-to-build; in blue dots, we compare a project's time-to-permit to its time-to-build.[7] A project that requires five years to complete spends around six months in pre-filing preparation and additional 18 months under review before issuance. By measuring the project start at site acquisition, we account for time spent preparing to file permits, as well as the time in unobserved regulatory processes.

Time-to-build for our standardized apartment building also varies considerably by neighborhood within Los Angeles County (see Panel B of Figure 2). In its densest and most desirable areas, such as Santa Monica and Westwood, such projects typically take around five years. Yet in other neighborhoods, time-to-build is comparable to Raleigh and Fort Worth. These patterns suggest substantial variation in permitting costs, both across and within cities.

2.2 Preapproval in Los Angeles

Our focus is on a segment of the Los Angeles market in which properties have "ready-to-issue" (RTI) permits. In this submarket, landowners sell not only the physical property, but also a vested legal right to build a specific project that is transferable at sale.[8] Landowners obtain these rights by taking a complete set of building plans fully through their local permitting process.

Upon closing, land buyers can "pull" permits and begin construction immediately, without

[7] Our across-city analyses use permit data only, due to data limitations. Within Los Angeles, our permit-matched listings data allow us to use site acquisition as the start date for time-to-build.

[8] While the RTI market's scale likely reflects high permitting costs in Los Angeles, it also relates to a larger shift in the U.S. homebuilding industry toward specialized intermediaries that bear the same regulatory risk. For example, the two largest U.S. homebuilders control about 80 percent of their combined acreage through options intended for exercise after entitlements are obtained, according to their 2024 10-K filings.

Figure 2: Permitting and Development in Los Angeles

(a) Time-to-Permit Versus Time-to-Build

(b) Years to Complete Apartment Buildings

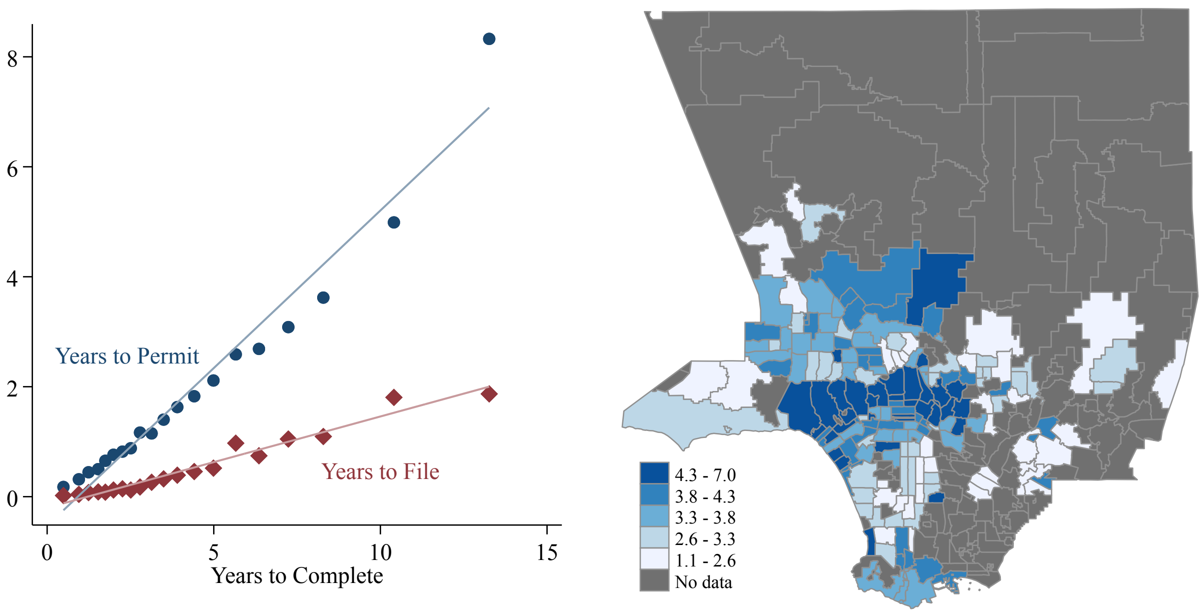

Notes: Panel A displays binned scatterplots of the time from site acquisition to building permit filing (red diamonds) and from acquisition to permit issuance (blue circles) against the total time-to-build (acquisition to completion). All durations are measured in years. The sample is of all completed developments that began with a site acquisition. Acquisition refers to the closing date on the land sale. Completion refers to the issuance date for the certificate of occupancy. Panel B displays the estimated mean time (by zipcode) from land acquisition to completion for a 30-unit multifamily apartment building. See Appendix C for estimation details.

any further approvals. However, this right is tied to a precise project specification, including the building footprint, floor plans, and exterior materials. While the buyer is free to pursue alternative plans, doing so would forfeit the preapproved status and trigger new reviews. The option value lost under preapproval implies the approval premium would understate permitting costs.

Institutional context suggests that the value of an approval primarily reflects time and resource cost savings, rather than direct financial transfers to the government. Permit fees paid before issuance are modest (around 1.2 percent of raw-land value, see Appendix C for this calculation). While substantially larger impact fees are typically required at the moment of issuance, these remain the obligation of the buyer and thus, in principle, would not be capitalized into the implicit price of a "ready-to-issue" permit.

To obtain a permit, developers also incur predevelopment expenses, such as architecture and engineering fees. A portion of these costs would likely be incurred even in the absence of permitting regulations. Using R.S. Means Company data, we estimate predevelopment imposes average resource costs of 12.6 percent of raw-land value. The time costs associated with these activities appear modest in Panel A of Figure 2, where we assume that the observed time-to-file

Figure 3: The When and Where of Preapproval

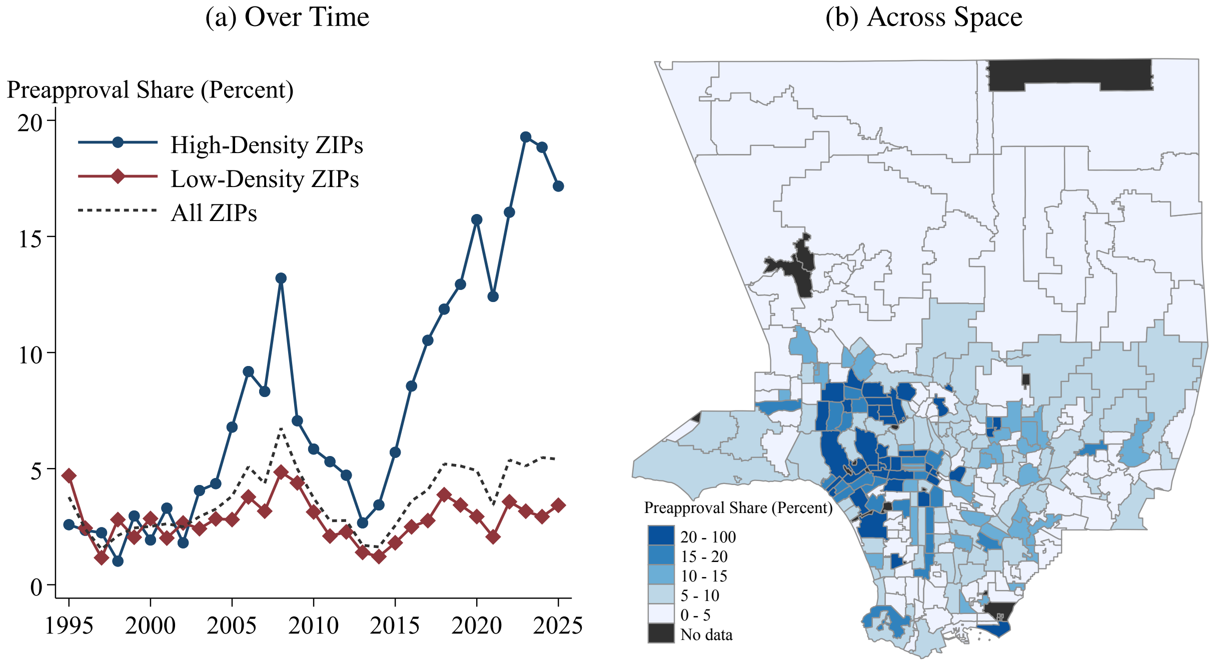

Notes: This figure displays the share of vacant land listings or sales with preapproved permits, over time (Panel A) and by zipcode (Panel B). Panel A splits zipcodes at the median population density. This figure uses the "narrow" definition of preapproval. Appendix Figures A3 and A4 respectively show maps and time series of preapproval rates that include incomplete permitting, and for nonvacant properties.

reflects exclusively necessary architecture and engineering work. Overall, the value of approved plans appears to lie in the approval and not in the plans. Section 5.3 considers other explanations of the approval premium.

Preapproval is also used on properties with existing structures, although it is more common on vacant land. For nonvacant properties, the content of preapprovals varies from minor additions to an existing structure (e.g., a pool or carport), to accessory dwelling units (ADUs), to a new primary dwellings. For both vacant and nonvacant properties, the scale of preapproved primary dwellings ranges from single-family homes to large apartment complexes.

Figure 3 shows the prevalence of preapproved listings over time and across space. From 1995 to 2025, preapproval became increasingly common in dense areas, approaching one in four land listings. Geographic variation in preapproval for non-vacant properties resembles that of vacant land, consistent with linkages between the land and teardown markets (see Appendix Figure A3).

2.3 Data

Property Listings. Our primary dataset is the universe of property listings from 1995 to 2024 recorded on the Multiple Listing Service (MLS) for Los Angeles County, obtained via Cotality (previously CoreLogic). We split the data into vacant land and nonvacant properties, that is, those with existing structures when listed.[9] The MLS data contain each property's administrative parcel number (APN, a tax identifier), dates and prices of listings and sales, the location (street address and geographic coordinates), and basic characteristics (e.g., lot size and floorspace).

Central to our analysis is the MLS "remarks" field, a description provided by the real estate agent. Listing portals, such as Zillow and LoopNet, make the remarks highly visible to the public. Remarks are typically paragraph-length summaries of properties' desirable attributes. For vacant land, remarks often discuss sites' physical characteristics (topography and views), proximity to amenities, utility hookups, zoning classification, and permit status.

We process the remarks using a variety of machine-learning methods, detailed below, to determine the permit status of properties and measure other attributes. In a key precursor to our analysis, Nowak and Smith (2020) use characteristics extracted from MLS remarks to adjust repeat-listing home-price indexes for quality improvements, and they show that the presence of inferred improvements explains relative appreciation of the same house.

Building Permits. We have collected building permit microdata for Los Angeles County. For the City of Los Angeles, permits came directly from the City's open-data portal. For all other localities, we obtained permit records from the construction-data company Shovels. For our time-to-build analysis (Section 5), we study only new-construction permits, dropping additions, alterations, and site work (demolition and grading) that did not result in a new building.

We used APNs to match listings to permits, a match enhanced by the County's administrative records on parcel splits and mergers over time. To validate match quality, we confirm that permit timing aligns with the expected sequence of development (see Appendix Figure A5). For properties coded as preapproved, we find permits are typically submitted prior to to the listing or sale date, opposite to the pattern in properties without preapprovals. However, matching permit records to property listings is notoriously difficult and occasionally yields "wrong-sequenced" timings.

Other Data. Supporting analyses use historical property tax assessments (2009–2024, from the Los Angeles County Assessor, via Cotality), harmonized zoning classifications (as of 2019, from the Southern California Association of Governments), and neighborhood characteristics (from the 2023 tract tabulations of the 5-year American Community Survey). Appendix C provides further information on these supplementary data sources.

[9] For our vacant data, properties must be reported to the MLS as a "lot/land" listing. We also require these properties have zero floor area, be an arm's-length transaction, and have a sale or listing price of at least $50,000. For our nonvacant data, we also only retain parcels with any mention of redevelopment potential in the listing description. See Appendix Table A2 for a step-by-step documentation of sample restrictions.

Table 1: Classification of Permit Status Based on MLS Remarks

Status	Definition and Example
Ready to Issue	*Definition:* The description explicitly states all necessary plans and permits are fully approved, paid for, and issued. *Example:* Street to street lot with city lights and mountain views. Price includes plans for 3,400sq. ft. post & beam home. Planning permission, Mulholland Scenic Hwy permission. Sewer and street improvement engineering plans. Ready to build immediately.
In Progress	*Definition:* The description shows clear evidence of active engagement with a planning department, such as submitted plans or applications that are awaiting corrections or approval. *Example:* Multi Million Dollar Homes Adjacent-Great Views! Water and Electric in Street. Has Conceptual House Plans. Property is close to final geology and grading plan approval.
Not Ready	*Definition:* The description lacks any definitive permit status, contains hard negatives ("expired plans"), or relies solely on describing potential. *Example:* Spectacular upslope site overlooking Lake Hollywood with big future potential. Raw site with no utilities or street improvements. 8350 square feet. Seller may assist. Motivated seller.

Notes: This table presents our classification scheme for permit status, which uses the remarks field of property listings. For each status, we quote the definition provided to the LLM, as well as the remarks text for a characteristic example.

2.4 Inferring Property Characteristics from Listing Remarks

Table 1 shows our classification of listings into three categories: ready to issue (fully approved), permitting in progress, and not permitted (no apparent progress on permits). For each category, we give the definition and the MLS remarks for an example lot.

We took two approaches to implement this classification at scale. First, we applied a naive keyword-based approach, using a restrictive dictionary of keywords to flag listings as ready-to-issue (e.g., "fully permit" or "approved plans") if any keywords were present and otherwise treating listings as not permitted. There is no "in-progress" category for the keyword classification. Second, we used a large language model (Open AI's GPT-4.1, 2025-04-14 release) to classify the listings and to extract additional information. Unless explicitly noted, we use the LLM-based measure of permitting status. Appendix C includes the keyword list and LLM prompts.

We took several steps to improve and assess LLM output quality. First, we refined our definitions through iterative piloting and human review of LLM output. Second, we elicited "chain-of-thought" reasoning and a subjective confidence rating for each listing, used in later analyses. Third, the prompt also gave instructions about ambiguous or conflicting text. Fourth, we employed humans to hand-review a sample of approximately 7,400 listings, yielding benchmarks for LLM performance. We report results from data-quality checks below and in Section 4.

Using the LLM, we also elicited a summary of the "content" of any development approvals, both as a categorical variable and as free text. Our (non-exclusive) categories are: additions to an existing structure, ADUs, single-family primary dwellings, multi-family primary dwellings, or other structures. For only the nonvacant properties, we also obtain a "value proposition" categorical variable: whether the remarks emphasize the land's existing use, its redevelopment potential, or a hybrid of these factors. We use this variable to retain all listings of nonvacant properties if they are ever, across any listing, classified as hybrid or redevelopment value.[10]

We summarize disagreement between our classification methods through "confusion matrices." First, we find the keyword approach is highly conservative. In the vacant sample, keywords identify only 69 percent of properties the LLM classifies as fully preapproved, and just 26 percent of properties the LLM classifies as permitting in progress (see Appendix Table A5). In the other direction, the LLM classifies as not permitted just 3 percent of properties the keyword approach flags as preapproved, with the remainder split between fully preapproved and in-progress classifications. Results are similar for the nonvacant sample. Second, we find strong agreement between humans and LLMs (see Appendix Table A6): For properties the LLM classifies as fully preapproved, humans agree in 92.2 percent of cases, although more disagreement exists for "in-progress" approvals.

2.5 Summary Statistics

Table 2 presents summary statistics on property and neighborhood characteristics of our data, split by permit status and the presence of an existing structure on the property.

Our main dataset contains 95,724 unique properties, of which 34,321 (36 percent) are vacant land at listing and 61,403 have existing structures at listing. Due to repeat sales and listings of the same property, our data contain considerably more observations (186,462) than unique properties. We observe 2,667 vacant properties and 4,379 nonvacant properties that are preapproved, which amounts to 7.7 percent of vacant properties and 7.1 percent of nonvacant properties.

The sample medians in Panel A immediately suggest an approval premium. Among vacant properties, the median price of approved land is \$948,000 (in constant 2024 dollars), as compared to \$299,000 for properties without permits or permits in progress. Notably, the approved properties

[10] Appendix Table A3 provides definitions and illustrative examples. In the nonvacant sample, 1.3 percent of listings are classified as hybrid or redevelopment value, but these listings contain 97.7 percent of preapprovals (see Appendix Table A4). Our sample restriction is intended to exclude properties for which the real option to redevelop is irrelevant and thus which almost never pursue preapproval.

Table 2: Summary Statistics by Permit Status and Presence of Existing Structure

	Vacant Land		Nonvacant Property	
	Non-RTI	RTI	Non-RTI	RTI
	(1)	(2)	(3)	(4)
Panel A: Median Property Characteristics at Listing				
Price	299,000	948,000	772,000	1,150,000
Price Per Sqft. Land	15.2	71.3	96.2	131.8
Living Area (sqft.)	0	0	1,743	1,797
Lot Area (sqft.)	21,880	9,975	7,410	7,500
Panel B: Median Neighborhood Characteristics				
Population Density	2,755	7,190	9,998	9,293
% Poor	8.9	8.9	10.4	9.5
% Non-Hispanic White	43.4	43.4	25.0	33.4
% Hispanic	34.2	24.7	35.2	26.2
% College Graduate	35.4	50.5	37.4	46.5
% Renters	27.7	37.2	44.9	42.2
% Vacant Units	6.9	6.7	5.7	6.7
Avg. Land Price ($/sqft.)	110.0	156.1	144.3	156.1
Panel C: Mean Structure Characteristics in Permit				
Floorspace (sqft.)	9,188	7,903	5,731	3,883
% Any Residential	60.3	63.8	21.1	22.1
% 5+ Units	6.3	11.4	3.8	4.6
Construction Cost	981,952	1,147,015	924,381	527,573
Height (ft.)	29.0	29.5	25.3	24.0
Full Sample				
Observations	63,672	3,705	113,253	5,832
Unique Properties	31,678	2,643	57,035	4,368
Matched to Permit				
Observations	18,390	2,390	61,703	4,006
Unique Properties	9,980	1,697	28,994	2,990

Notes: This table reports summary statistics on the characteristics of properties, splitting according to permit status and whether it has an existing structure on the property at the time of listing. Panels A–B report unweighted medians. Panel C reports unweighted means among the matched sample with positive permit-reported construction costs. The final panel provides counts of observations and unique properties.

are smaller on average, yielding striking differences in prices per square foot of land. The same premium appears for properties with existing structures.

Summary statistics on the neighborhood characteristics (Panel B) indicate that, for vacant land, preapproved properties are generally in areas with higher population density, levels of schooling, and land prices than not-preapproved properties. Among nonvacant properties, however, we do not find the same differences in neighborhood characteristics.

There are also key differences between between vacant and nonvacant properties in the content of any preapproved permits, as inferred from both the listing text (see Appendix Table A7) and submitted permits (Panel C). These two sources provide complementary data about projects. Also, the permit data are reported by developers but reflect only the matched subsample.

Using the listings for vacant properties, 97 percent of preapprovals pertain to primary dwelling units, of which 38 percent are for multiple dwellings. By contrast, only 55 percent of preapprovals on nonvacant property pertain to dwellings, with the remainder concentrated in additions to existing structures or ADUs. These differences suggest one should not anticipate similar approval premia across vacant and nonvacant listings on average, either in percent or dollar terms.

Panel C of Table 2 reports the characteristics of structures in permits submitted for the properties. These averages apply to the subsample of listings that can be matched to permits, and they importantly include post-listing permits, unlike Appendix Table A7. For both vacant and nonvacant properties, preapproval is consistently associated with more-intensive land use.[11]

We have also examined the characteristics of switchers and non-switchers of permit status in our data (see Appendix Table A9).[12] Switchers are somewhat distinct from all approved or all land listings, having smaller lots and lower per-square-foot land prices before approval. Such differences in switchers from populations of interest may raise external-validity concerns for our repeat-listing specification, which exploits switchers to identify the approval premium. To address this concern, we re-estimate premia in samples reweighted for observable differences (see Appendix Figure A9).

2.6 Preapproval and Intermediaries in Property Markets

Why do landowners sell preapproved properties rather than develop themselves? Two analyses suggest most preapprovals are by investors specialized in this activity, acquiring raw land specifically to permit and resell, and thereby intermediating between long-term landowners and developers seeking "shovel-ready" sites (consistent with Lovo and Spaenjers, 2018).

First, we show the preapproved share of sales follows a "hump" shape with respect to holding period (see Panel A of Figure 4). The share is very low for holding periods under six months, likely too short to complete permitting, and peaks around a holding period of three years, consistent with permitting timelines but not long-term land investment. Using preapproval rates at long holding periods as a baseline, we estimate specialists are involved in approximately 71 percent of preapproved new primary dwellings (i.e., excluding ADUs, pools, etc.).[13]

[11] Appendix Table A8 examines preapproval determinants by multivariate regression and reaches similar conclusions.

[12] Following Miller et al. (2023), we also calculate the implicit weights on observations in our repeat-listing design. Appendix Table A9 therefore further reports the average characteristics of the "identifying" population.

[13] Among vacant properties, 7.4 percent of sales with holding periods under ten years are preapproved, versus 2.1

Figure 4: The Who of Preapproval

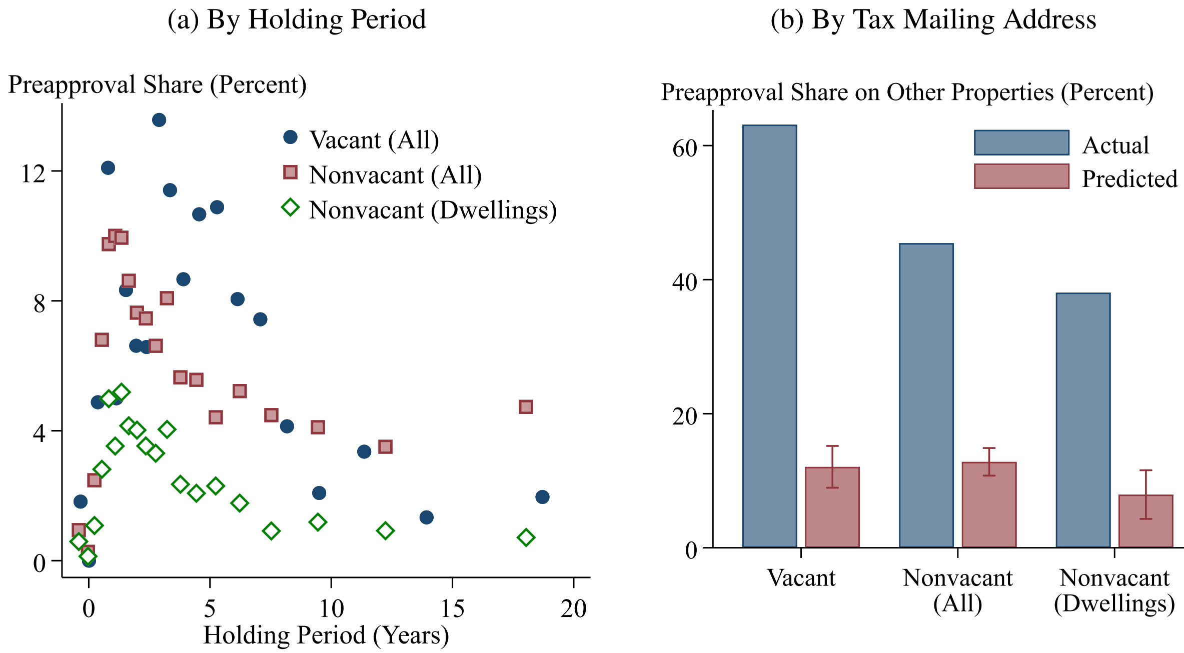

Notes: This figure examines the role of specialized intermediaries in preapproval. Panel A shows a binned scatterplot of the share of sales with preapprovals by holding period, separately for vacant properties (blue solid dots), all nonvacant properties (red squares), and nonvacant properties for primary dwelling approvals only (green diamonds). Panel B shows, for preapproved listings, the preapproval rate among *other* listings linked to the same tax mailing address (in blue). We also show (in red) the predicted preapproval rate from a regression of preapproval on fixed effects for the year of listing, zipcode of the listed property, and zipcode of the tax mailing address. Results are shown for vacant properties as well as for the two versions of the nonvacant-property analysis. Bars reflect 95-percent confidence intervals with two-way-clustered standard errors clustered by property and zipcode of the tax mailing address.

To further distinguish these permitting specialists, we identify sellers using property tax mailing addresses in the three years prior to listing. Following a recent literature on institutional real estate investors (Ganduri et al., 2023), this approach allows us to "see through" legal entities with the same ultimate owners. Appendix C provides futher details on this data linkage.

We find that, on average, sellers of preapproved land preapproved 61 percent of their other sales. This preapproval share is 49 p.p. more than predicted using characteristics of the sold property and of the tax mailing address (see Panel B of Figure 4). Under the assumption that non-specialist sellers preapprove at the predicted baseline rate, the excess 49 p.p. is attributable entirely to specialists, implying a specialist share of around 80 percent ($0.49/0.61 = 0.80$), close to the holding-period estimate. A similar share applies for nonvacant properties.

This seller specialization has two main implications for our analysis. First, and reassuringly, it suggests the approval premium represents an equilibrium return on a permitting investment,

percent of those with holding periods exceeding ten years. Short-hold transactions account for 97.1 percent of all transactions in the sample. If no specialist holds for more than ten years, the share of preapprovals attributable to specialists is $[0.971 \times (0.074 - 0.021)]/[0.971 \times 0.074 + 0.029 \times 0.021] \approx 0.71$.

alleviating the potential concern that many preapproved-land sales may be distressed sales of troubled projects. Second, the role of specialists suggests unobserved property-level heterogeneity, with two distinct consequences. Heterogeneity in levels (i.e., specialist selection of higher- or lower-quality land) poses a threat to the internal validity of cross-sectional estimates of the approval premium and time savings. Heterogeneity in gains (i.e., specialist selection of land with higher potential upside) suggests that our research design captures a local effect for high-return projects, posing an external-validity issue. We examine both selection concerns in Sections 4 and 5.

3 A Model of Regulated Development in Equilibrium

This section proposes a theoretical model of developer choice over permitting effort and housing capital investment. We use the model to interpret the two key empirical objects: the approval premium and the effect of preapproval on time-to-build. Finally, we explicitly derive our estimating equations from the model.

3.1 Environment

The housing market is populated by developers and households.

Land and Developers. Units of land ("parcels") are indexed by $j = 1, \ldots, J$. Each parcel is subject to a quantity limit $\bar{h}_j \in [0, \infty)$ on housing capital. Land exists in one of three states: raw (undeveloped and without an approved permit), approved but undeveloped, or developed.

Developers buy, sell, hold, and improve land, trading it in a perfectly competitive market. They are risk-neutral and maximize present-value profits using a uniform per-period discount factor $\beta \in (0, 1)$. If they hold a raw land parcel, they choose the intensity of permitting effort, expressed as the arrival probability of a permit $q \in [0, 1]$ per period, subject to an increasing flow cost $c_j(q)$. We treat these costs as proportional to the continuation value of the project.[14]

Once permits are approved, the developer chooses housing capital $h_j \in [0, \bar{h}_j]$ subject to a one-time cost function $k_j(h)$. Development itself is instantaneous. Once built, developments yield a present-value income $R_j(h) = r_j(h)h$, where r_j is the capitalized rent price.

To support a well-defined stationary equilibrium, we assume that an exogenous fraction $\delta \in (0, 1]$ of developments, selected uniformly at random, has their capital fully demolished each period. This land reverts to the raw state but otherwise remains available in the market.

Households. Households, indexed by $i \in [0, 1]$, maximize period-utility functions $u(j; i)$ each period by choosing developed parcels $j(i)$. Each inelastically demands one unit of housing at their

[14] This "iceberg" formulation captures soft costs (e.g., legal fees) and regulatory scrutiny that roughly scale with project value. It can also represent endogenous project completion risk.

chosen parcel. Period utility takes a quasilinear form over parcels:

$$u(j;i) = \alpha_j - r_j + \varepsilon_j(i), \tag{1}$$

where α_j denotes parcel j's exogenous amenity and $\varepsilon_j(i)$ is household i's idiosyncratic taste for j. Tastes are draws from an arbitrary distribution, yielding a demand system for all locations.

3.2 Land Valuation and Permitting Effort

We present and then solve the developer's optimization problem, yielding the value of approved and raw land, as well as the optimal level of permitting effort.

Approved Land. After permit approval, the developer's problem is static. Developers solve:

$$\pi_j(\bar{h}_j) = \max_h \left\{ R_j(h) - k_j(h) \right\} \quad \text{s.t.} \quad h \leq \bar{h}_j. \tag{2}$$

We let $\lambda(\bar{h}_j) = \max\{0,\ R'_j(\bar{h}_j) - k'_j(\bar{h}_j)\} \in [0, \infty)$ denote the shadow price of the quantity limit.

Raw Land. In each period before permit arrival, the developer chooses an arrival rate q to solve their Bellman equation:

$$v_{0,j}(\bar{h}_j) = \beta \cdot \max_q \left\{ q\pi_j(\bar{h}_j) + (1 - q - c_j(q))v_{0,j}(\bar{h}_j) \right\}, \tag{3}$$

which, when the developer plays their profit-maximizing strategy ($q = q^*$), yields the following value of raw land:

$$v_{0,j}(\bar{h}_j) = \frac{\beta q_j^*}{1 - \beta(1 - q_j^* - c_j(q_j^*))} \pi_j(\bar{h}_j), \tag{4}$$

where, for parsimony, we suppress the dependence of the developer's chosen arrival rate q^* on the quantity limit $\bar{h}$. Relative to approved land, the value of raw land is discounted for the time delay and permitting costs to be paid before development profits are realized.

Permitting Effort. We characterize the optimal level of permitting using the first-order condition for the developer Bellman equation (Equation 3):

$$c_j(q_j^*) = \frac{1 - \beta}{\beta} \cdot \frac{1}{\varepsilon_c(q_j^*) - 1}, \tag{5}$$

where the cost elasticity is defined to be $\varepsilon_c(q_j^*) = q_j^* c'_j(q_j^*)/c_j(q_j^*)$. Effort is decreasing in the discount factor (less costly to wait) and in the cost elasticity (similar to a markup).

Our key observation is that Equation 5 omits land-value objects, such as $\pi_j(\bar{h}_j)$ or $v_{0,j}(\bar{h}_j)$. It depends solely on the discount parameter and the effort cost elasticity. Section 3.4 refers to this observation in considering the vulnerability of our empirical results to unobservables.

3.3 Equilibrium Prices and Objects of Interest

We first state a lemma on equilibrium land prices. Its role is to link developer primitives to price data. Three propositions then derive our empirical objects of interest in the model. Appendix B contains all proofs and a formal definition of equilibrium.

Lemma 1. *Equilibrium prices of raw and approved land are given respectively by* $p_{0,j}(\bar{h}_j) = \frac{\beta q_j^*}{1-\beta(1-q_j^*-c(q_j^*))} \int_0^{\bar{h}_j} \lambda_j(s)\,ds$ *and* $p_{1,j}(\bar{h}_j) = \int_0^{\bar{h}_j} \lambda_j(s)\,ds$.

Lemma 1 establishes that land prices fully capitalize permitting cost in competitive land markets.

Proposition 1. *Relative to its counterfactual raw price, approved land trades at a premium of*

$$\theta_j(\bar{h}_j) = \frac{p_{1,j}(\bar{h}_j) - p_{0,j}(\bar{h}_j)}{p_{0,j}(\bar{h}_j)} = \frac{1-\beta}{\beta q_j^*} + \frac{c_j(q_j^*)}{q_j^*},$$

and we express the dollar-value premium as $\tau_j(\bar{h}_j) = p_{1,j}(\bar{h}_j) - p_{0,j}(\bar{h}_j)$.

Proposition 1 derives the approval premium. When developers are less patient ($\beta \to 0$), or when permitting is slower in equilibrium ($q^* \to 0$), the costlier is waiting, and thus the more developers would pay to skip the wait. Costlier effort ($c_j(q_j^*) \to \infty$) also raises the premium.

Our expression for the premium features two components, "pure wait" and "capitalized hassle." The former means that raw land should trade at a discount to approved land even in a model where permit arrival was exogenous and costless. The latter means that, when developers can exert effort to hasten permit arrival, the present value of the effort cost is capitalized into land prices.

Combining Proposition 1 with Equation 5 yields a markup rule for the approval premium: $\theta_j = \frac{1-\beta}{\beta q_j^*} \cdot \frac{\varepsilon_c(q_j^*)}{\varepsilon_c(q_j^*)-1}$. Inverting this expression allows us to estimate the cost elasticity of the permitting hazard, the primitive parameter in our model governing the relative importance of pure wait versus effort. The more rapidly diminishing are returns to effort ($\varepsilon_c \to \infty$), the more the premium reflects pure wait. Intuitively, the preapproved land market offers a "make-or-buy" decision: Developers make their own permit only if their effort cost is below the approval premium, such that the availability of preapproved land enforces a required return on effort.

Next, we turn to the second object of interest: the effect of preapproval on the cumulative hazard of permit arrival (here equivalent to "time-to-build" due to instantaneous construction). In the model, permitting effort and thus the permit arrival rate are both constant over time. The probability that a project on raw land is completed by period t is then $F_{0,j}(t) = 1 - (1 - q_j^*)^t$, whereas preapproved land is ready immediately ($F_{1,j}(t) = 1$ for $t \geq 0$). The effect of preapproval on time-to-build is the cumulative difference in distributions.

Proposition 2. *Reducing time-to-build through preapproval increases the present value of rental income $r_j(\bar{h}_j)$ by the following proportion:*

$$\frac{\sum_{t=0}^{\infty} \beta^t \left(F_{1,j}(t) - F_{0,j}(t)\right)}{\sum_{t=0}^{\infty} \beta^t F_{0,j}(t)} = \frac{1-\beta}{\beta q_j^*}.$$

Proposition 2 shows the source of the approval premium in a model where preapproval purely reduces wait (i.e., no effort cost of permitting).[15] The empirical analysis below uses Propositions 1 and 2 in two ways. First, given a calibrated discount parameter β, we can decompose the approval premium into pure-wait and capitalized-hassle terms, essentially by subtracting the time-to-build effect from the approval premium. Second, the discount parameter β is identified by the approval premium and the time-to-build effect only in a pure-wait model. With costly permitting effort, that ratio would also contain the period hassle cost $c(q_j^*)$. Interpreting our estimates in a pure-wait model should thus yield upwardly-biased estimates of discount rates.

Having established and decomposed the approval premium, we now connect it to the aggregate housing cost wedge. Following Glaeser and Gyourko (2003), we define this wedge as the gap between housing prices and physical construction costs.[16] We use our framework to decompose this wedge into permitting costs and a residual term that reflects the scarcity of housing in specific locations. This scarcity results from regulations (quantity limits $\bar{h}_j$) as well as the intrinsic scarcity of differentiated locations (idiosyncratic preferences $\varepsilon_j(i)$).

Proposition 3. *The housing cost wedge decomposes into permitting costs and raw land value:*

$$R_j(\bar{h}_j) - k_j(\bar{h}_j) = \tau_j(\bar{h}_j) + p_{0,j}(\bar{h}_j) = p_{1,j}(\bar{h}_j),$$

such that the permitting share of the wedge can be derived using only land-market data:

$$\frac{\tau_j(\bar{h}_j)}{R_j(\bar{h}_j) - k_j(\bar{h}_j)} = \frac{\theta_j(\bar{h}_j)}{1 + \theta_j(\bar{h}_j)}. \quad (6)$$

Proposition 3 means that, in a competitive land market, the equilibrium profit from developing approved land either offsets permitting costs or is capitalized into the price of raw land. This proposition raises two points that motivate our analyses in Section 6.

First, it provides a cost-accounting framework. Following the left hand side of Equation 6,

[15] Proposition 2 is a result about the gain in present-value rental income, whereas the approval premium in Proposition 1 pertains to land value. In analysis that combines these objects, we adjust for the ratio of rental income to land value. Due to construction cost, a one-percent increase in rental income raises land value by more than one percent. The appropriate adjustment is dampened to the extent that permitting also delays construction costs.

[16] In their quantitative implementation, Glaeser and Gyourko (2003) contrast the price and construction cost of single-family residences, so that average and marginal cost are definitionally equivalent.

one can divide the dollar-cost premium τ_j by the Glaeser and Gyourko (2003) estimate of the total wedge, derived from house prices and external measures of construction costs. However, a practical drawback of this approach is that it measures premia at site acquisition, home prices later if at all, and construction costs only through imputations. These multiple measurements raise the prospect of significant estimation error in permit shares of the wedge.

Second, the proposition provides an alternative approach to the permitting share of the housing cost wedge. Following the equation's right hand side, we can use exclusively land prices, dividing the approval premium (in percent) by one plus this premium. We proceed this way, using land prices to obtain a more-credible estimate of the permit share more credibly. We then scale up our shares by our estimates of the wedges ($(R_j - k_j)/k_j$) to report contributions to the wedge.

3.4 From Theory to Data

We close our theoretical analysis by connecting the model's equilibrium conditions to our empirical specifications. In particular, we state our identifying assumptions for the approval premium and time-to-build effect as formal moment conditions.

Approval Premium. In Section 4, our main approach to identifying the approval premium is a repeat-listing regression (Equation 12). To derive this strategy from the model, we express parcel-level changes in (log) prices over time as changes in fundamental value and in permit status:

$$\Delta \log p_j(\bar{h}_j) \;=\; \Delta \log \pi_j(\bar{h}_j) \;+\; \log(1+\theta_j) \cdot \Delta \mathrm{RTI}_j. \tag{7}$$

Identification of the approval premium θ_j requires that changes in permit status $\Delta\mathrm{RTI}_j$ are independent of other changes in fundamental value $\Delta \log \pi_j(\bar{h}_j)$, conditional on a set of control variables X_j. Formally, this moment condition is

$$\mathrm{E}[\,\Delta \log \pi_j(\bar{h}_j) \cdot \Delta \mathrm{RTI}_j \mid X_j\,] \;=\; 0. \tag{8}$$

Our cross-sectional specification (Equation 11) requires an arguably stronger assumption than Equation 8, in which fundamental value must be conditionally independent (in levels) of approval.

Time-to-Build Effect. The model can also be used to interpret the time-to-build effects in Section 5. To connect the model more closely to the data, we include an arbitary, exogenous distribution of construction time, rather than instantaneous completion after permitting. This modification allows us to redefine q as a completion hazard, consistent with Section 5. For simplicity, we also assume here the effort-cost function takes an isoelastic form, specifically:

$$c_j(q) \;=\; \psi_j q^{\varepsilon_c}, \quad \log \psi_j \;=\; \log \psi_{0,j} \;+\; \eta \mathrm{RTI}_j.$$

By Equation 5, the equilibrium hazard q^*_{jt} is:

$$\log q^*_j \;=\; \alpha_0 \;-\; \frac{1}{\varepsilon_c - 1}\left(\log \psi_{0,j} \;+\; \eta \mathrm{RTI}_j\right), \tag{9}$$

where α_0 collects model parameters that are constant across j.[17] By Equation 9, identification of the time-to-build effect $\eta/(\varepsilon_c - 1)$ requires the absence of selection into preapproval according to the effort-cost parameter $\psi_{0,j}$, conditional on controls X_j:

$$\mathrm{E}[\,\log \psi_{0,j} \cdot \mathrm{RTI}_j \mid X_j\,] \;=\; 0. \tag{10}$$

We offer two observations about this moment condition. First, in a richer model with developer heterogeneity in effort costs, it would be a high-level restriction on that heterogeneity. That is, the extent of potential gains from trade in permitting (which motivate preapproval) could not be related to the average cost of producing a permit across potential developers of the same property. An intuitive example would be if unobservably hard-to-permit properties were to select into preapproval (see Section 5). Second, the estimation of time-to-build effects does not, in theory, require assumptions on the profit function $\pi_j(\bar{h}_j)$, following from Equation 5.

4 The Approval Premium

This section estimates the approval premium. We first outline our empirical approach, progressing from a cross-sectional regression to our main repeat-listing design. We then present the results, first for vacant land, then for properties with existing structures, and finally heterogeneity in premia.

4.1 Methods

We take three approaches to estimating the approval premium. First, a cross-sectional hedonic regression yields a benchmark estimate. As this approach is vulnerable to selection on unobserved lot characteristics, our next (and preferred) approach is a repeat-listing difference-in-differences design that exploits changes in permit status for identification, thereby removing time-invariant unobservables. Third, to address dynamic concerns such as the anticipation of approval, we augment this approach with an event-study design.

Cross-Sectional Regression. We project lot-level land prices onto permit status (RTI_{it}) and other

[17] While Equation 9 specifies a proportional-hazards relationship between approval and the completion hazard, we estimate the hazards more flexibly using a series of horizon-specific linear probability models (see Equation 15).

observed characteristics (X_{it}):

$$\log p_{it} = \beta \mathrm{RTI}_{it} + X_{it}\gamma + u_{it}, \tag{11}$$

where the coefficient β is the cross-sectional approval premium. As controls, we include fixed effects for zipcode and year, along with the logarithm of lot area as a continuous covariate. We cluster standard errors by parcel throughout the paper and present inference adjusted for spatial correlation in robustness checks.

We measure preapproval using the keyword and LLM approaches, both of which are introduced in Section 2. For the LLM measure, we proceed in three ways: a "narrow" binary indicator (completely preapproved, versus not or in progress), a "broad" binary indicator (preapproval complete or in progress, versus not), and a "joint" specification with three categories (completely preapproved, permits in progress, or not).

Equation 11 is a useful benchmark but is clearly vulnerable to selection bias, as we discuss in Section 3.4. For example, landowners may be more likely to seek preapproval for properties with unobserved flaws to reduce buyer risk, or for unobservably valuable properties to maximize sale price, either of which would bias the estimate (Greenstone, 2017).

Measurement error in preapproval is another concern with Equation 11 and the repeat-listing analysis below. We leverage the human-validation sample to assess potential bias. In this sample, we re-estimate Equation 11 using hand-coded permit status. We also instrument for the human classification using the keyword-based and LLM-based classifications as instruments.

Repeat-Listing Difference-in-Differences. Our second specification leverages lot-level changes over time in approval for identification:

$$\log p_{it} = \alpha_i + \alpha_t + \beta \mathrm{RTI}_{it} + u_{it}, \tag{12}$$

where α_i is a lot fixed effect and α_t is a year fixed effect.[18] We refer to β from this regression as the repeat-listing approval premium. This design faces two key threats to identification, as well as a third concern about external validity.

The first threat is time-varying unobserved lot attributes. We state the identification assumption formally in Section 3.4. To yield a biased estimate, the average price on properties that obtain preapproval must have a different counterfactual trajectory in the absence of preapproval than prices on properties that remain unpermitted. One plausible reason for such a bias is neighborhood-level change, such as gentrification, that may correlate with preapproval. We address this concern by testing sensitivity to a rich set of time-varying controls. Furthermore, our event study transparently

[18] We use a fixed-effect, rather than first-difference, specification to accommodate properties with multiple listings and to account more transparently for variation in the timing of listings.

addresses this point, insofar as sharp price changes before approval are inconsistent with smoothly-trending confounders like gentrification.

A second threat is the capitalization of anticipated approval. When prices of unapproved properties are influenced by expectations of future approval, this anticipation bias implies the approval premium should not be interpreted as a willingness to pay for permits (Bishop and Murphy, 2011). Our event-study design has an additional role here, in that we can restrict the reference group to observations where the capitalization of anticipated approval is unlikely.

Third, the repeat-listing design inherently relies on properties that we observe switching permit status. As noted previously, this subsample of properties differs observably from larger populations of interest (see Appendix Table A9). We address this point in two ways. First, we follow Miller et al. (2023) in reweighting our approval-premium regressions to address observable imbalance. Second, to address selection on unobservables into preapproval, we estimate heterogeneity in premia by whether the seller is a specialist investor.[19]

Event Study. We estimate the following specification to expose capitalization dynamics:

$$\log p_{it} = \alpha_i + \alpha_t + \sum_{s \in \mathcal{S}} \beta_s \mathrm{RTI}_{i,t+s} + u_{it}, \tag{13}$$

where α_i and α_t are parcel and year-month fixed effects. The set $\mathcal{S} = \{0, 1, 2, 4, \ldots, 9\}$ contains indicators for contemporaneous permit approval ($s = 0$) and leads up to nine years prior to approval. We omit the indicator for the third lead ($s = 3$), which serves as the reference year. We also limit the sample to properties that obtain preapproval while in sample and remove properties immediately following approval. This ensures that the coefficients β_s are identified exclusively from comparisons among not-yet-preapproved properties and properties just preapproved in the current year.

To determine the appropriate reference year, we first estimate Equation 13 using an indicator for whether permits are "in progress" as the outcome. This supporting analysis provides direct evidence on the arrival of information about permit status, and thus the plausible start of any land-price capitalization. We use those findings to set the event-study reference year, so that the estimates of the approval premium plausibly capture the full value of permits.

4.2 Results for Vacant Land

We present our results separately for vacant and nonvacant properties. Vacant land arguably provides a cleaner setting, as it is inherently not vulnerable to unobserved characteristics of existing

[19] This test is motivated by viewing the selection issue as in Melitz (2003): Specialist investors may pay high fixed costs to build capacity for complex approvals with high price returns. In this view, specialists should secure a higher premium than the typical preapproval, one that might sharply overstate citywide average permitting costs. We ultimately find little evidence for this view (see Appendix Figure A9 and Table A15).

Figure 5: Estimates of the Approval Premium

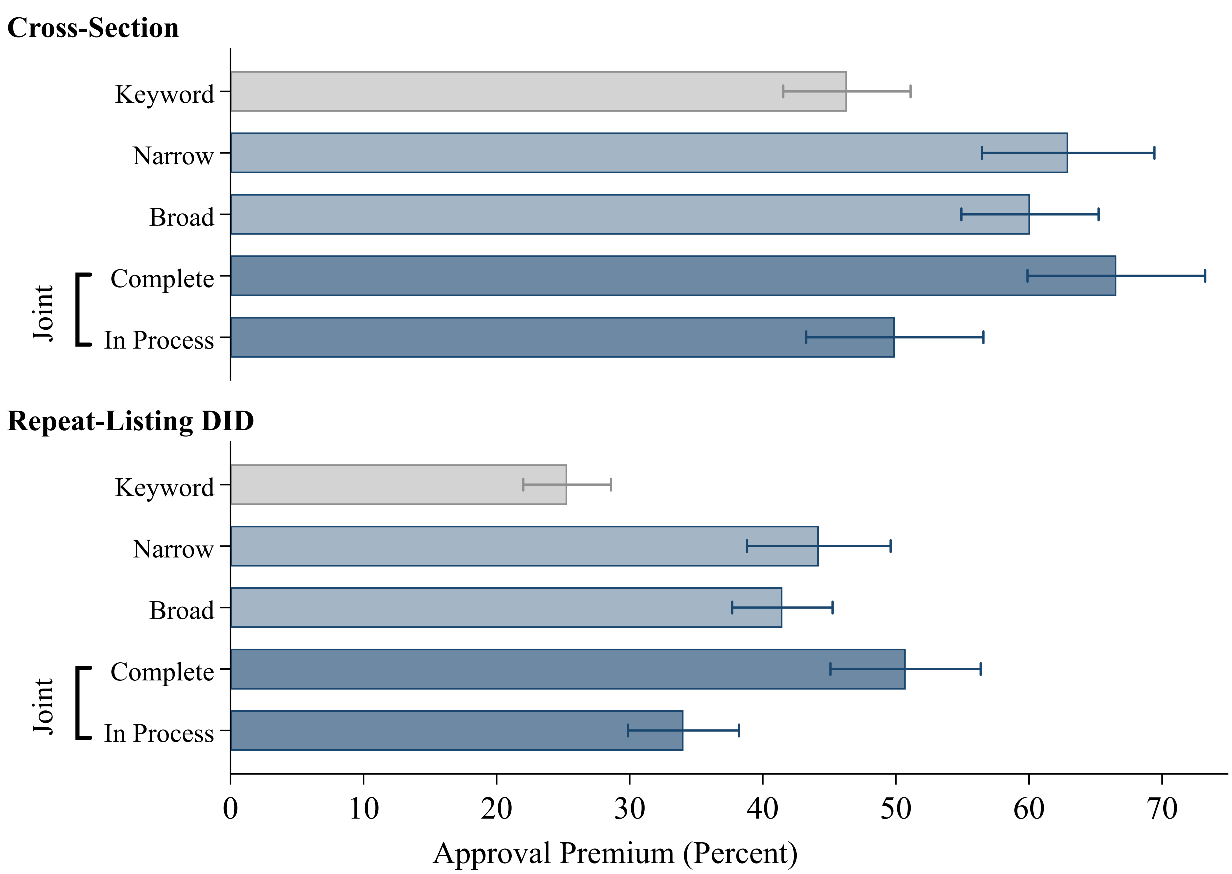

Notes: This figure displays estimated approval premia from Equations 11 (cross-section) and 12 (repeat-listing difference-in-differences). All coefficients are exponentiated, showing $100[\exp(\beta) - 1]$. See the main text for definitions of the measure of permit status. Bars reflect 95-percent confidence intervals, calculated using the delta method and with standard errors clustered by property.

structures. Nonvacant properties represent a larger market, especially in a densely-built city (Gedal and Ellen, 2018). That setting also allows us to explore heterogeneity in approval content.

Main. Figure 5 presents our estimates of the approval premium across specifications, grouped as the cross-sectional regression (Equation 11) and repeat-listing DID (Equation 12). Within these groups, we vary the measure of preapproval. In the cross-sectional results, the keyword-based measure of preaproval yields an approval premium of 50 percent, somewhat lower than the LLM-based measures. Taken alone, both the "narrow" and "broad" LLM-based measures yield similar estimates, but when included in the same regression, properties with fully-completed permits command a larger premium to properties with permitting in progress.

Moving from the cross section to the repeat-listing design, both keyword-based and LLM-based estimates shrink, although more so the former than the latter. This is consistent with reduced measurement error in the LLM-based measure, and the greater sensitivity to measurement error in fixed-effect specifications. We again find evidence that complete permits are valued above those in progress. The second-to-last row shows our preferred estimate of the approval premium, 50 percent (with a standard error of 5 p.p.).

Among properties with preapproved permits, the median implicit price of an approval is approximately \$48 per square foot. Similarly, the median dollar premium for an approval is approximately \$770,000 (see Appendix Table A10). Using construction-cost estimates from R.S. Means Company as a benchmark, permitting is approximately 36 percent of these costs (which include materials, labor, equipment, and site work; see Figure 9).

Event Study. Figure 6 presents estimated "event-study" coefficients from Equation 13. The coefficients in Panel A capture the dynamic path of land-price capitalization of anticipated permit approval. We find that price capitalization of approval begins approximately two years before, with no evident pretrend through nine years before approval.

We can also see that the time path of capitalization is consistent with the observed arrival of permit information. Panel B replaces the price outcome with an indicator for whether the listing indicates that permitting is in progress but not yet completed. Consistent with the price path, at two years before permitting is complete, listings begin stating that permitting is in progress. Of course, they stop indicating permitting is in progress when permits are ready-to-issue, yielding the negative final coefficient in the graph.

Measurement Error. We consider the reliability of our measures of preapproval and the sensitivity of our estimates to measurement error. We lack "ground truth" measures of preapproval, and so our focus is on inter-rater agreement. We estimate bivariate regressions of one rater's (a human or the LLM) binary classification on another rater's classification of the same listing. Human-versus-human comparisons yield a pooled reliability coefficient of 0.75, similar to the human-versus-LLM coefficient of 0.80 (see Appendix Table A11). These findings suggest our LLM ratings are of similar quality to our human ratings. Splitting the sample by raters' subjective confidence, we also find disagreement is higher among listings where permit status is perceived as unclear.

Within our human validation sample, we can attempt to adjust for measurement error by instrumenting the human ratings using the keyword- and LLM-based measures (see Appendix Table A12). In both cross-section and repeat-listing specifications, IV estimates of approval premia are approximately 40 percent larger than corresponding OLS estimates and are precisely estimated. We find similar inflation in the nonvacant sample (Appendix Table A13).

We present these as supporting analyses rather than main results for two reasons. First, human review is prohibitively costly to perform on all listings. Second, the validity of the IV strategy relies on an exclusion restriction that is unlikely to be precisely satisfied in this setting. If LLM

Figure 6: News Arrival and Dynamic Capitalization of Permit Approval

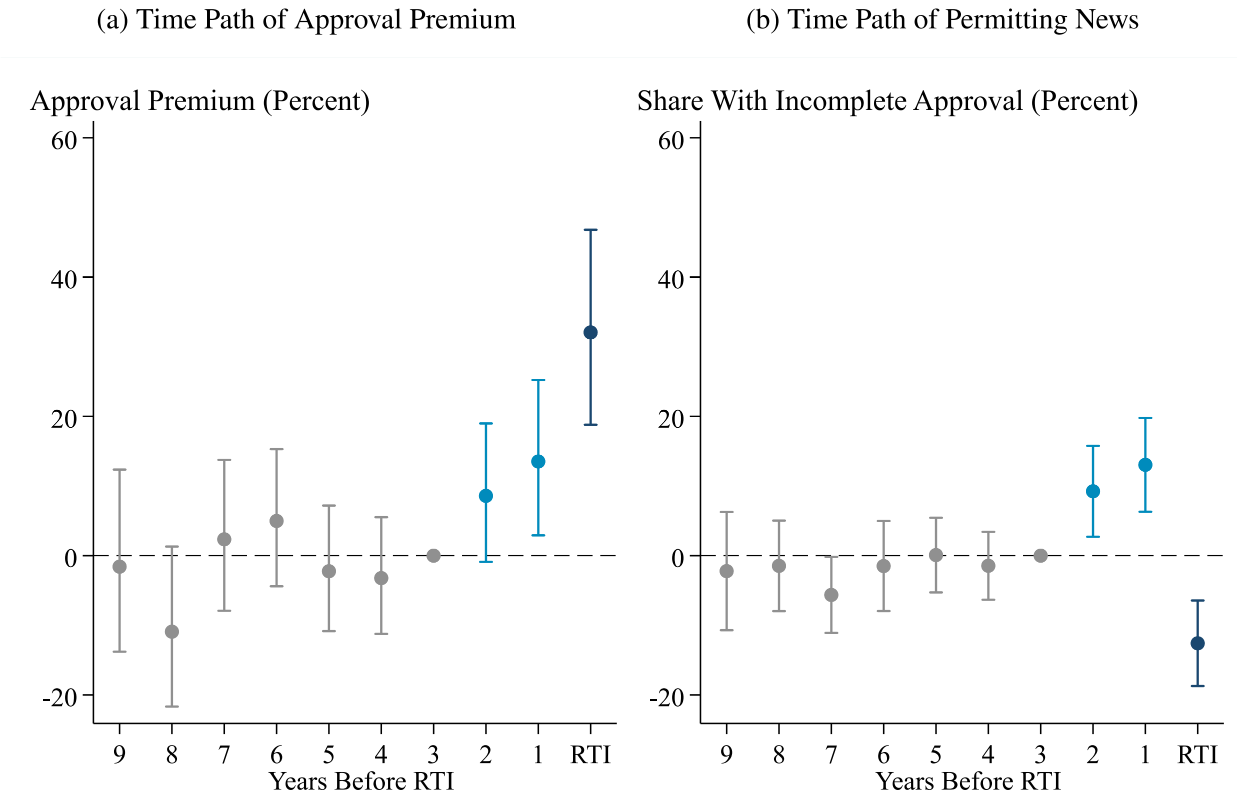

Notes: This figure displays event-study coefficients from Equation 13. In Panel A, the outcome is the log price. In Panel B, the outcome is an indicator for whether the listing indicates permitting is in progress. Coefficients in Panel A are exponentiated, with the figure showing delta-method estimates of $100[\exp(\beta) - 1]$. Both panels include nine leads of an indicator for the property being fully approved. The reference period is three years prior to the RTI listing. All regressions are estimated among vacant properties (see Appendix Figure A6 for the nonvacant-property event study). Bars indicate 95-percent pointwise confidence intervals, with standard errors clustered by parcel.

classification errors are positively correlated with unobserved quality (e.g., "hype" in the listing remarks), this bias is amplified in IV relative to OLS. Consequently, we view the IV estimates as a plausible upper bound and the OLS estimates as a lower bound.

Robustness Checks. We have performed extensive robustness checks with respect to controls for time-varying characteristics as well as to alternative sample definitions (see Appendix Figure A7). First, the inclusion of zipcode–year fixed effects or year-interacted tract characteristics (e.g., poverty rate) reduces the estimated approval premium by around 6 percentage points. Second, we use machine-learning methods (RoBERTa, see Liu et al., 2019) to generate a vector embedding of the listing text. Controlling for the first ten principal components of this embedding also modestly reduces the approval premium. Third, our results are robust to dropping subsamples with potential problems, such as sales within two years of approval (to rule out anticipation), distressed sales, top-priced land, land that trades more than twice in sample, or listings with ambiguous permit status. Fourth, we can also exclude unsold listings from our analysis (i.e., limiting to completed sales).[20] Finally, estimates of approval premia are unaffected by controls for other listing attributes (lot size, listing month, remarks length in characters; see Appendix Table A14).

We also compute standard errors on approval premia using methods that account for spatial correlation (Conley, 1999; Müller and Watson, 2023). Preapproval is nearly spatially independent across properties within neighborhood. By consequence, estimated premia remain statistically significant even under conservative inference assumptions (see Appendix Figure A8).

Following Miller et al. (2023), we also reweight our data so that the identifying samples resemble larger populations of interest on observed characteristics. For the repeat-listing specification, reweighting the data to resemble the cross-sectional identifying sample, all preapproved properties, or all not-preapproved properties has little effect on approval premia (see Appendix Figure A9). "Switcher" properties are thus not observably selected on predictors of higher or lower premia.

Heterogeneity in Premia. To consider potential selection-on-gains, we test whether specialized investors achieve higher approval premia than non-specialists (see Section 2.6). We find no significant difference in premia for fully-approved properties, though there is limited evidence of a difference for projects still in progress (see Appendix Table A15). These results suggest that the returns to permitting are relatively uniform across seller types, providing some justification for applying our estimates to the broader Los Angeles land market.

We also inspect heterogeneity over time and by zoning class. We find the 2024-specific premium is slightly larger than our pooled estimate (see Appendix Figure A13). Estimates of heterogeneous premia by zoning class provide inconclusive results (see Appendix Figure A14).

4.3 Results for Properties with Existing Structures

In the model, the value of a property with an existing structure but without new permits is

$$\frac{1}{1-\beta(1-\delta)}\left[\,\pi_j(\bar{h}_j) + \beta\delta\cdot p_{0,j}(\bar{h}_j)\,\right], \tag{14}$$

reflecting per-period income π_j and demolition probability δ that returns the property to raw land. The dollar-value approval premium on nonvacant property is thus $\beta\delta\tau_j/(1-\beta(1-\delta))$.

Equation 14 informs our interpretation of approval premia on nonvacant property in several ways. First, all else equal, the premium should decrease in percentage terms with the value of the existing structure (higher π_j). Second, the premium increases in the probability that existing structures are soon demolished (higher δ). Third, if existing structures are immediately demolished

[20] We have also varied the minimum-price threshold used to screen out potentially erroneous and non-arm's-length transactions (see Appendix Figure A10). We find modestly larger premia without any threshold. Premia are robust, if smaller, even with thresholds we view as highly conservative (e.g., $200,000).

($\beta\delta \approx 1$), then premia should be similar in dollars between the vacant and nonvacant samples.[21]

We first report estimates of the average approval premium in this sample: Preapproval raises the price of property with existing structures by 10 percent (see Appendix Figure A11).[22] Event-study estimates in this sample are of a similar magnitude (see Appendix Figure A6). To explain the divergent premia between the two samples, we consider the three hypotheses above.

First, some existing structures on these properties carry value. If a preapproved project's value is fixed in level terms, one expects to find a smaller percentage premium on a larger base. Panels A and B of Figure 7 validate this intuition, showing first that approval premia are larger for "teardowns" than for structures with ongoing value in their current use, and second that premia decline monotonically in floor area ratio of the existing structure (floorspace over lot area).

Second, on properties with existing structures, "what" has been preapproved also differs systematically from vacant land. For vacant land, preapproval entails almost exclusively new housing, whereas many nonvacant properties have preapproved additions and ADUs (see Appendix Table A7). Panel C of Figure 7 shows that approval premia vary systematically with the content of approvals. This variation does not track floor area, but instead it aligns suggestively with where permitting approval is most uncertain. Approval premia are largest for multifamily housing, for instance, and the premium for ADUs exceeds that of primary dwellings or additions.[23]

[21] Relatedly, the model casts property value as additively separable between existing structures and approved not-yet-built structures to be built. This view seems sensible for redevelopment but not smaller improvements to existing structures, such as a pool or garage. Our DFL weights assign almost zero weight to such properties.

[22] Cross-sectional estimates of the approval premium are approximately zero for properties with existing structures, even with controls (see Appendix Table A16), consistent with confounding by unobserved structure attributes.

[23] A similar analysis for vacant properties, however, shows limited evidence of a larger premium for multifamily approvals relative to all other approvals, mostly single-family housing (see Appendix Table A19).

Figure 7: Heterogeneity in Approval Premia for Nonvacant Lots

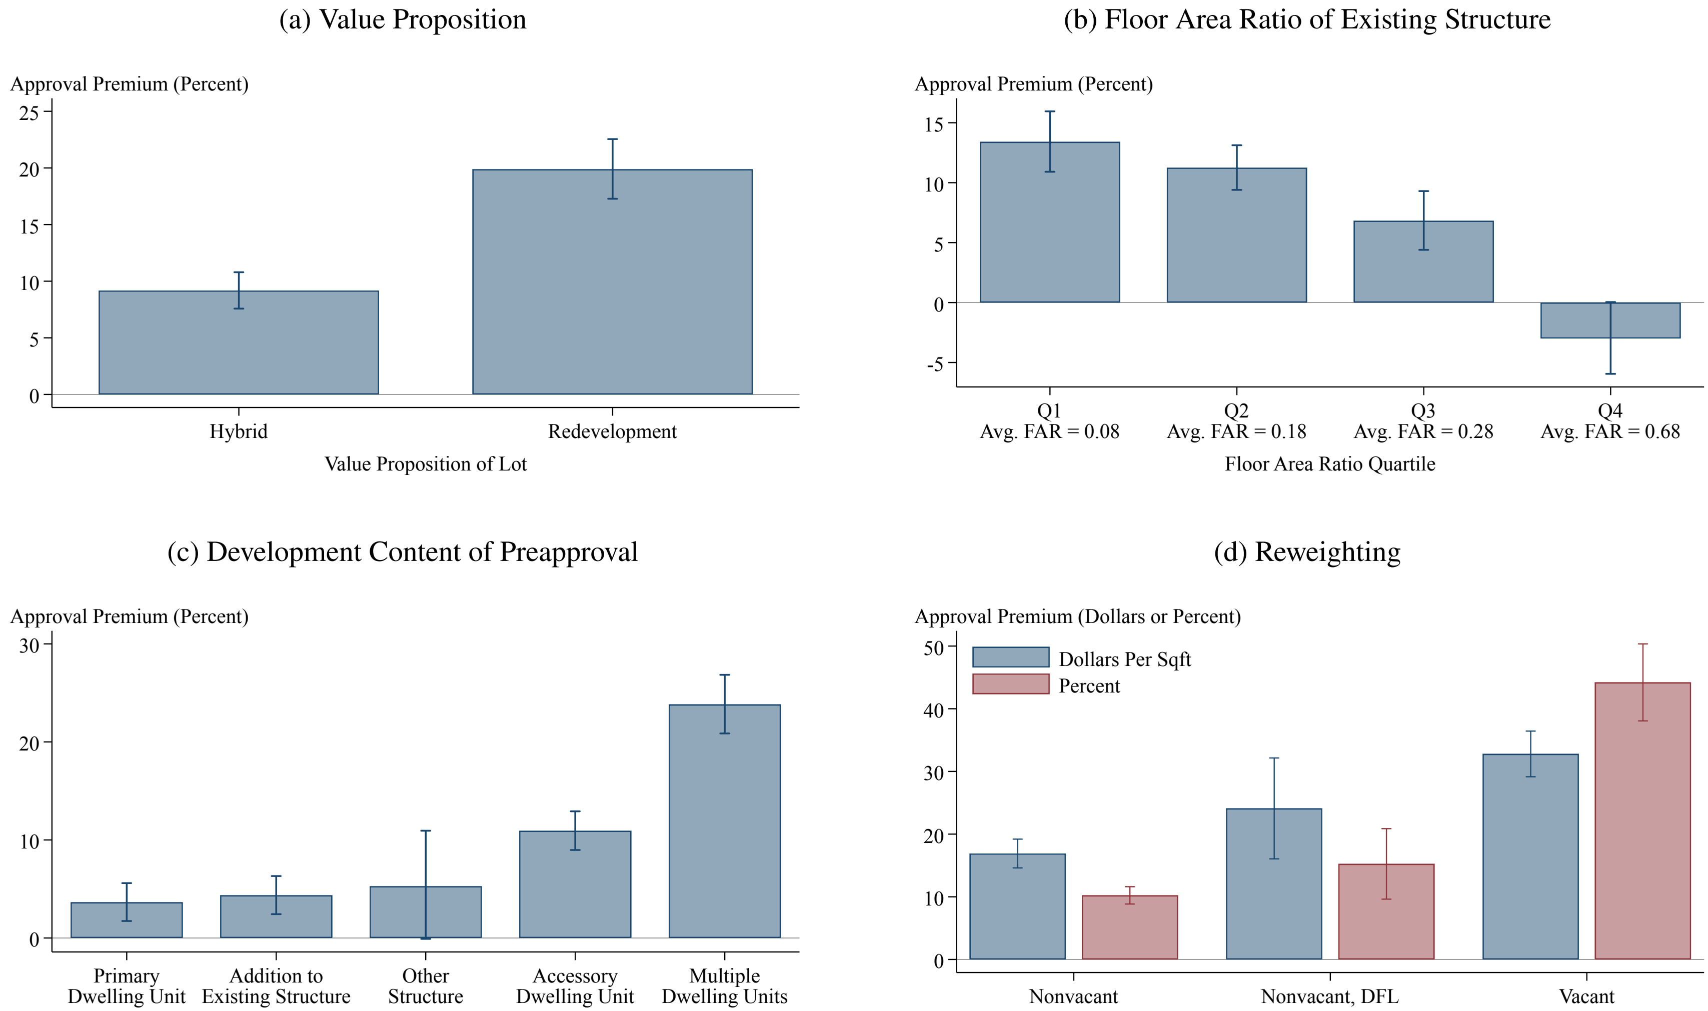

Notes: This figure displays several analyses of approval premia of properties with existing structures. All estimates reflect exponentiated coefficients on ready-to-issue status, estimated using our repeat-listing specification (Equation 12), potentially interacted with other property characteristics. Bars reflect 95-percent confidence intervals from standard errors clustered at the property level. Panel A also includes controls for the main effect of lot value proposition. "FAR" is the floor area ratio, and "DFL" is the DiNardo et al. (1996) reweighting method. See the main text for further explanation. Standard errors in Panels A–C are calculated by delta method, and by the bootstrap in Panel D.

Panel D then tests whether the approval premia are equal across the two samples. We implement this test by reweighting, as in DiNardo et al. (1996, DFL), the nonvacant sample to resemble the vacant one on several characteristics: zipcode, sale year, lot size, and the content of approvals. In the first and third groups of bars, we report the dollar and percentage approval premia in both samples. The middle set of bars reports these premia in the DFL-reweighted nonvacant sample. We cannot reject the dollar-value approval premia are equal when comparing the vacant and the reweighted nonvacant samples.

5 The Time-to-Build Effect of Preapproval

Why do developers pay a premium for preapproved land? This section examines a natural potential source of its value to developers: reducing time-to build, conditional on completion, as well as completion risk. Next, we consider alternative explanations of the approval premium. Throughout this section, we rely on a sample of listings matched to permits, in which we can compare the land uses and construction timelines of properties by permit status.

5.1 Approach

Background. We estimate discrete-time hazard models to flexibly account for right-censoring (Kalbfleisch and Prentice, 2002), which is important given long development timelines. Let s denote the time elapsed since sale, measured in quarterly intervals, and let t denote calendar time in quarters (e.g., 2024 Q4). We measure our outcome of interest, project completion, using the first issuance date on any permanent certificate of occupancy linked to a new-building permit on the property. Our estimand is the effect of preapproval on the cumulative completion hazard function of preapproved lots, relative to a counterfactual in which they were not preapproved.

Our empirical strategy is to compare preapproved properties to ones that are observably similar but lack approved permits when sold. While this approach faces the challenge of selection into preapproval, we see no alternative in this setting. The same lot is rarely developed twice within a short time interval, so there is no sensible "repeat" design for construction outcomes. To assess the threat of confounding unobservables, we show the sensitivity of our estimates to different sets of controls. We also compare our results to the observed time-to-permit of preapproved properties.

Specification. We model the probability of a project i completing in within a quarterly interval s using horizon-specific linear probability models. For a given horizon s, we let

$$q_{is} = \beta_s \mathrm{RTI}_i + \mathbf{X}_{it}\gamma_s + u_{ist}, \tag{15}$$

where q_{is} is an indicator for project i completing in horizon interval s. In our baseline specification,

$\mathbf{X}_{it}$ includes fixed effects by zipcode and year. In robustness checks, we consider three sets of controls: lot and neighborhood attributes (lot size and tract characteristics), project attributes (permitted floorspace, unit count, building height, permit-reported construction cost), and listing attributes (embedding vectors from the listing remarks).

The cumulative empirical hazard function for preapproved lots, $F_1(t)$, is constructed by chaining the quarterly empirical mean hazards $\bar{q}_{1,s}$ of this population: $F_1(t) = 1 - \prod_{s=0}^{t-1}\left(1 - \bar{q}_{1,s}\right)$. For the counterfactual cumulative hazard, $\hat{F}_0(t)$, we adjust each quarterly hazard using the respective horizon-specific estimate $\hat{\beta}_s$, as in $\hat{q}_{0,s} = \bar{q}_{1,s} - \hat{\beta}_s$. Horizon-specific estimates of the average effect of preapproval on preapproved properties are thus:

$$\hat{\Delta}(t) = F_1(t) - \hat{F}_0(t). \tag{16}$$

We use a nonparametric bootstrap to form standard errors on $\hat{\Delta}(t)$. For properties that are sold repeatedly, we build our sample to include only their final observed sale.

The direction of the potential bias from selection into preapproval is ambiguous (see our formal discussion in Section 3.4). On one hand, suppose "fast-to-build" developers usually outbid "slow-to-build" developers for preapproved land. If such developers complete projects faster for reasons that go beyond preapproval, we would overstate preapproval's effect. On the other hand, landowners might use preapproval to improve flawed properties that will have high time-to-build for reasons we cannot observe. These issues are difficult to adjudicate empirically in our setting, due to an absence of credible property-level selection instruments.

We estimate Equation 16 in two matched samples. First, we use listings that match to any submitted permit, even if not issued or completed. This sample, which provides our baseline results, is thus limited to cases where land buyers had apparent intention to build—mitigating selection concerns, but ruling out impacts of preapproval at the initial stage of the development process. Second, as a robustness check, we limit the sample to lots where a permit was issued for new construction, demolition, or grading, thus proceeding further along the stages of development. The contrast of these results sheds some light on selection into preapproval.

5.2 Results

Preapproval appears to reduce the time-to-build remaining after site acquisition, relative to otherwise-similar properties without approved permits at sale. Panel A of Figure 8 reports our estimates of dynamic treatment effects (Equation 16) through ten years after site acquisition.

At four years after site acquisition, preapproved properties are between 8 and 12 percentage points (depending upon controls) more likely to have completed construction. Panel B shows the counterfactual completion rate at this horizon is approximately 35 percent, suggesting that preapproval increases the probability of completion within five years by roughly 30 percent. This counterfactual uses the estimates with all controls.

The vertical gap between the curves in Panel B represents the change in the probability of completion at a fixed point in time, whereas horizontal distances give estimates of time savings. For example, projects on preapproved land reach the milestone of a 40-percent completion probability one to two years earlier than if instead they started from raw land (see Appendix Figure A19).

The estimated effects of preapproval emerge from one to three years after site acquisition. This timing aligns with a plausible pace of construction, if preapproval allows development to begin immediately after site acquisition, while other buyers must wait for approvals.

Furthermore, our effects suggests preapproval reduces both time delays and completion risk. Even with our full set of controls, we find preapproval raises completion probabilities at ten years after site acquisition, a horizon that seems beyond most construction timelines. Our short-horizon estimates appear quite robust to controls, although there is some sensitivity at longer time horizons. Panel C reports the sensitivity to controls in greater detail. In particular, this long-horizon sensitivity emerges when we include controls for project attributes; other controls matter little.

Are these estimates plausible in magnitude? One natural benchmark is the time-to-permit for properties that sell as preapproved land, measured from either permit submission or site acquisition (see Appendix Figure A20). We find median times-to-permit of 2.2 and 3.0 years.[24] These times are roughly consistent with horizontal distances between the actual and counterfactual completion curves (see Appendix Figure A19). By implication, the implicit market for development approvals lets developers trade time with landowners on an approximately one-for-one basis—that is, shifting holding times across owners without changing the total time-to-build.

Robustness Checks. We summarize two robustness checks related to sample definition. First, estimated time-to-build effects are strikingly similar, if modestly larger, in the alternative sample that conditions inclusion on a matched permit issuance (Appendix Figure A21). The counterfactual completion rate at four years after site acquisition is similar. In this sense, our main results may somewhat understate the importance of permitting for time-to-build.

Second, our permit analyses combine data from the City of Los Angeles with data from 87 other County municipalities that we suspect to be of lower quality. Split the data into City and rest-of-County subsamples, we find our time-to-build effects are driven by the City (see Appendix Figure A22). Estimates on the rest-of-County subsample are similar but imprecise.

[24] This benchmark provide a lower bound on potential time savings if, due to specialization, the landowners are faster at securing approvals than the developers who purchase the site.

Figure 8: Dynamic Effects of Preapproval on the Timing of Completion

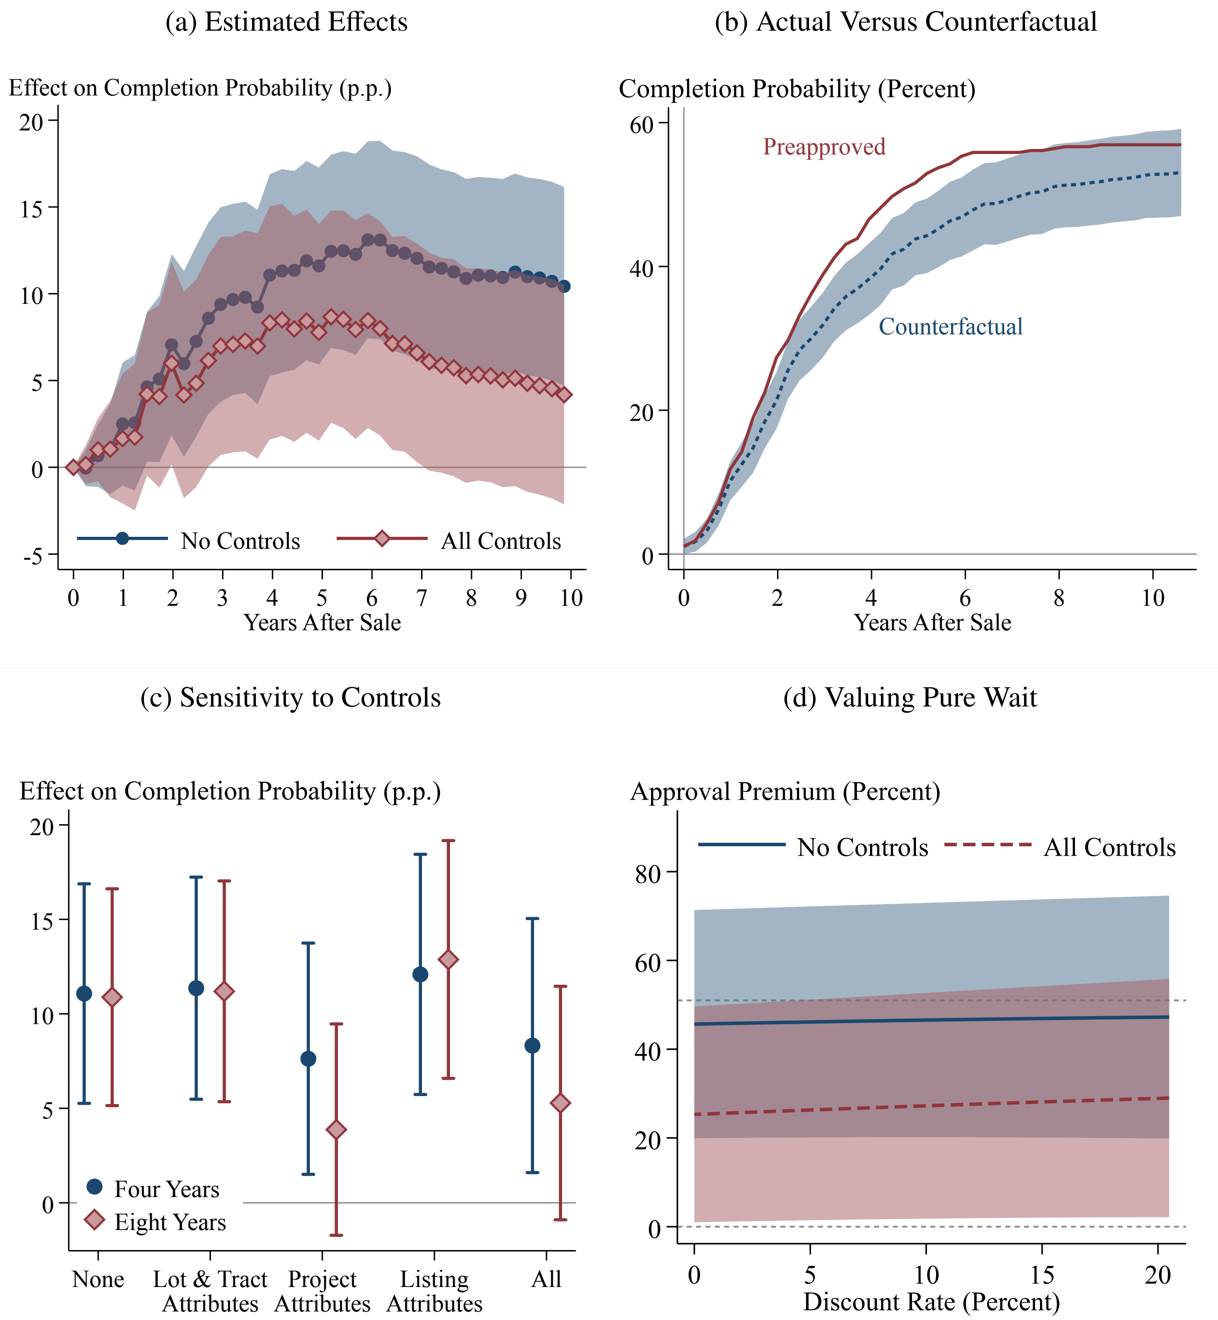

Notes: This figure displays estimated effects of preapproval on time to completion (Panels A and B) as well as sensitivity and heterogeneity analyses of these effects (Panels C and D). Panel A plots quarterly estimates $\hat{\Delta}(t)$ from Equation 16 from zero to ten years after site acquisition, showing results with and without controls. Panel B plots the empirical cumulative completion hazard for preapproved properties (red solid line) and then subtracts off the estimated effects of preapproval to obtain the counterfactual completion profile in time (blue dashed line) The sample is all vacant properties that ever submit a permit. Panel C reports estimated effects on the cumulative completion probability at two years (blue circles) and ten years (red diamonds) after site acquisition for different sets of controls. Panel D reports the land-price premium consistent with reduced time-to-build at various discount rates. Color bands indicate 95-percent pointwise confidence intervals, with standard errors calculated by a nonparametric bootstrap clustered by property.

5.3 Alternative Explanations of the Approval Premium

Preapproval could be valuable for other reasons beyond time savings and completion-risk reduction. Two plausible explanations are examined here, by returning to land-price regressions.

Permitted Use. Preapproval can essentially commit the municipality to allow a property to have a more-valuable, more-intensive land use than is typical. Consistent with this hypothesis, we find that on average, preapproved properties submit permits with greater floor space, unit count, construction cost, and building height than similarly-sized properties without preapproval in the same neighborhood and year (see Appendix Table A18).

Might these differences drive the approval premium? We address this concern by augmenting the repeat-listing specification with permit controls—floorspace, unit count, valuation, building height—imputing each when it is not observed (see Appendix C for details). Our estimates suggest the approval premium is not readily explained by more-intensive use (see Appendix Table A17).

Resolving Uncertainty. Permitting may also resolve uncertainty about risks to property value, such as soil contamination, wildfire vulnerability, or neighborhood opposition to development. Importantly, these risks might raise development costs or reduce completed-structure value beyond their effects on time-to-build, completion risk, or observed dimensions of land use and thus present a distinct mechanism for the approval premium.

We examine this hypothesis by estimating heterogeneity in approval premia along several proxies of salient neighborhood-level risks: a lead-pollution index from the California Department of Public Health, the per-capita rate of CEQA environmental impact reports, and a wildfire risk score from the Federal Emergency Management Agency (FEMA).

Approval premia appear higher in areas with greater lead risk. This heterogeneity in premia is not readily explained by neighborhood covariates, such as poverty or population density (see Appendix Figure A17). Using areas with no measured lead risk to provide an upper bound on the uncertainty channel suggests it could explain up to 40 percent of the approval premium.[25] There is no evidence of notable heterogeneity in premia with respect to CEQA filings or wildfire risk.

6 The Approval Premium as a Regulatory Tax

This section interprets the approval premium as a measure of regulatory taxation. First, we decompose our estimate of the premium into pure-wait and capitalized-hassle components. Second, we estimate the permitting share of the total regulatory tax on new housing, considering both Los Angeles County on average as well as variation across its neighborhoods.

[25] Such a conclusion assumes that time-to-build effects (and all other mechanisms for the approval premium) do not covary with lead risk. Our time-to-build analysis is inadequately powered to usefully test that hypothesis.

6.1 Pure Wait and Capitalized Hassle

Proposition 1 expresses the approval premium as the sum of our two sources of permitting cost, pure wait and capitalized hassle. We now use our estimates from Sections 4 and 5 to consider to what extent preapproval is valuable because permitting involves burdensome procedures, versus because the process imposes time delays and completion risk.

At a discount rate of 10 percent per annum, we find developers would be willing to pay premia of 27 percent or 45 percent, depending upon controls (see Panel D of Figure 8). At discount rates of 5 and 15 percent, the expected premia differ only slightly. This insensitivity arises because preapproval mostly retimes completion at horizons that are short and thus insensitive to discounting. Using our estimated 50-percent premium (in Figure 5), around half of the total cost of permitting reflects pure wait and the remainder can be attributed to capitalized hassle. However, we note these shares are estimated imprecisely. At plausible discount rates (see Gormsen and Huber, 2025), or without controls, we cannot reject that the entire premium reflects pure wait.

From a policy perspective, the distinction between pure wait and capitalized hassle bears importantly on permitting reform proposals. If pure wait, reforms intended to reduce time-to-build or completion risk—such as approval "shot clocks," limiting potential contestation and litigation of proposed development, or expanding review capacity in permitting offices—seem aimed at the underlying source of regulatory cost (Liscow, 2025). On the other hand, if instead the approval premium were to reflect capitalized hassle, then policymakers might instead consider reforms intended to reduce compliance activities, such as the complexity of regulatory filings or the share of projects that require multi-agency reviews (see O'Neill et al., 2019a,b).

We also ask whether the implied estimates of capitalized hassle are themselves reasonable. To fit empirically, we expect predevelopment costs of 4 to 23 percent of the price of raw land. We consider fiscal charges and non-fiscal predevelopment costs separately in this calculation, as we have data specific to Los Angeles on the former but not the latter. For each proposed project, we compute permit fees according to current City regulations (see Appendix C for details). For non-fiscal predevelopment costs, we use estimates from R.S. Means.

We estimate that average permit fees for preapproved projects are 1.2 percent of raw-land value (i.e., prices adjusted for the approval premium). Larger fees on development are paid when permits are finally issued, but these would not be capitalized into the approval premium. Imputed predevelopment expenses (architecture and engineering) are on average 12.6 percent of raw-land value. These figures appear broadly consistent with our estimates of capitalized hassle.

Figure 9: Permitting Versus the Housing Cost Wedge

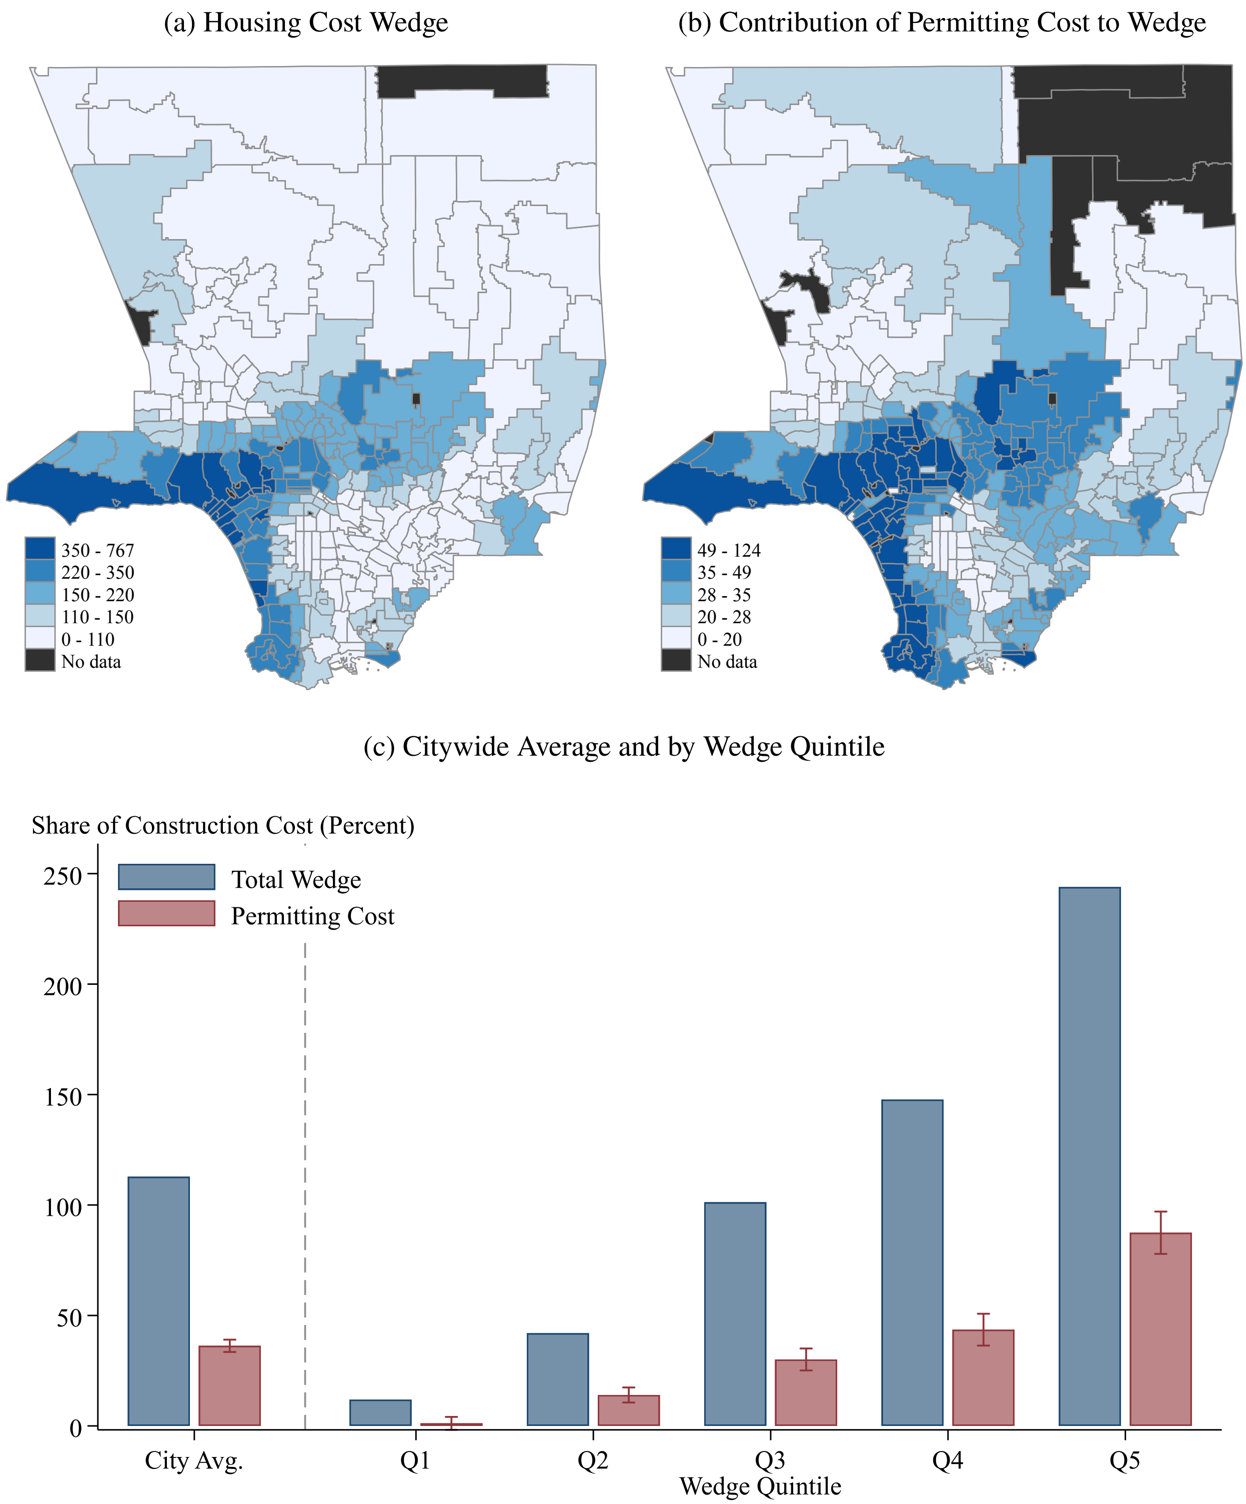

Notes: This figure displays zipcode-level estimates of the median wedge between rents and construction costs (Panel A), the model-implied contribution of permitting costs to this wedge (Panel B), and a summary of the wedges and permitting's contribution on citywide average and by zipcode wedge quintile (Panel C). See Appendix C for details on wedge measurement and the spatially-weighted regression used to obtain zipcode-level estimates of permitting costs.

6.2 Permitting and the Housing Cost Wedge

Methods. We adapt methods in Glaeser and Gyourko (2003, 2018) and Gyourko and Saiz (2006) to impute construction costs at the property level. Differences between actual sale prices and these imputed costs yield these papers' preferred measure of the housing cost wedge.

We follow preceding work in using R.S. Means data to calibrate a construction-cost function. However, rather than American Housing Survey microdata as in these papers, we draw upon listings to fill in the arguments of the cost function: property locations, sale prices, and structure characteristics. The benefit for our research question is that we can obtain neighborhood-level wedges and make finer adjustments for characteristics.[26]

Imputations in hand, we compute neighborhood-level wedges as medians of property-level differences between sales prices and costs, expressed as a share of costs. To obtain the contribution of permitting to the wedge, we leverage Proposition 3, which shows the permitting share of the wedge is $\theta/(1+\theta)$, where θ is the approval premium. We then multiply the shares by the wedges.

Results. We find that, for the median home sale in Los Angeles in 2024, prices were approximately 120 percent above construction costs. This estimate aligns with that of Glaeser and Gyourko (2018) for an earlier period. From 1995 to 2024, the median wedge declined on net as a fraction of construction cost, albeit with striking cyclicality due to the 2000s housing boom and Great Recession (see Appendix Figure A16).

Panel A of Figure 9 documents geographic variation in the housing cost wedge. The wedges are largest in high-price areas such as Santa Monica and Westwood, implying construction costs vary less than home prices across the county. In high-price areas, home prices are more than four times construction cost. Wedges are smallest in areas where supply is unlikely to be constrained, such as the Mojave Desert or economically distressed parts of South Los Angeles.

Panel B shows geographic variation in the contribution of permitting to the housing cost wedge, reported as a percentage of construction cost. We estimate these effects using a geographically-weighted version of Equation 12 (see Appendix C).[27] On citywide average, permitting adds 36 percent to construction cost, or approximately 32 percent of the total housing cost wedge in Los Angeles. Neighborhood variation in permitting cost is visibly positively associated with wedges, time-to-build (Panel B of Figure 2), and the preapproval rate (Panel B of Figure 3).

[26] Appendix C explains our approach. In summary, we account for project scale (square footage, number of stories), construction materials (wall material, roof type), specific amenities (e.g., baths, basements, pools, fireplaces, fences, porches, garages, and heating systems), and general "build quality" (inferred from listing keywords). We use a repeat-sale design to adjust for depreciation and remove the hedonic value of land semiparametrically. Due to the presence of existing structures, whose value enters the property price but not the permitting cost, performing an equivalent wedge-share calculation in the nonvacant sample is not straightforward.

[27] Appendix Figure A18 presents two other maps of the approval premium, as a percent of counterfactual raw-land value and in dollars per square foot of land. Appendix Figure A12 shows heterogeneity in premia according to specific neighborhood characteristics.

Panel C summarizes these geographic patterns. Splitting neighborhoods according to their wedges, we find permitting contributes essentially nothing to construction cost in the bottom quintile and adds around 85 percent to construction cost in the top quintile (see Appendix Table A20 for estimates). Such variation is associated with, but clearly smaller than, the variation in total wedges. By consequence, permitting alone cannot explain why prices dramatically exceed costs in the city's most-desirable areas.

We close our analysis by examining how much opportunity might exist to reduce permitting costs. Assuming the economic cost of permitting scales linearly with time-to-build, raising Los Angeles to the speed of Fort Worth and Raleigh (see Panel B of Figure 1) would yield a gain equivalent to 21 percent of construction costs. Further assuming these savings are fully capitalized into land, permitting reform might raise citywide-average land prices in Los Angeles as much as 25 percent. By contrast, the city's latest permitting reforms had little apparent effect on construction (Kestelman, 2025), suggesting that objective might require far-reaching changes.

7 Conclusion

Housing cost pressures in leading U.S. cities have set economists, governments, and the public on a search for policy answers. While permitting is often cited as a burden on development, considerable measurement challenges have mostly limited these debates to the realm of anecdote.

This paper provides the first systematic estimates of permitting costs and their share of the aggregate housing cost wedge in a major U.S. city. We do so by exploiting a mature, but previously unstudied, submarket in Los Angeles for land with preapproved permits. Using a repeat-listing difference-in-differences design, we find that land with approved permits holds an average premium of 50 percent relative to raw land. We further leverage the permit data to show that preapproval appears to reduce remaining time-to-build for acquiring developers. Overall, our results suggest permitting is a major barrier to housing supply in Los Angeles, putting in new focus the importance of de-facto burdens versus de-jure regulatory constraints.

References

Anagol, Santosh, Fernando Ferreira, and Jonah Rexer, "Estimating the Economic Value of Zoning Reform," *American Economic Journal: Economic Policy*, forthcoming.

Autor, David H, "Work of the Past, Work of the Future," in "AEA Papers and Proceedings," Vol. 109 American Economic Association 2019, pp. 1–32.

Bajari, Patrick, Jane Cooley Fruehwirth, Kyoo Il Kim, and Christopher Timmins, "A Rational Expectations Approach to Hedonic Price Regressions with Time-Varying Unobserved Product Attributes: The Price of Pollution," *American Economic Review*, 2012, *102* (5), 1898–1926.

Banzhaf, H Spencer, "Difference-in-Differences Hedonics," *Journal of Political Economy*, 2021, *129* (8), 2385–2414.

Bartik, Alexander, Arpit Gupta, and Daniel Milo, "The Costs of Housing Regulation: Evidence from Generative Regulatory Measurement," Working Paper 2024.

Baum-Snow, Nathaniel and Gilles Duranton, "Housing Supply and Housing Affordability," Working Paper 33694, National Bureau of Economic Research 2025.

Bishop, Kelly C and Alvin D Murphy, "Estimating the Willingness to Pay to Avoid Violent Crime: A Dynamic Approach," *American Economic Review*, 2011, *101* (3), 625–629.

Brooks, Leah and Zachary Liscow, "Infrastructure Costs," *American Economic Journal: Applied Economics*, 2023, *15* (2), 1–30.

Brueckner, Jan K and Kala Seetharam Sridhar, "Measuring Welfare Gains from Relaxation of Land-Use Restrictions: The Case of India's Building-Height Limits," *Regional Science and Urban Economics*, 2012, *42* (6), 1061–1067.

Büchler, Simon and Elena Lutz, "Making Housing Affordable? The Local Effects of Relaxing Land-Use Regulation," *Journal of Urban Economics*, 2024, *143*, 103689.

Buck, Steven, Maximilian Auffhammer, and David Sunding, "Land Markets and the Value of Water: Hedonic Analysis Using Repeat Sales of Farmland," *American Journal of Agricultural Economics*, 2014, *96* (4), 953–969.

Conley, Timothy G, "GMM Estimation with Cross Sectional Dependence," *Journal of Econometrics*, 1999, *92* (1), 1–45.

Davis, Lucas W, "The Effect of Health Risk on Housing Values: Evidence from a Cancer Cluster," *American Economic Review*, 2004, *94* (5), 1693–1704.

Diamond, Rebecca and Enrico Moretti, "Where Is Standard of Living the Highest? Local Prices and the Geography of Consumption," Working Paper 29533, National Bureau of Economic Research 2021.

_ , Zong Huang, and Tim McQuade, "The Unequal Effects of Upzoning: Evidence from Cook County," Working Paper 2025.

DiNardo, John, Nicole M Fortin, and Thomas Lemieux, "Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach," *Econometrica*, 1996, *64* (5), 1001–1044.

Duranton, Gilles and Diego Puga, "Urban Growth and Its Aggregate Implications," *Econometrica*, 2023, *91* (6), 2219–2259.

Fu, Yuming and C Tsuriel Somerville, "Site Density Restrictions: Measurement and Empirical Analysis," *Journal of Urban Economics*, 2001, *49* (2), 404–423.

Gabriel, Stuart and Edward Kung, "Development Approval Times and New Housing Supply: Evidence from Los Angeles," Working Paper 2025.

Ganduri, Rohan, Steven Chong Xiao, and Serena Wenjing Xiao, "Tracing the Source of Liquidity for Distressed Housing Markets," *Real Estate Economics*, 2023, *51* (2), 408–440.

Ganong, Peter and Daniel Shoag, "Why Has Regional Income Convergence in the US Declined?," *Journal of Urban Economics*, 2017, *102*, 76–90.

Gedal, Michael and Ingrid Gould Ellen, "Valuing Urban Land: Comparing the Use of Teardown and Vacant Land Sales," *Regional Science and Urban Economics*, 2018, *70*, 190–203.

Glaeser, Edward and Joseph Gyourko, "The Impact of Building Restrictions on Housing Affordability," *Economic Policy Review*, 2003, (Jun), 21–39.

_ and _ , "The Economic Implications of Housing Supply," *Journal of Economic Perspectives*, 2018, *32* (1), 3–30.

Glaeser, Edward L, Joseph Gyourko, and Raven E Saks, "Why Have Housing Prices Gone Up?," *American Economic Review*, 2005, *95* (2), 329–333.

_ , _ , and Raven Saks, "Why Is Manhattan So Expensive? Regulation and the Rise in Housing Prices," *The Journal of Law and Economics*, 2005, *48* (2), 331–369.

Gormsen, Niels Joachim and Kilian Huber, "Corporate Discount Rates," *American Economic Review*, 2025, *115* (6), 2001–2049.

Greenstone, Michael, "The Continuing Impact of Sherwin Rosen's 'Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition'," *Journal of Political Economy*, 2017, *125* (6), 1891–1902.

Gyourko, Joe and Jacob Krimmel, "The Impact of Local Residential Land Use Restrictions on Land Values Across and Within Single Family Housing Markets," *Journal of Urban Economics*, 2021, *126*, 103374.

Gyourko, Joseph and Albert Saiz, "Construction Costs and the Supply of Housing Structure," *Journal of Regional Science*, 2006, *46* (4), 661–680.

_ and Raven Molloy, "Regulation and Housing Supply," in "Handbook of Regional and Urban Economics," Vol. 5, Elsevier, 2015, pp. 1289–1337.

— , **Jonathan S Hartley, and Jacob Krimmel**, "The Local Residential Land Use Regulatory Environment across US Housing Markets: Evidence from a New Wharton Index," *Journal of Urban Economics*, 2021, *124*, 103337.

Hsieh, Chang-Tai and Enrico Moretti, "Housing Constraints and Spatial Misallocation," *American Economic Journal: Macroeconomics*, 2019, *11* (2), 1–39.

Jackson, Kristoffer, "Do Land Use Regulations Stifle Residential Development? Evidence from California Cities," *Journal of Urban Economics*, 2016, *91*, 45–56.

Kahn, Matthew E, Ryan Vaughn, and Jonathan Zasloff, "The Housing Market Effects of Discrete Land Use Regulations: Evidence from the California Coastal Boundary Zone," *Journal of Housing Economics*, 2010, *19* (4), 269–279.

Kalbfleisch, John D and Ross L Prentice, *The Statistical Analysis of Failure Time Data*, John Wiley & Sons, 2002.

Kestelman, Stephanie, "Effects of Transit-Oriented Incentive Programs: Evidence from Los Angeles," Working Paper 2025.

Klein, Ezra and Derek Thompson, *Abundance*, Simon and Schuster, 2025.

Kohlhase, Janet E, "The Impact of Toxic Waste Sites on Housing Values," *Journal of Urban Economics*, 1991, *30* (1), 1–26.

Kulka, Amrita, Aradhya Sood, and Nicholas Chiumenti, "Under the (Neighbor) Hood: Understanding Interactions Among Zoning Regulations," Working Paper 2024.

Liscow, Zachary, "Getting Infrastructure Built: the Law and Economics of Permitting," *Journal of Economic Perspectives*, 2025, *39* (1), 151–180.

Liu, Yinhan, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov, "RoBERTa: A Robustly Optimized BERT Pretraining Approach," *arXiv preprint arXiv:1907.11692*, 2019.

Lovo, Stefano and Christophe Spaenjers, "A Model of Trading in the Art Market," *American Economic Review*, 2018, *108* (3), 744–774.

Ludwig, Jens, Sendhil Mullainathan, and Ashesh Rambachan, "Large Language Models: An Applied Econometric Framework," Working Paper 33344, National Bureau of Economic Research 2025.

Lyons, Ronan C and Éamonn Sweeney, "Time to Build: Rules-Based Planning and Construction Project Duration in Dublin," *Journal of Housing Economics*, 2025, p. 102089.

Manville, Michael, Paavo Monkkonen, Nolan Gray, and Shane Phillips, "Does Discretion Delay Development? The Impact of Approval Pathways on Multifamily Housing's Time to Permit," *Journal of the American Planning Association*, 2023, *89* (3), 336–347.

Mawhorter, Sarah and Carolina Reid, "Terner California Residential Land Use Survey," Report, University of California, Berkeley 2018.

Mayer, Christopher J and C Tsuriel Somerville, "Land Use Regulation and New Construction," *Regional Science and Urban Economics*, 2000, *30* (6), 639–662.

Melitz, Marc J, "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 2003, *71* (6), 1695–1725.

Miller, Douglas L, Na'ama Shenhav, and Michel Grosz, "Selection into Identification in Fixed Effects Models, with Application to Head Start," *Journal of Human Resources*, 2023, *58* (5), 1523–1566.

Moretti, Enrico and Harrison Wheeler, "The Traffic Noise Externality: Costs, Incidence and Policy Implications," Working Paper 34298, National Bureau of Economic Research 2025.

Müller, Ulrich K and Mark W Watson, "Spatial Correlation Robust Inference in Linear Regression and Panel Models," *Journal of Business & Economic Statistics*, 2023, *41* (4), 1050–1064.

Nowak, Adam D and Patrick S Smith, "Quality-Adjusted House Price Indexes," *American Economic Review: Insights*, 2020, *2* (3), 339–356.

O'Neill, Moira, Giulia Gualco-Nelson, and Eric Biber, "Comparing Perceptions and Practice: Why Better Land Use Data Is Critical to Ground Truth Legal Reform," *UC Berkeley Terner Center Working Paper Series*, 2019.

_ , _ , and _ , "Examining the Local Land Use Entitlement Process in California to Inform Policy and Process," Working Paper 2, Berkeley Law Center for Law, Energy, and the Environment 2019.

Palmquist, Raymond B, "Measuring Environmental Effects on Property Values without Hedonic Regressions," *Journal of Urban Economics*, 1982, *11* (3), 333–347.

Quigley, John M and Larry A Rosenthal, "The Effects of Land Use Regulation on the Price of Housing: What Do We Know? What Can We Learn?," *Cityscape*, 2005, pp. 69–137.

Rosen, Sherwin, "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition," *Journal of Political Economy*, 1974, *82* (1), 34–55.

Turner, Matthew A, Andrew Haughwout, and Wilbert Van Der Klaauw, "Land Use Regulation and Welfare," *Econometrica*, 2014, *82* (4), 1341–1403.

Appendices for Online Publication

A Additional Tables and Figures

Figure A1: Transaction Volume of Preapproved Properties

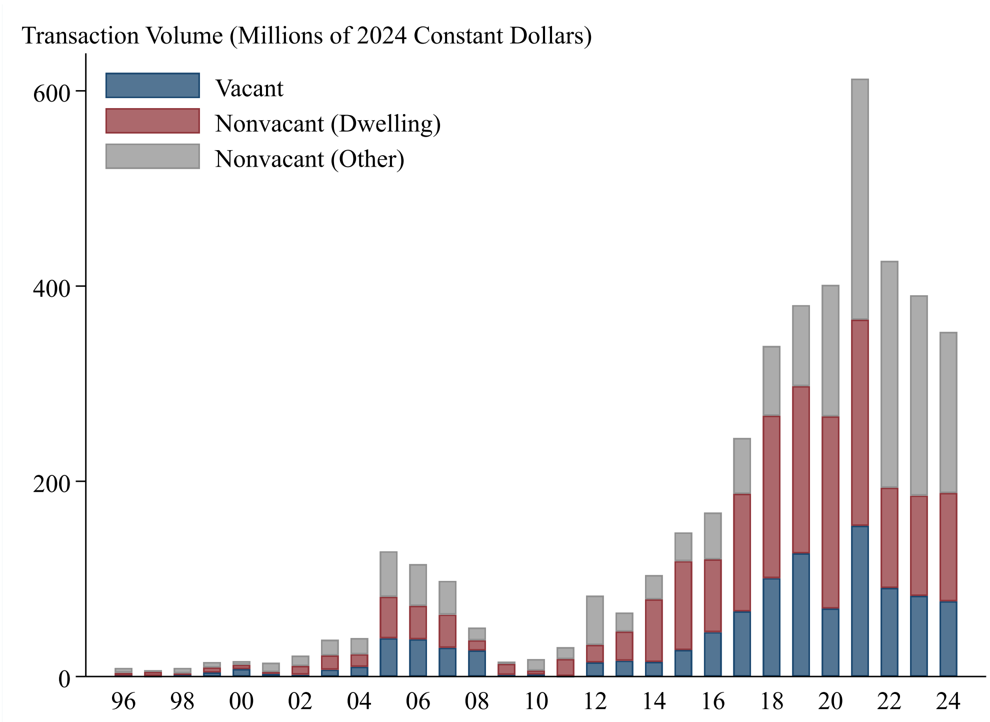

Notes: This figure displays the time series of the total transaction volume of preapproved properties. We split the data by the presence of an existing structure and the content of approvals (dwelling or other). All values are in millions of 2024 constant dollars, adjusted using the Consumer Price Index.

Figure A2: Time-to-Build Versus Time-to-Permit Across Cities

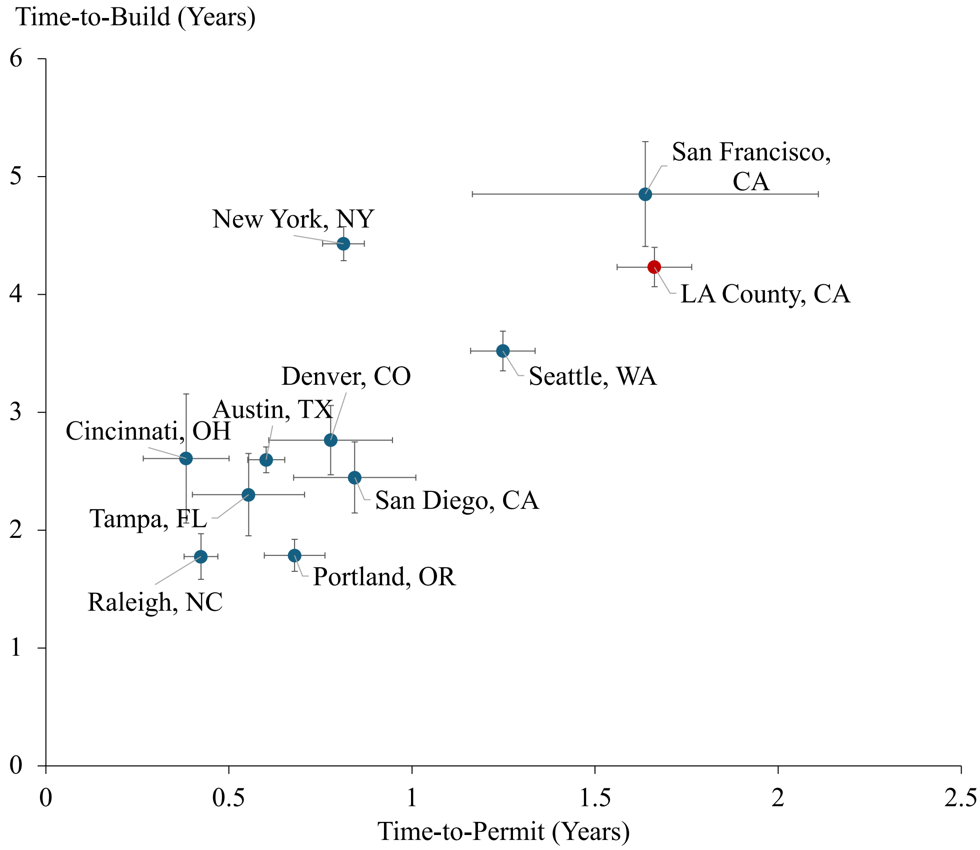

Notes: This figure displays estimates of the time-to-build and time-to-permit for a standardized apartment building across U.S. cities standardized apartment building (30 units, 2019 permit submission, 6000 people per square mile tract density; see Figure 1 for further details).

Figure A3: Additional Maps of Preapproval Rates

(a) Include Permits in Progress (Vacant Lots)

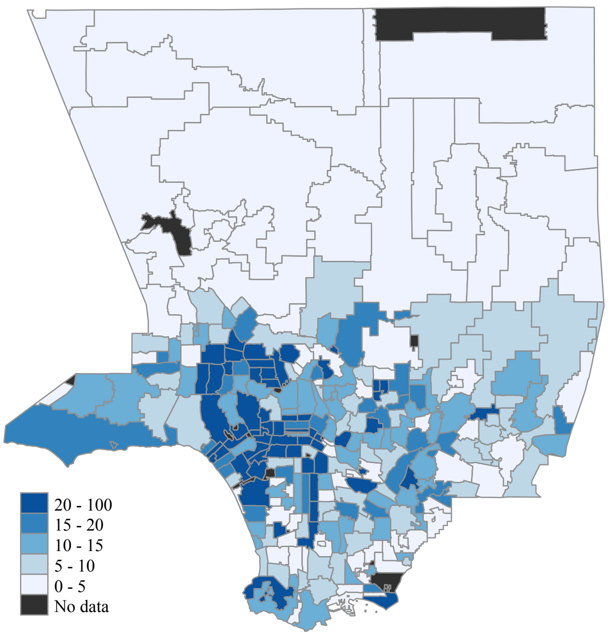

(b) Nonvacant Lots

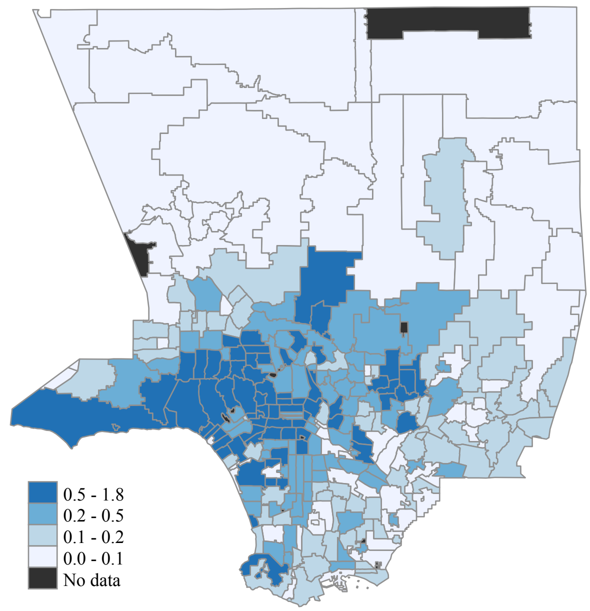

Notes: This figure displays two maps of preapproval rates across zipcodes in Los Angeles County. Panel A maps preapproval rates including properties that report their permitting to be in progress, using the vacant sample. Panel B maps preapproval rates in the nonvacant sample. Units are percentages of listed properties.

Figure A4: Additional Time Series of Preapproval Rates

(a) Including Permits in Progress (Vacant Land) (b) Nonvacant Properties

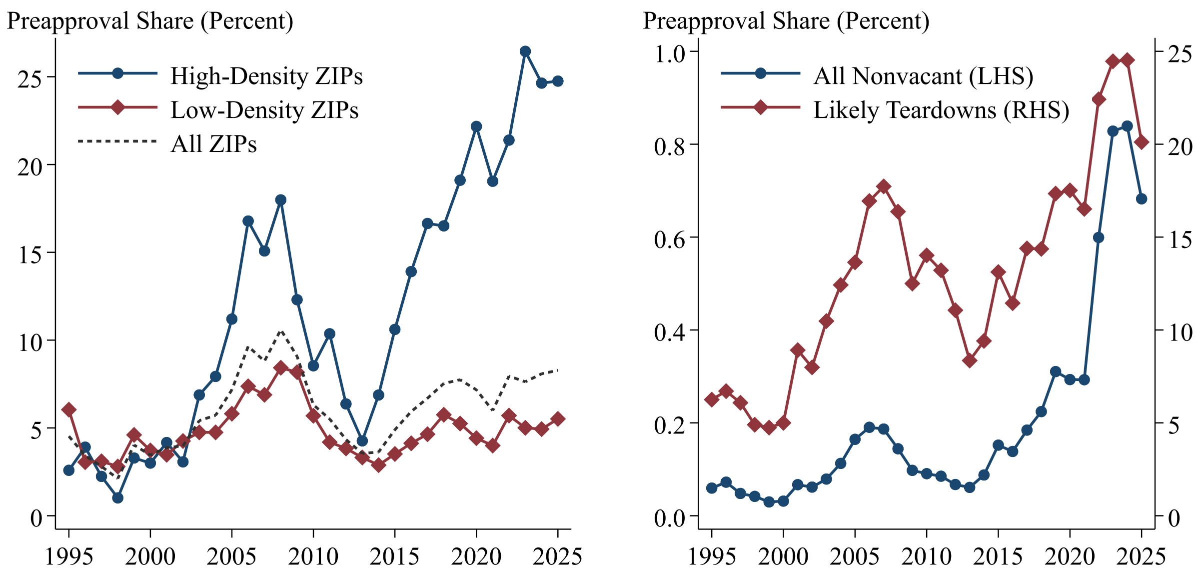

Notes: This figure displays the share of listings with preapproved permits over time. Panel A differs from Panel A of Figure 3 only in the definition of preapproval: Here we include properties with permits in progress as preapproved. Panel B displays two time series of preapproval shares for nonvacant properties: all nonvacant properties (blue circles) and limiting the sample to "likely teardowns" (red diamonds) based on the LLM classification of the property's value proposition as inferred from the remarks. All other nonvacant results in this paper refer to the likely-teardown sample.

Figure A5: Preapproval in the Permitting Process

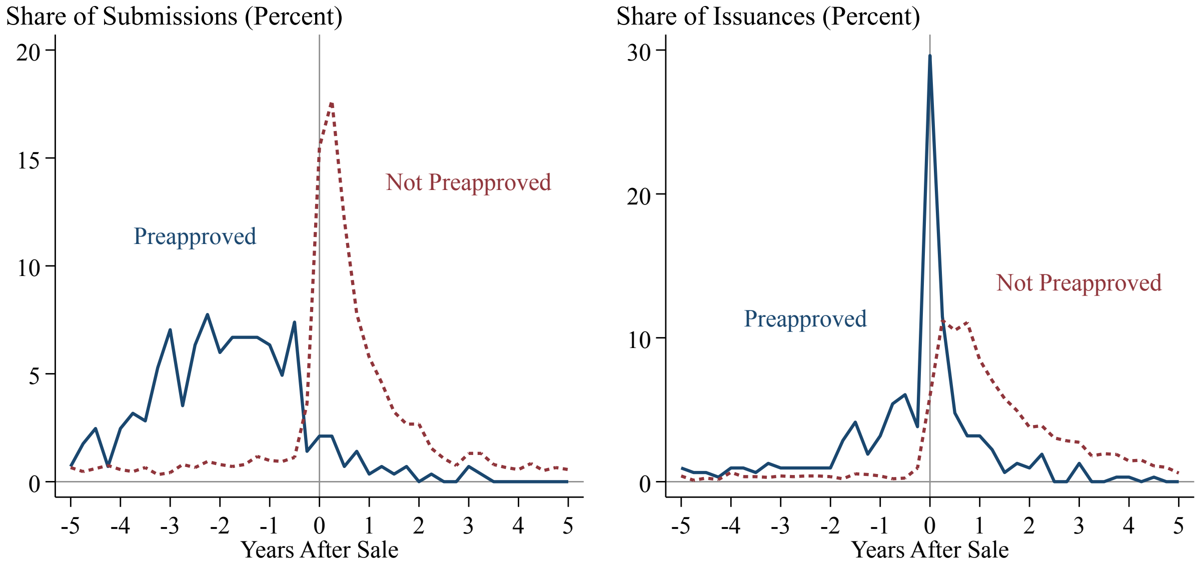

Notes: This figure displays histograms of the timing of permit submittal (left panel) and issuance (right panel) relative to sale. Shares are expressed in quarters of years, with shares computed conditional on any submittal or issuance within a range of five years around the sale. This window captures 74 percent of matched permit submissions and 83 percent of matched permit issuances. Histograms are plotted for properties according to whether they are preapproved (solid blue line) or not preapproved (dashed red line). The sample is all vacant properties that ever submit a permit.

Figure A6: Event-Study Estimates of the Approval Premium (Nonvacant Properties)

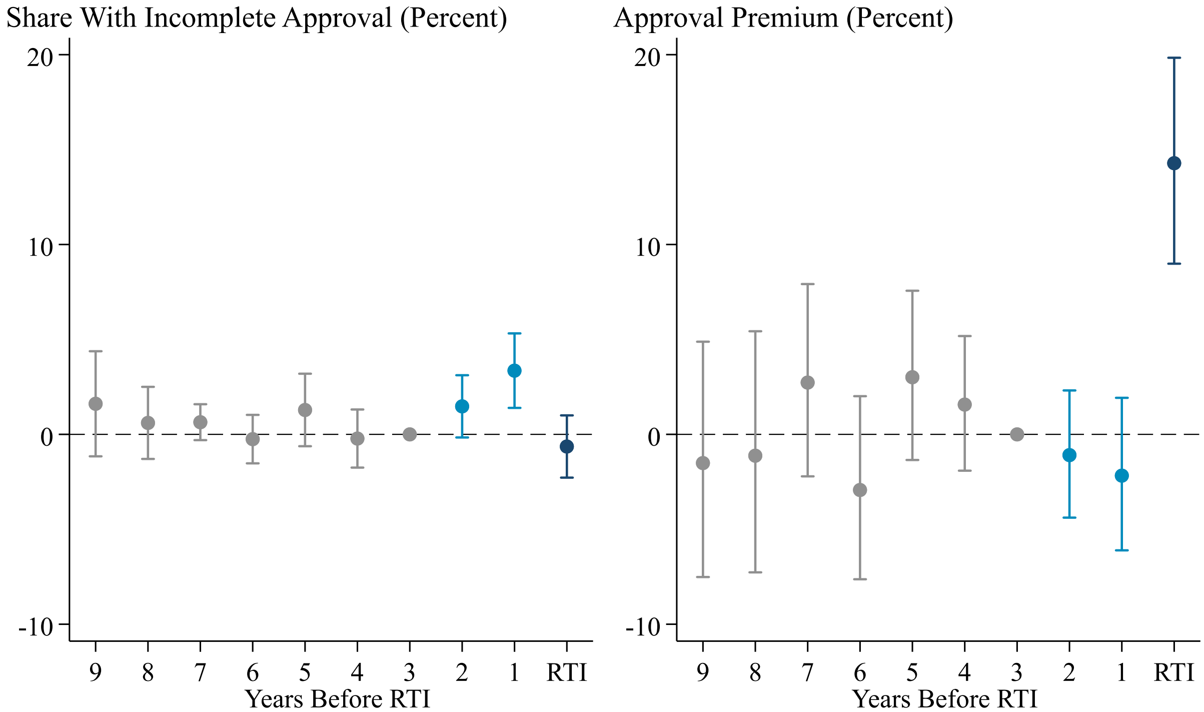

Notes: This figure displays event-study coefficients from Equation 13. It is equivalent to Figure 6 but for using the nonvacant-property sample. In Panel A, the outcome is an indicator for whether the listing indicates permitting is in progress. In Panel B, the outcome is the log price. Coefficients in Panel B are exponentiated, with the figure showing delta-method estimates of $100[\exp(\beta) - 1]$. Both panels include nine leads of an indicator for the property being fully approved. Bars indicate 95-percent pointwise confidence intervals, with standard errors clustered by parcel.

Figure A7: Robustness to Sample Definition and Controls

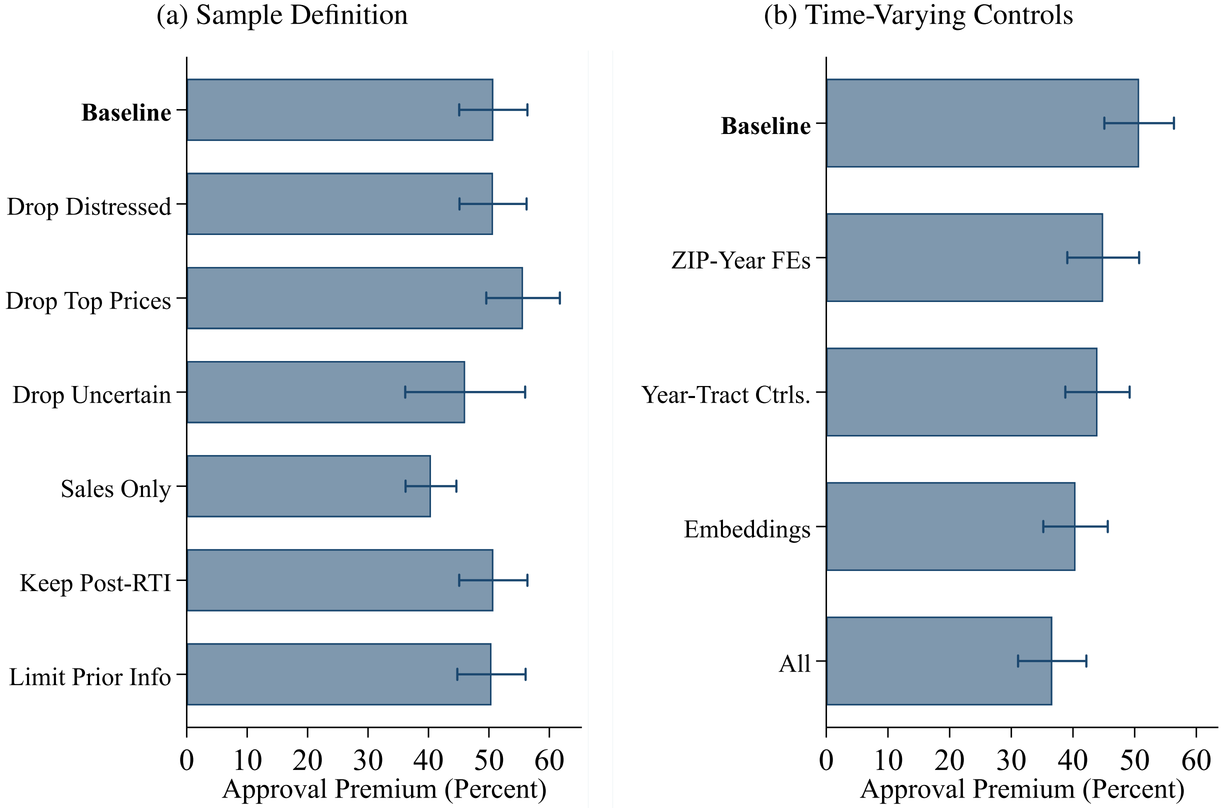

Notes: This figure reports estimates of the approval premium for the sample of vacant land (from Equation 12) under alternative sample definitions (Panel A) and with additional controls (Panel B). "Drop top prices" excludes land with a price above \$5 million. All coefficients are exponentiated, with the figure showing delta-method estimates of $100[\exp(\beta) - 1]$. Bars indicate 95-percent pointwise confidence intervals, with standard errors clustered by parcel.

Figure A8: Adjustments to Standard Errors for Spatial Correlation

(a) Conley (1999)

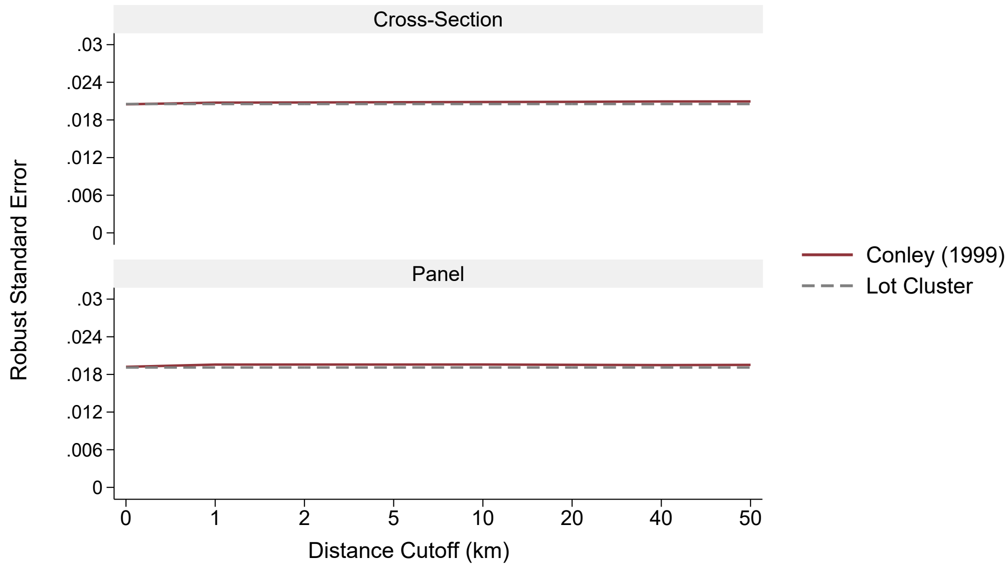

(b) Müller and Watson (2023)

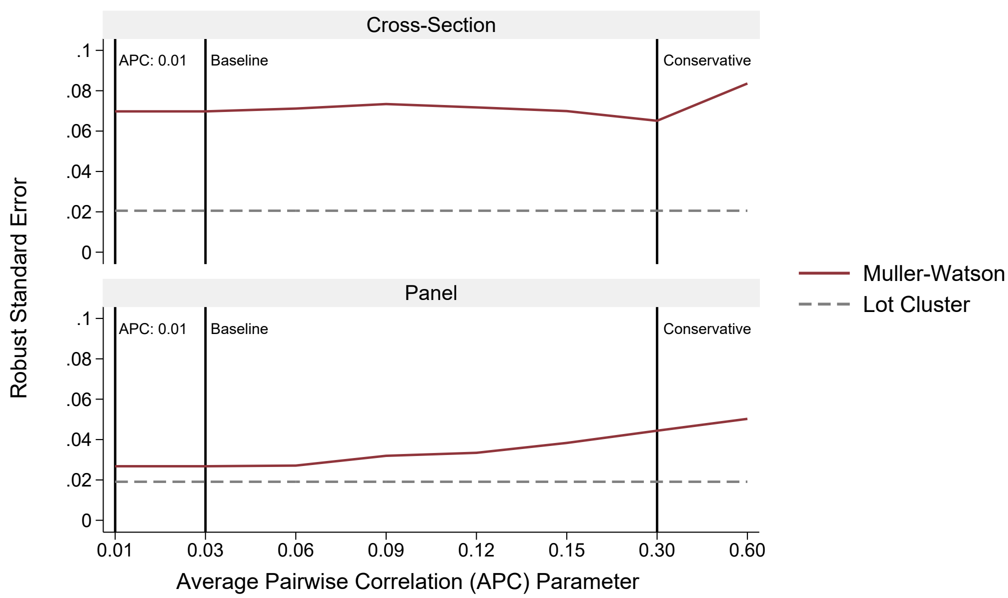

Notes: This figure displays two sensitivity analyses for our estimates of the approval premium with respect to spatial correlation. Panels A and B respectively implement methods from Conley (1999) and Müller and Watson (2023). Both panels plot estimates of standard errors on estimated premia from our cross-sectional and repeat-listing specifications (Equations 11 and 12) as functions of tuning parameters in these methods that determine the extent of spatial correlation. Dashed horizontal lines plot estimated lot-clustered standard errors for both specifications. In Panel A, the tuning parameter is the bandwidth of the Bartlett kernel. In Panel B, the tuning parameter is related to the average pairwise correlation (APC) between lot-level residuals. Solid vertical lines plot empirical APCs, as well as "baseline" and "conservative" values recommended in Müller and Watson (2023).

Figure A9: Reweighted-on-Observables Estimates of Approval Premia

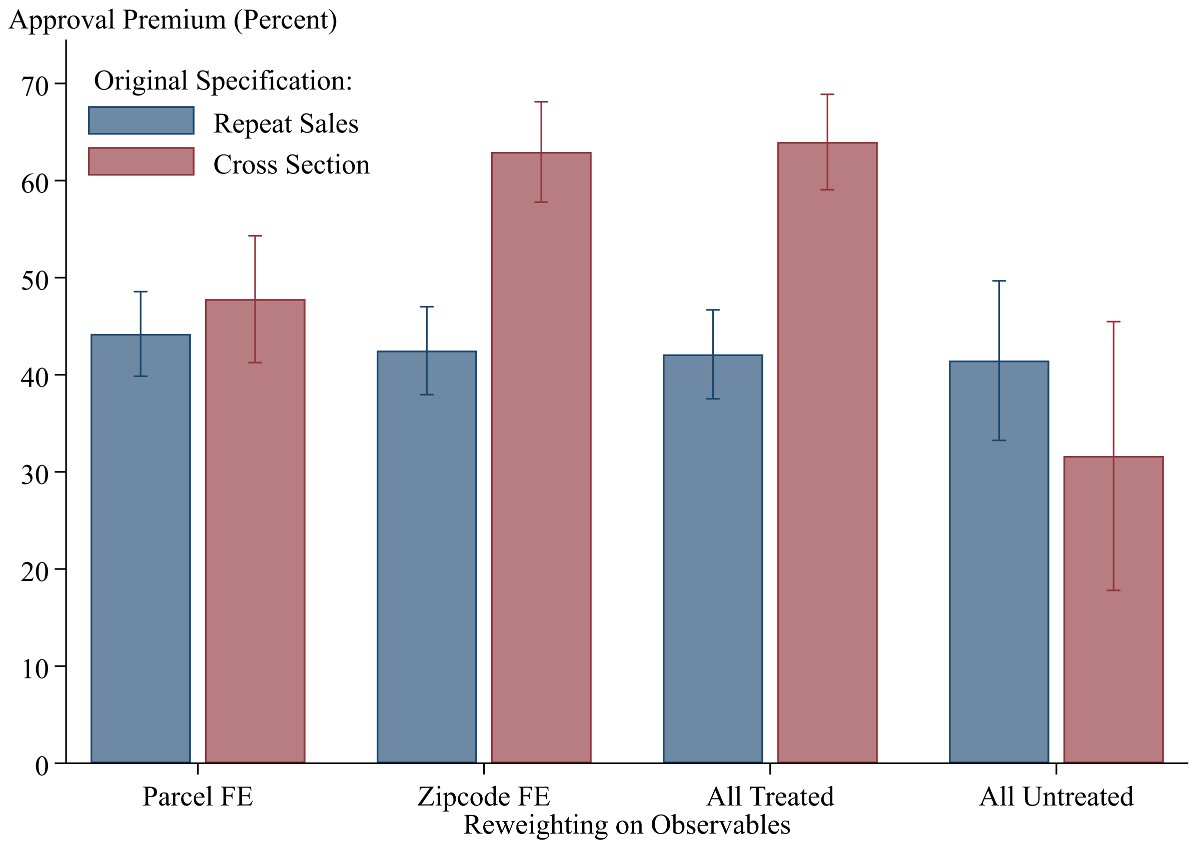

Notes: This figure displays estimates of approval premia reweighted such that the sample providing identifying variation resembles various populations of interest. Red and blue bars respectively report estimates from Equations 11 (cross-section) and 12 (repeat-listing difference-in-differences). Labels indicate the target distribution for each reweighting exercise: "Repeat Listing Sample" reweights the data to match observables of properties with multiple sales; "Cross Section Sample" reweights to match the broader cross-sectional identifying sample. "All Treated" and "All Untreated" reweight to match the observables of all listings with and without preapprovals, respectively. For each estimator's native identifying sample, the reweighting is trivial (weights equal to one). All coefficients are exponentiated, showing delta-method estimates of $100[\exp(\beta) - 1]$. We report 95-percent confidence intervals obtained with standard errors clustered by property.

Figure A10: Sensitivity of Approval Premium to Minimum Price Threshold

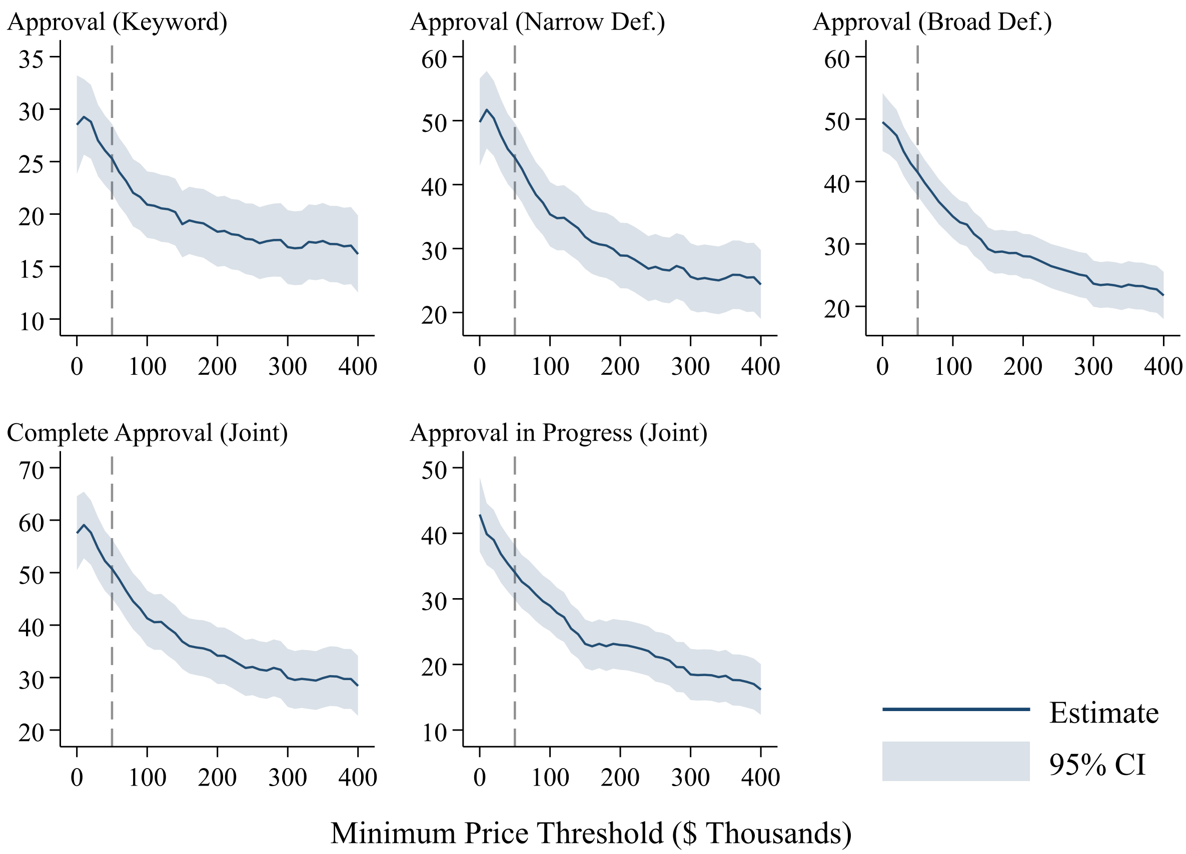

Notes: This figure displays estimated approval premia from Equation 12 and their 95% confidence intervals (CIs), with the horizontal axis varying the minimum-price threshold used to remove outlier or non-arm's-length transactions from the sample. Each panel of the figure shows estimates for a different definition of permit status (see the main text). All estimates are exponentiated, showing $100[\exp(\beta) - 1]$. Dashed gray lines in all panels indicate $50,000, the threshold used throughout the paper. Confidence intervals are obtained via delta method and reflect standard errors clustered by property.

Figure A11: Estimates of the Approval Premium for Nonvacant Properties

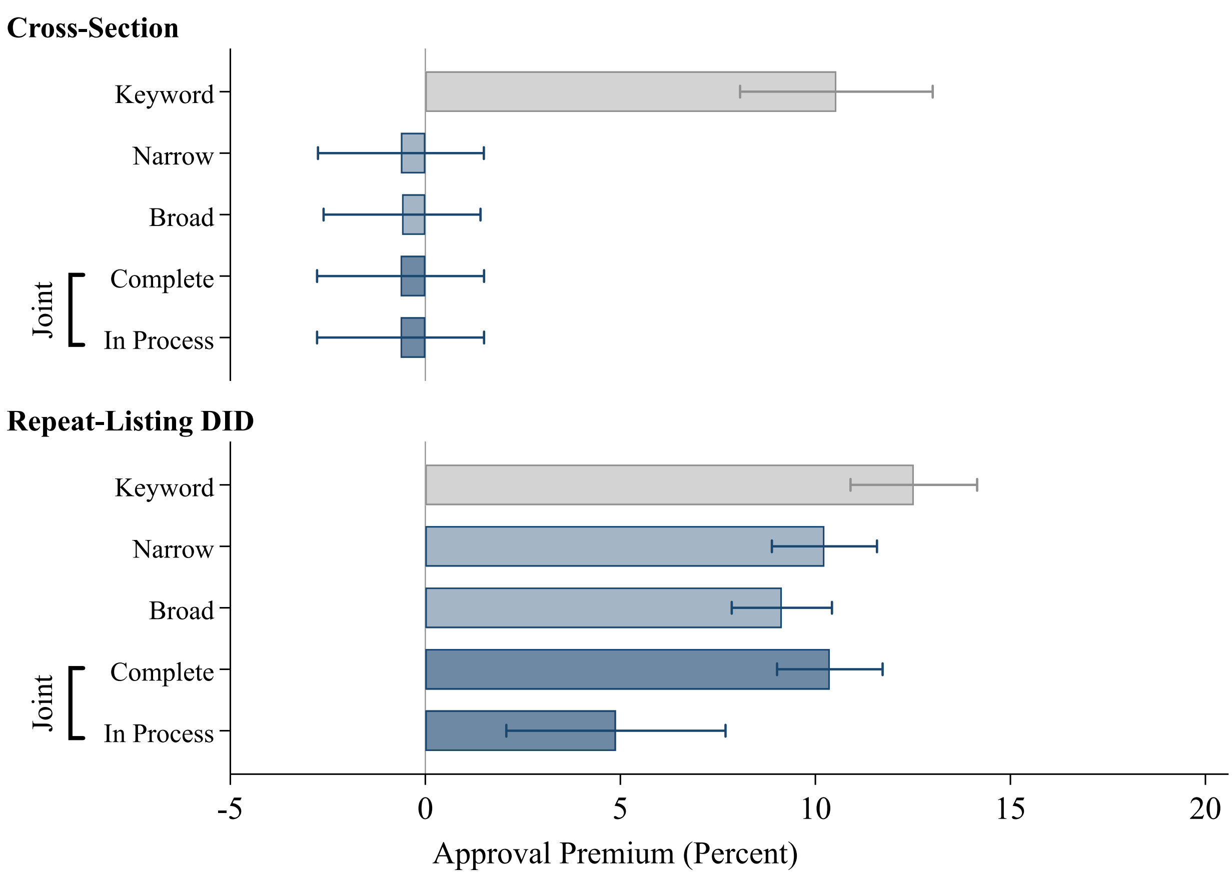

Notes: This figure displays estimated approval premia from Equations 11 (cross-section) and 12 (repeat-listing difference-in-differences). All coefficients are exponentiated, showing delta-method estimates of $100[\exp(\beta) - 1]$, and reflect the sample of nonvacant lots. See the main text for definitions of the measure of permit status. Bars reflect 95-percent confidence intervals with clustering by property.

Figure A12: Heterogeneity in Approval Premia by Neighborhood Characteristics

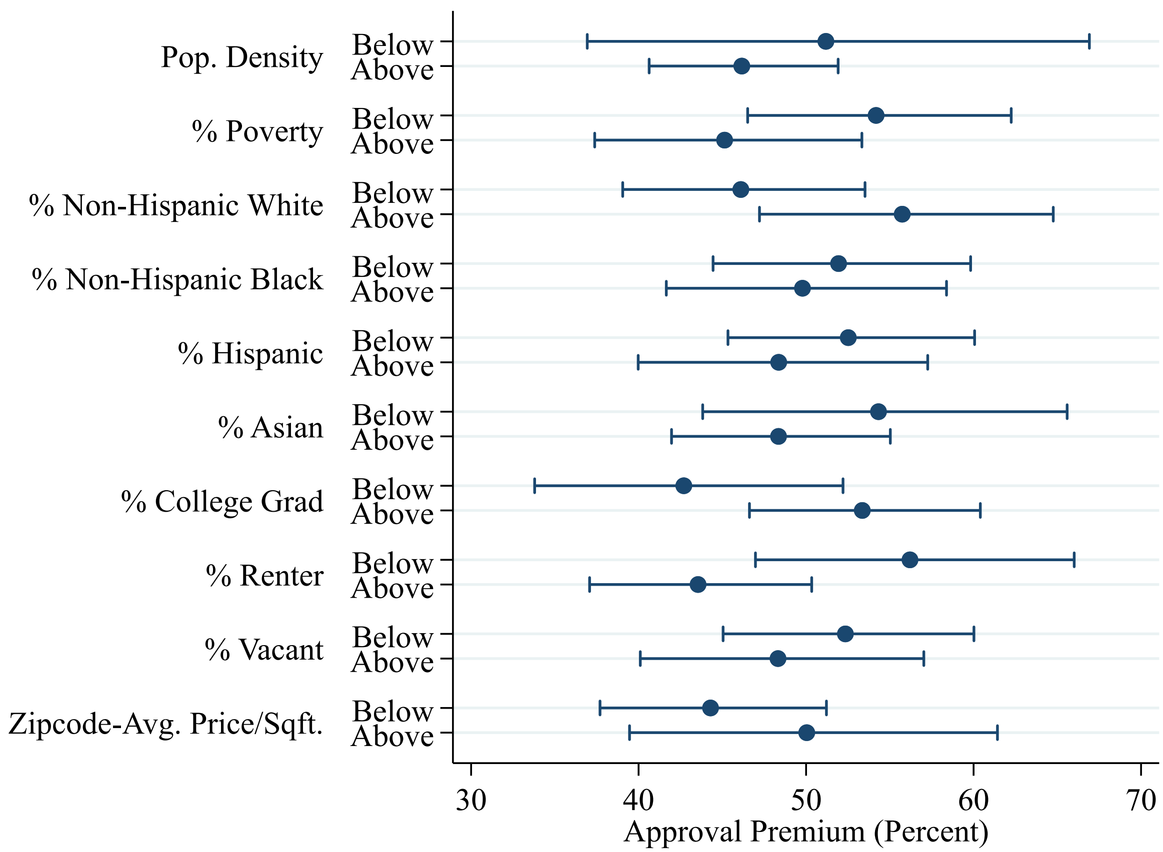

This figure displays estimated effects of approval on land prices, modifying Equation 12 to allow for heterogeneous effects by neighborhood characteristics. The specification includes indicators for whether the parcel's specific neighborhood characteristic is above the median and interacts these with permit status. Bars indicate 95-percent confidence intervals, with standard errors clustered at the parcel level.

Figure A13: Approval Premia Over Time

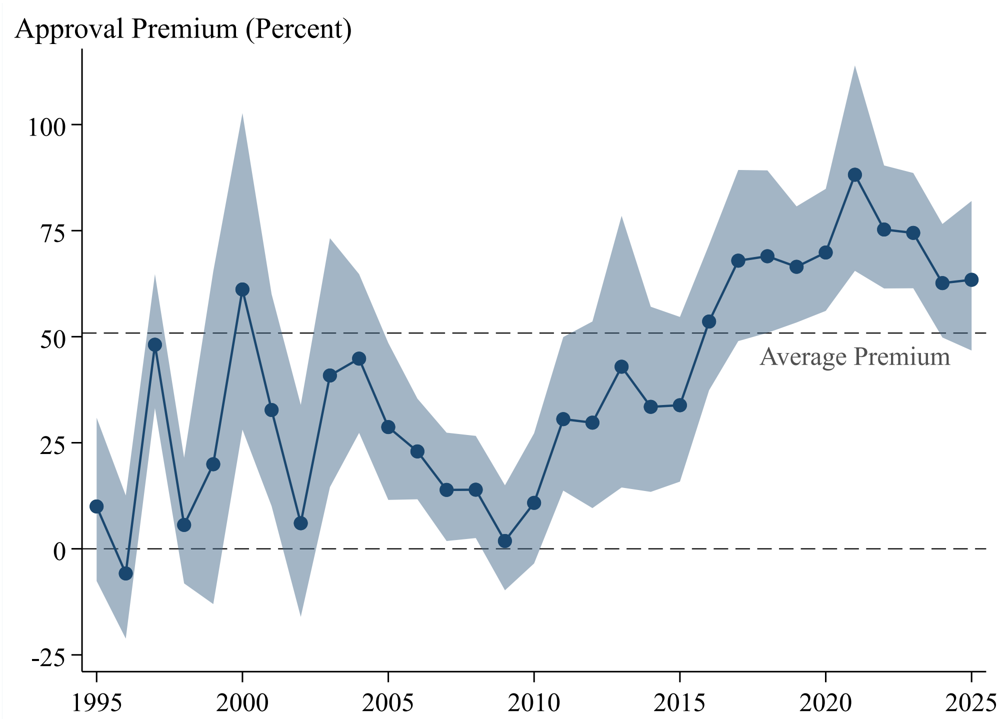

Notes: This figure displays year-specific estimates of the approval premium. Estimates are obtained by augmenting Equation 12 with interactions of permit status with year indicators. All coefficients are exponentiated, showing $100[\exp(\beta) - 1]$, and reflect the vacant-property sample. The color band reflects 95-percent pointwise confidence intervals, estimated by the delta method, with standard errors clustered at the property level.

Figure A14: Approval Premium by Zoning Class

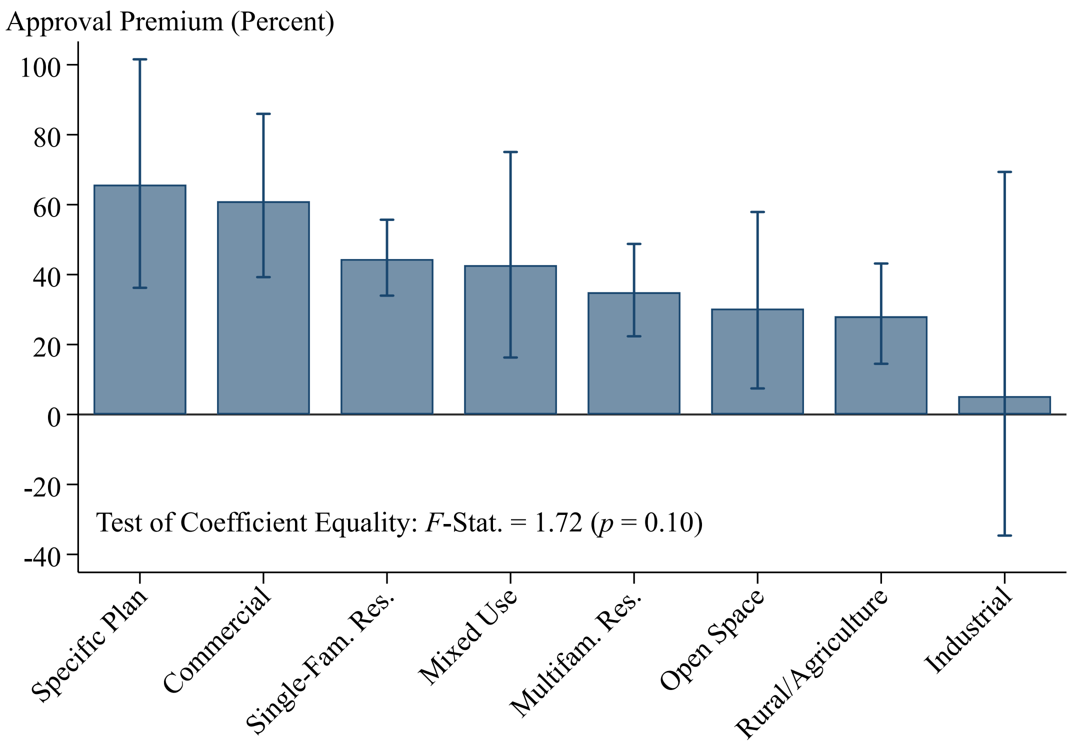

Notes: This figure displays estimates of approval premia by zoning classification. All bars respectively report estimates from a variant of Equation 12 (repeat-listing difference-in-differences) where permit status is interacted with zoning-class dummies. See Appendix C for definitions of zoning classes. All coefficients are exponentiated, showing delta-method estimates of $100[\exp(\beta) - 1]$. We report 95-percent confidence intervals with standard errors clustered by property.

Figure A15: Noncompletion Probability for Proposed Apartment Buildings

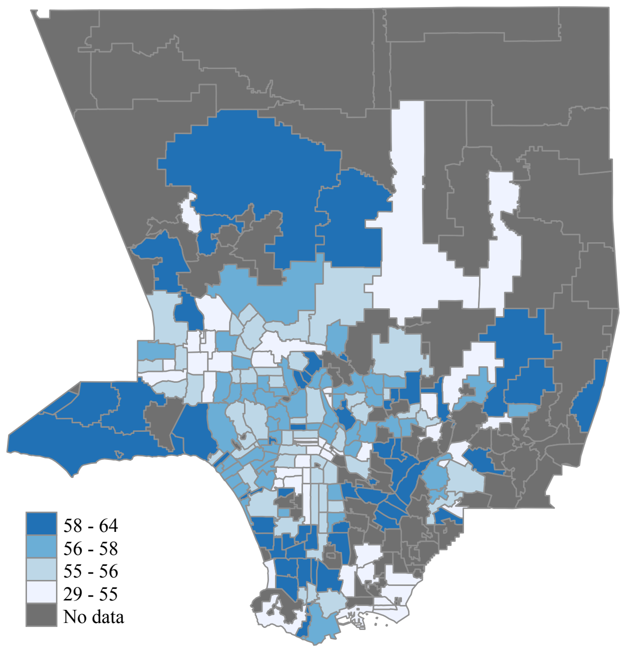

Notes: This figure displays the predicted share of proposed apartment buildings that do not complete construction, given permit submission or issuance, by zipcode. Probabilities are estimated using a 30-unit apartment building on the subsample of multifamily housing permits. See Appendix C for estimation details.

Figure A16: Overview of Imputed Construction Costs and Housing Cost Wedge

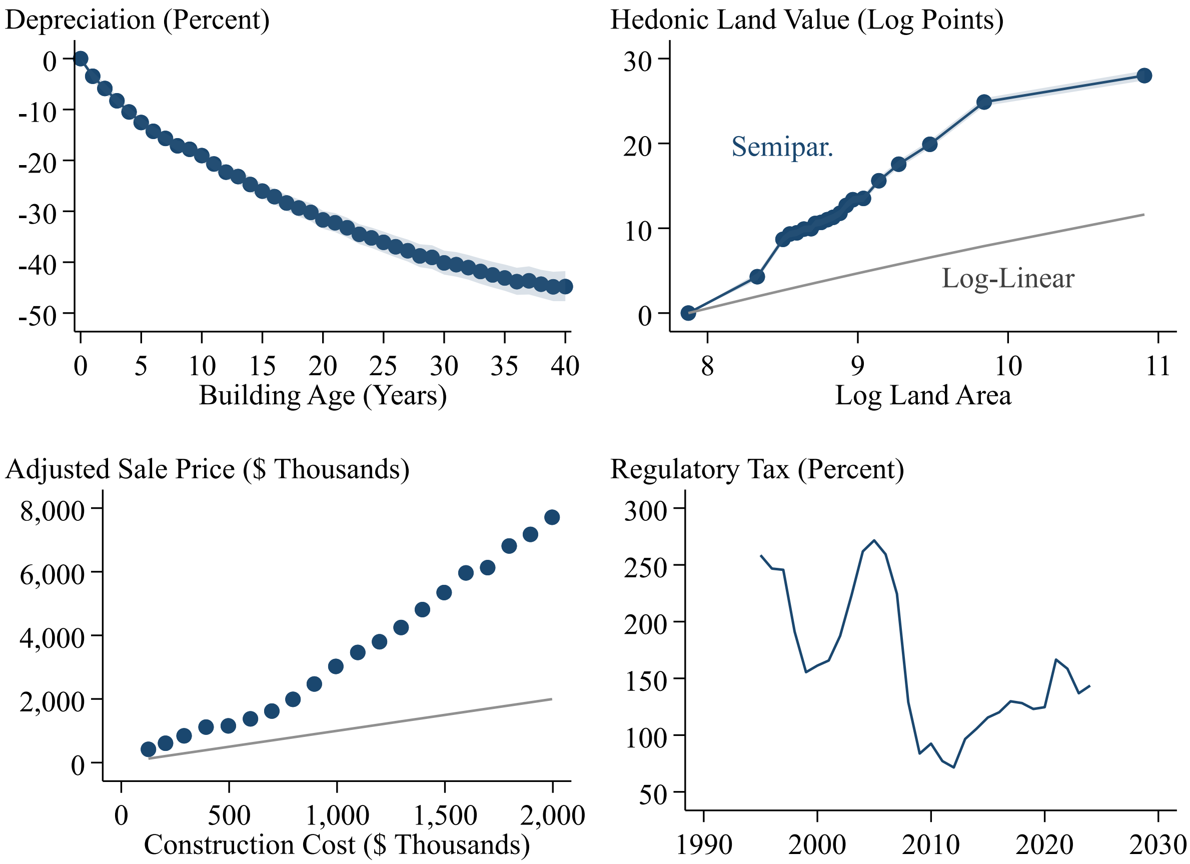

Notes: This figure displays several outputs of the imputation process for construction costs, by which we estimate the housing cost wedge. Panel A displays the depreciation schedule, estimated via a repeat-listing regression we use to adjust the sale prices of existing housing into their "as-if-new" values. Panel B displays the semiparametric estimates of the hedonic value of land, expressed in log points; we compare these values to the estimated hedonic value of land if it is assumed to be linear in log area. Panel C shows a binned scatter plot of sale prices, adjusted for depreciation and hedonic land value, against construction cost. Panel D shows annual medians of the housing cost wedge. See Appendix C for further details on construction-cost imputation.

Figure A17: Heterogeneity in the Approval Premium by Property Risk Factor

(a) Lead Pollution

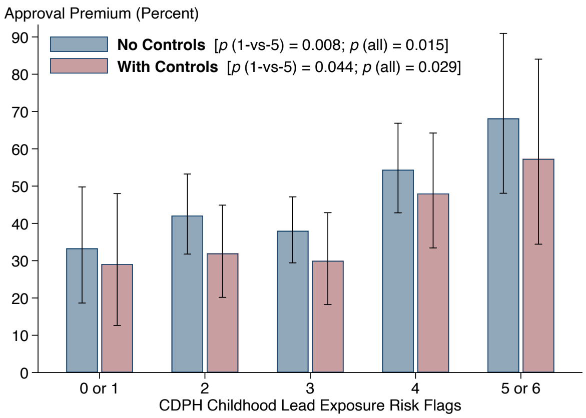

(b) Environmental Review Intensity

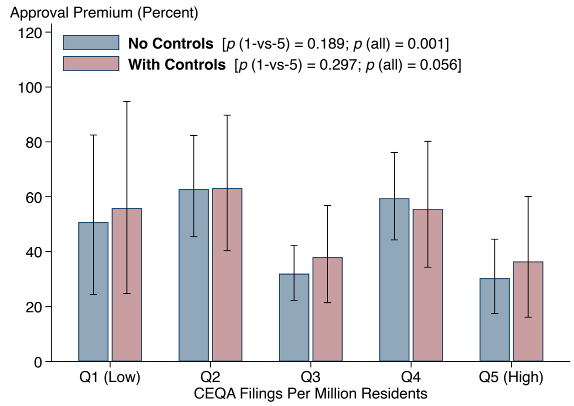

(c) Wildfire

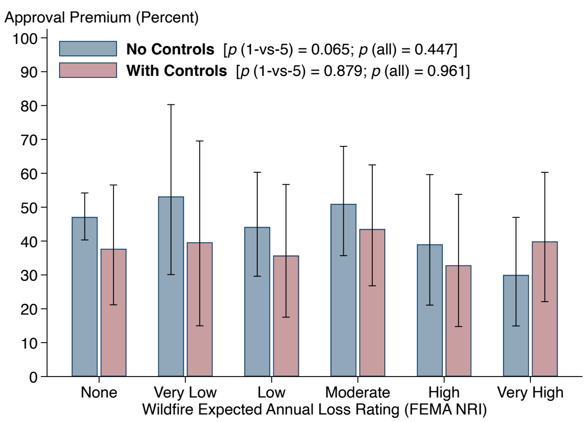

Notes: Panels A–C report estimates of heterogeneous approval premia from a variant of the repeat-listing specification (Equation 12). Blue bars present estimates without other controls. Red bars include interactions between permit status and a vector of other neighborhood-level covariates (including population density and median household income), all demeaned to preserve the interpretation of the main effects. Both sets of bars display 95-percent confidence intervals with standard errors clustered by property. In Panel A, the lead pollution index is the California Department of Public Health (CDPH)'s count of geospatial indicators of childhood lead exposure risk. In Panel B, the environmental-review index is the count of California Environmental Quality Act (CEQA) environmental impact reports (EIR, FIN, JD, and SIR codes) per resident. In Panel C, the wildfire index is taken from the U.S. Federal Emergency Management Agency (FEMA) National Risk Index (March 2023 edition). For each specification, we report two p-values from tests of coefficient equality: the highest category (t-test) versus the lowest category, and all categories being equal (F-test).

Figure A18: Geographic Variation in the Approval Premium (Alternative Units)

(a) Percent Premium on Raw-Land Value (b) Dollar Value per Square Foot of Land

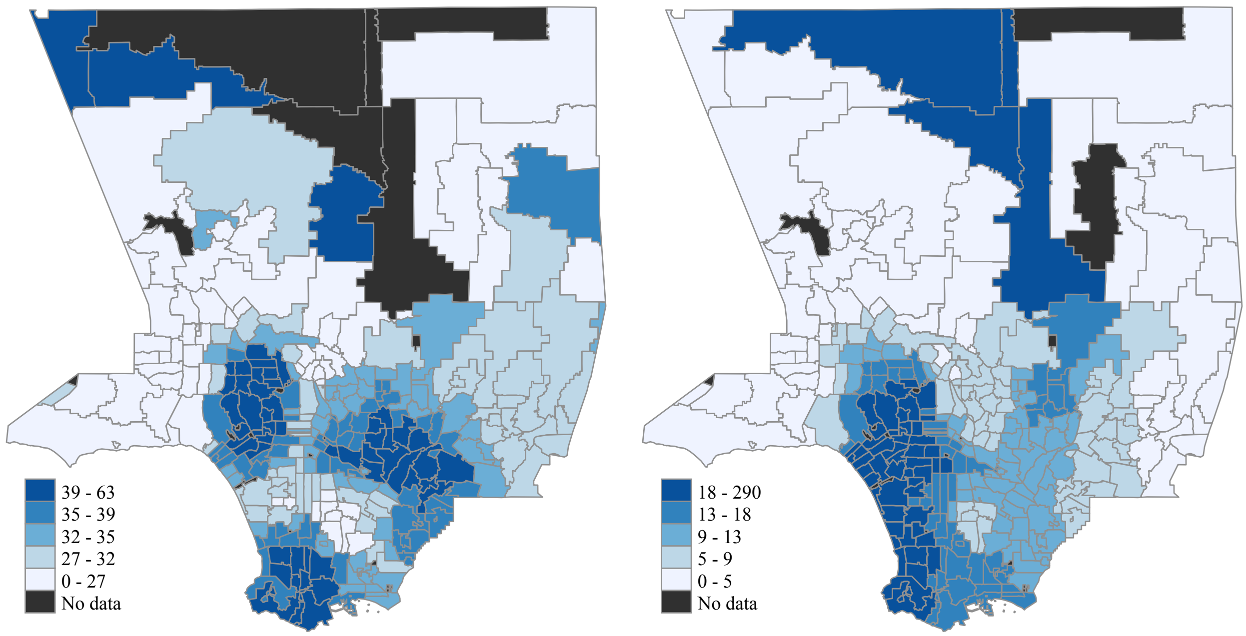

Notes: This figure displays local-area estimates of the approval premium. Estimates are constructed using geographically weighted regressions (see Appendix C). Panel A shows exponentiated coefficients $100[\exp(\beta) - 1]$, representing the percentage markup relative to the counterfactual raw-land value. Panel B transforms these into the estimated implicit dollar price of a permit per square foot of land, calculated by applying the local percentage premium to the local predicted raw-land value.

Figure A19: Estimated Time Savings by Completion Probability

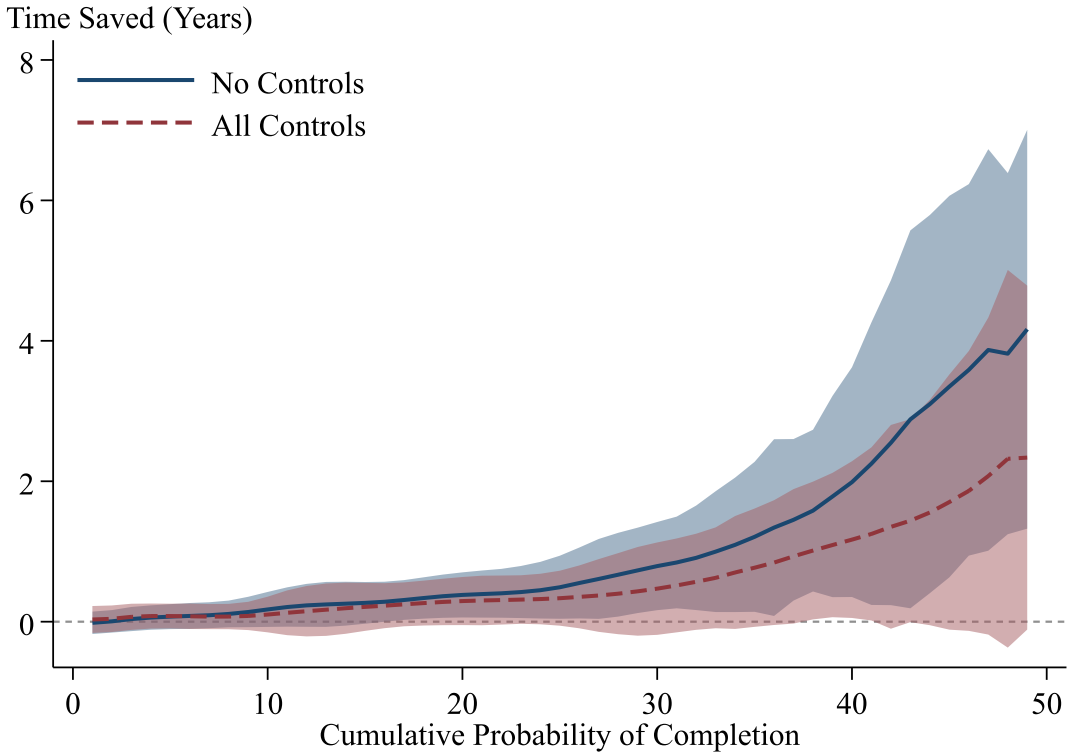

Notes: This figure plots the estimated pull-forward effect of preapproval on completion, expressed in years saved. The horizontal axis reports the cumulative probability of completion, and the vertical axis reports the horizontal distance (in years) between the actual and counterfactual completion curves. For each percentage-point increment in completion probability from 1 to 50 percent, we solve for the time t such that $\Pr(\text{Completion} \leq t) = p$, using linear interpolation between estimated quarterly horizons. This inverse is not well-defined beyond the maximum counterfactual completion probability (see Figure 8); we do not attempt to estimate time savings beyond that level. The navy solid line displays results for a baseline specification without controls. The maroon dashed line includes the full set of lot, tract, project, and listing attributes. Shaded bands indicate 95-percent pointwise confidence intervals, with standard errors calculated by a nonparametric bootstrap with clustering by property.

Figure A20: Distribution of Time to Permit Issuance for Preapproved Properties

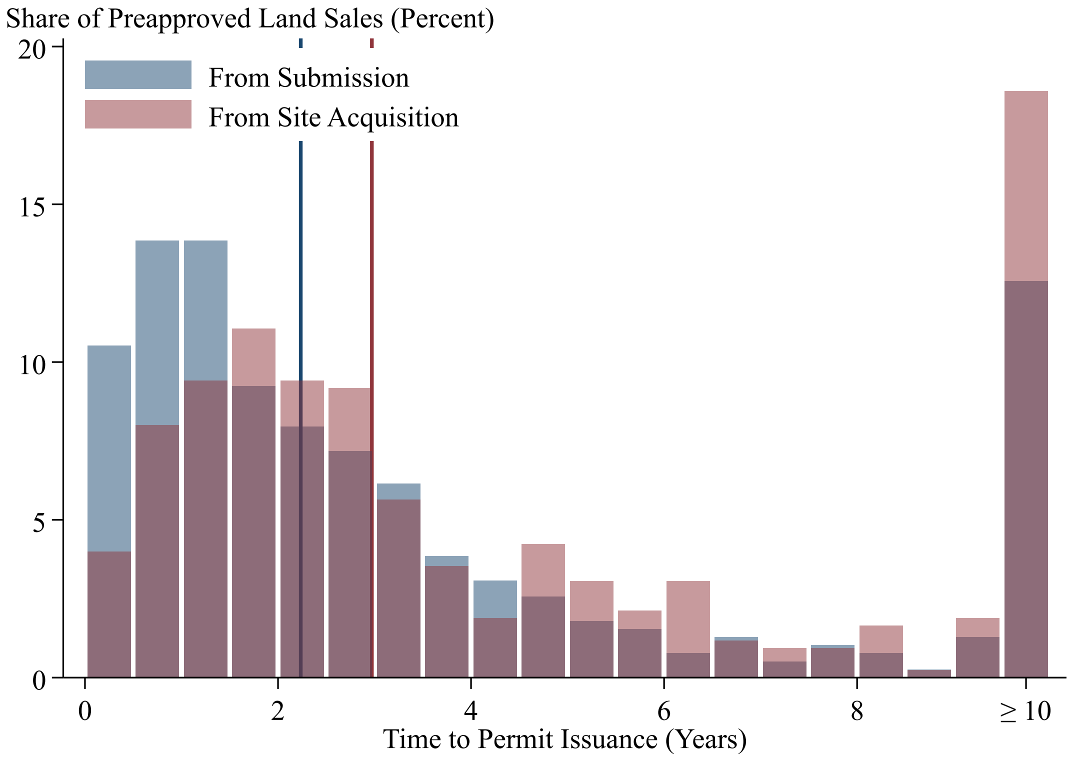

Notes: This figure displays distributions of time elapsed (in years) in the permitting process for vacant land that eventually sells after approval. We measure duration starting either from permit submission (blue bars) or site acquisition (red bars). For each respective distribution, solid lines in the corresponding color indicate median durations.

Figure A21: Dynamic Effects of Preapproval on Completion (Alternative Sample Definition)

(a) Estimated Effects (b) Actual Versus Counterfactual

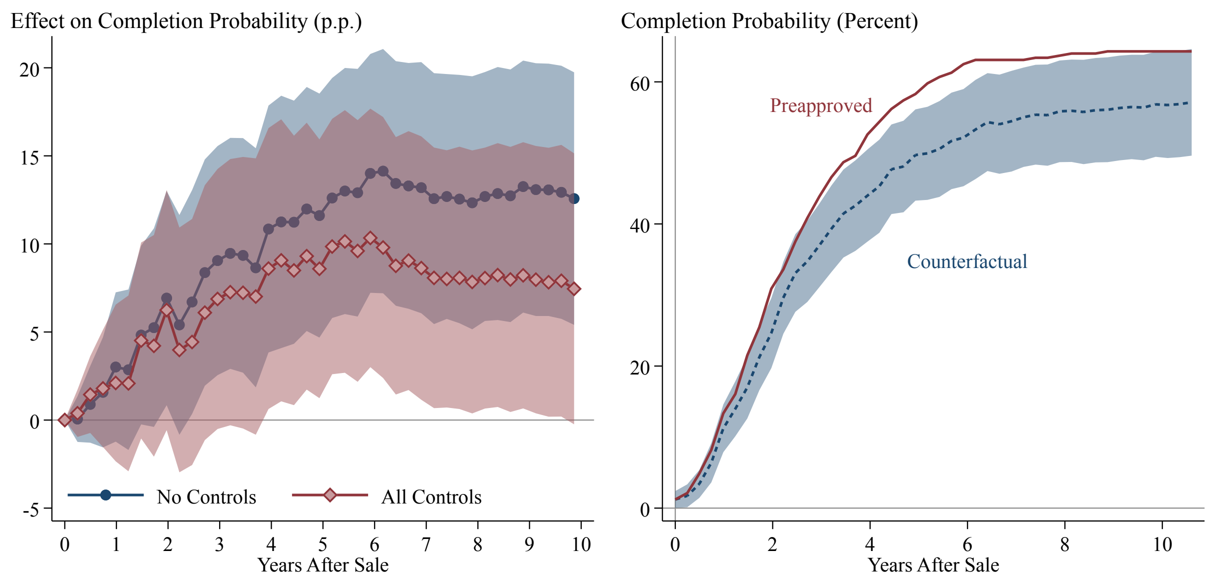

Notes: This figure displays estimated effects of preapproval on time to completion. It differs from Panels A and B of Figure 8 only in the sample definition: Here we limit the sample to listings with matched issued permits for new construction, demolition, or grading. Panel A plots quarterly estimates $\hat{\Delta}(t)$ from Equation 16 from zero to ten years after site acquisition, showing results with and without controls. Panel B plots the empirical cumulative completion hazard for preapproved properties (red solid line) and then subtracts off the estimated effects of preapproval to obtain the counterfactual completion profile in time (blue dashed line) The sample is all vacant properties that ever submit a permit. Color bands indicate 95-percent pointwise confidence intervals, with standard errors calculated by a nonparametric bootstrap clustered by property.

Figure A22: Time-to-Build Effects in the City of Los Angeles Versus Rest-of-County

(a) Dynamic Path of Treatment Effects (b) By Sample, Controls, and Time Horizon

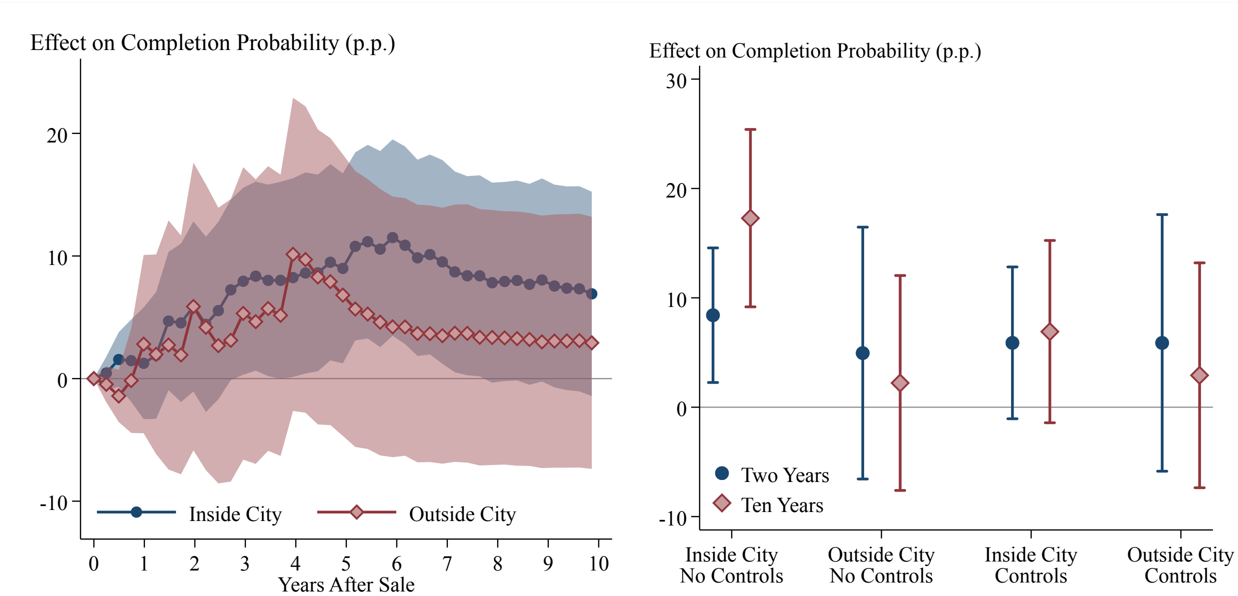

Notes: This figure displays estimated effects of preapproval on time to completion (Panels A and B). The panels differ from Figure 8 in splitting the sample according to location: whether the property is inside or outside the limits of the City of Los Angeles. Panel A plots quarterly estimates $\hat{\Delta}(t)$ from Equation 16 from zero to ten years after site acquisition, showing results with and without controls. Panel B plots the empirical cumulative completion hazard for preapproved properties (red solid line) and then subtracts off the estimated effects of preapproval to obtain the counterfactual completion profile in time (blue dashed line) The sample is all vacant properties that ever submit a permit. Color bands indicate 95-percent pointwise confidence intervals, with standard errors calculated by a nonparametric bootstrap clustered by property.

Table A1: Rents and Home Prices in Selected U.S. Cities

	American Community Survey (2019–23)			Zillow (Jan. 2025)	
	Gross Rent (Median)	Gross Rent (Mean)	Home Price (Median)	Typical Home Price (ZHVI)	Typical Rent (ZORI)
United States	$1,348	$1,448	$303,400	$360,183	$1,889
Boston-Cambridge-Newton, MA-NH	$1,940	$1,951	$610,900	$703,078	$2,907
	43.9%	*34.8%*	*101.4%*	*95.2%*	*53.9%*
Los Angeles-Long Beach-Anaheim, CA	$1,987	$2,087	$825,300	$950,074	$2,857
	47.4%	*44.2%*	*172.0%*	*163.8%*	*51.3%*
New York-Newark-Jersey City, NY-NJ	$1,780	$1,909	$587,400	$681,777	$3,146
	32.0%	*31.9%*	*93.6%*	*89.3%*	*66.5%*
San Diego-Chula Vista-Carlsbad, CA	$2,154	$2,234	$791,600	$935,140	$2,895
	59.8%	*54.3%*	*160.9%*	*159.6%*	*53.2%*
San Francisco-Oakland-Fremont, CA	$2,426	$2,506	$1,113,800	$1,124,852	$2,942
	80.0%	*73.1%*	*267.1%*	*212.3%*	*55.7%*
San Jose-Sunnyvale-Santa Clara, CA	$2,794	$2,815	$1,342,700	$1,586,416	$3,265
	107.3%	*94.5%*	*342.6%*	*340.4%*	*72.8%*

Notes: This table reports several measures of rents and home prices in six selected U.S. cities as well as the national average. Italicized figures report premia for that city and measure, in percentage terms, over the corresponding national average.

Table A2: Sample Construction

	Vacant		Nonvacant	
	Listings	Unique Parcels	Listings	Unique Parcels
	(1)	(2)	(3)	(4)
Raw data	4,920,848	1,236,897	4,920,848	1,236,897
Keep only sales	4,193,485	1,203,399	4,193,485	1,203,399
Keep listings with non-missing parcel ID	4,920,848	1,236,897	4,920,848	1,236,897
Keep listings with relevant property type	130,334	45,468	4,063,151	1,171,400
Keep listings with price and date	130,334	45,468	4,063,106	1,171,397
Keep non-duplicated listings	95,719	45,468	2,903,571	1,171,397
Keep listings with non-missing listing dates	95,715	45,466	2,903,539	1,171,391
Keep listings with only one RTI listing	94,961	45,366	2,902,732	1,171,308
Keep listings weakly preceding RTI	93,489	45,366	2,897,394	1,171,308
Keep listings with hybrid or redevelopment value	n.a.	n.a.	126,285	58,995
Keep vacant listings without living area	93,477	45,361	n.a.	n.a.
Keep listings above $50,000	67,377	33,280	119,085	58,155

Notes: This table reports the number of listings and unique parcels identified at each stage of the sample construction process. The relevant property type refers to vacant or non-vacant listings. The notation "n.a." indicates this sample restriction is not applied for this sample.

Table A3: Classification of Value Proposition Based on MLS Remarks

Type	Definition and Example
Current Use	*Definition:* Pitch is based on the property as it is today. *Example:* Immaculate Starter!!! Fully Remodeled 3 Bed, 1ba Home, In Cul-De-Sac, Beautiful Curb Appeal, Huge Gorgeous Yard W/Fruit Trees, Detached 1 Car Garage.
Hybrid	*Definition:* Pitch includes both current use and a future project as valuable paths. *Example:* Three separate small houses and one three garages building on one lot. It produces very steady good income, now, and it could be an ideal future development.
Development Value	*Definition:* Pitch is based on the value of the land and/or entitlements for a future project. *Example:* Development opportunity. R-3 zone. Can build 6 units condo (buyer to verify w/city). Property solid in ”as is” condition. Very convenient location. S.Pas school district.

Notes: This table introduces our classification scheme for a lot's value proposition from the remarks field of listings. For each status, we quote the definition provided to the LLM as well as the full remarks text for a characteristic example.

Table A4: Preapproval Rates by Value Proposition of Lot

	(1)	(2) Vacant	(3)	(4)	(5) Nonvacant	(6)
Preapproval Status	Yes	No	Total	Yes	No	Total
Value Proposition						
Development	61.4	38.6	100.0	20.9	79.1	100.0
	(3,576)	(2,246)	(5,822)	(15,106)	(57,176)	72,282)
Hybrid	57.1	42.9	100.0	15.9	84.1	100.0
	(133)	(100)	(233)	(15,788)	(83,585)	(99,373)
Current Use	35.7	64.3	100.0	0.0	100.0	100.0
	(5)	(9)	(14)	(184)	(484,451)	484,635)

Notes: This table reports preapproval rates by land vacancy and the value proposition of the lot, as inferred from the MLS listing text. "Development" means the listing indicates lot value mostly reflects future use when developed (if vacant) or redeveloped (if nonvacant), "current use" its current use, and "hybrid" both its current and future use. Cells report, for each combination of the value proposition and preapproval, listing counts (in parentheses) and preapproval rates.

Table A5: Preapproval Confusion Matrices: Keyword Versus Large Language Model

	(1)	(2) Vacant	(3)	(4)	(5) Nonvacant	(6)
Keyword Rating	Yes	No	Total	Yes	No	Total
LLM Rating						
Yes	68.6	31.4	100.0	60.0	40.0	100.0
	(2,567)	(1,174)	(3,741)	(3,532)	(2,355)	(5,887)
In Progress	26.3	73.7	100.0	19.4	80.6	100.0
	(623)	(1,744)	(2,367)	(236)	(981)	(1,217)
No	3.3	96.7	100.0	1.0	99.0	100.0
	(2,778)	(81,085)	(83,863)	(849)	(83,532)	(84,381)

Notes: This table reports confusion matrices that compare the preapproval status of lots as classified by keywords (across columns) and by a large language model (LLM, across rows). Columns 1 and 2 report the confusion matrix for vacant lots, and Columns 3 and 4 report the matrix for nonvacant lots. Cells report, for each keyword–LLM rating pair, counts (in parentheses) and shares of the total with respect to LLM classification.

Table A6: Human Validation of Preapproval Ratings by Large Language Model

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Vacant				Nonvacant			
LLM Rating	Yes	In Progress	No	Total	Yes	In Progress	No	Total
Human Rating								
Yes	92.2	58.3	7.6	100.0	85.0	11.0	0.5	100.0
	(1,256)	(448)	(398)	(2,102)	(902)	(11)	(46)	(959)
In Progress	3.3	28.7	3.0	100.0	0.3	78.4	0.0	100.0
	(30)	(245)	(100)	(375)	(27)	(62)	(15)	(104)
No	4.6	13.0	89.4	100.0	14.8	10.7	99.5	100.0
	(63)	(100)	(1,828)	(1,991)	(78)	(4)	(1,765)	(1,847)

Notes: This table reports confusion matrices for the preapproval status of lots, comparing classifications of human raters to those of a large language model (LLM). Human ratings vary by row and the LLM ratings vary by column. Columns 1–4 report the confusion matrix for vacant lots and Columns 5–8 report the confusion matrix for nonvacant lots. Cells report, for each human–LLM rating pair, raw counts (in parentheses) and shares of the total with respect to the LLM-based classification. The shares account for weights used to sample listings for human validation.

Table A7: Development Content of Preapprovals

	(1)	(2)	(3)	(4)
	Vacant		Nonvacant	
	Any	Priority	Any	Priority
Multiple Dwelling Units	38.1	38.1	22.5	22.5
	(2,150)	(2,150)	(1,569)	(1,569)
Primary Dwelling Unit	63.5	58.8	33.8	32.1
	(3,581)	(3,313)	(2,359)	(2,242)
Other Structure	11.7	2.8	12.0	7.1
	(659)	(160)	(835)	(496)
Accessory Dwelling Unit	8.2	0.1	29.8	23.6
	(464)	(4)	(2,081)	(1,646)
Addition to Existing Dwelling	36.2	0.2	45.8	14.8
	(2,078)	(48)	(3,199)	(1,033)

Notes: This table reports summary statistics on the development content of preapprovals. "Any" columns report the shares of lots that indicate, in their MLS listing text, any preapproval for the specified development. "Priority" columns report shares adjusted hierarchically, such that a preapproval only counts for any type if it does not count for any higher-priority type (higher row in the same column). Cells report counts (in parentheses) and content shares of preapproved listings.

Table A8: Neighborhood and Lot Determinants of Preapproval

	(1)	(2)	(3)	(4)	(5)	(6)
Neighborhood (Tract) Attributes						
Log Pop. Dens.	0.012***		0.006***	0.007***		
	(0.001)		(0.001)	(0.001)		
Share Poor	0.109***		0.050**	0.042		
	(0.025)		(0.024)	(0.026)		
Share White	0.040***		-0.005	0.055***		
	(0.009)		(0.009)	(0.015)		
Share Renter	0.076***		0.056***	0.056***		
	(0.010)		(0.010)	(0.012)		
Parcel Attributes						
Log Lot Area		-0.004***	0.000		-0.005***	-0.002***
		(0.000)	(0.000)		(0.001)	(0.001)
Log Floorspace		0.010	0.013**		0.012*	0.018***
		(0.007)	(0.007)		(0.006)	(0.006)
Log Units		0.007	0.001		-0.000	-0.010
		(0.008)	(0.008)		(0.008)	(0.008)
Log Cost		0.004	0.004		0.004	0.006
		(0.004)	(0.004)		(0.004)	(0.004)
Log Height		0.071***	0.059***		0.053***	0.049***
		(0.015)	(0.015)		(0.015)	(0.014)
Zipcode FE				X	X	
Tract FE						X
Year FE				X	X	X
Observations	63,553	63,910	63,553	63,551	63,854	63,501
Unique Lots	31,523	31,719	31,523	31,521	31,689	31,416
R^2	0.027	0.048	0.057	0.066	0.095	0.170
R^2 (within)				0.003	0.033	0.032

Notes: This table reports estimated coefficients from regressions of a binary lot-level indicator of preapproval on neighborhood and parcel attributes. When parcel attributes are included, we augment them with missing-value indicators but do not report estimated coefficients. All estimates reflect the vacant-lot sample. Standard errors are clustered at the parcel level. $^{*} = p < 0.10$, $^{**} = p < 0.05$, $^{***} = p < 0.01$.

Table A9: Selection into Identifying Sample for Repeat-Sales Design

	(1)	(2)	(3)	(4)	(5)	(6)
	Non-Switchers		Switchers		Identifying Pop. (Wt.)	
	Always RTI	Never RTI	Before	After	Treated	Control
Panel A: Property Characteristics at Listing						
Price	2,389,667	781,585	899,278	1,473,068	1,243,723	1,158,379
	(174,537)	(12,730)	(53,152)	(78,124)	(60,985)	(91,216)
Price Per Sqft. Land	159.9	261.6	57.9	99.1	82.5	96.1
	(5.7)	(130.3)	(3.2)	(4.6)	(3.2)	(17.0)
Log Price	13.970	12.712	12.971	13.519	13.384	13.168
	(0.036)	(0.008)	(0.042)	(0.035)	(0.034)	(0.044)
Proposed Floorspace	4,089	691	2,430	3,025	2,504	3,123
	(371)	(27)	(253)	(274)	(211)	(339)
Proposed % Any Residential	56.47	58.05	73.52	71.54	74.12	70.58
	(2.27)	(0.93)	(2.57)	(2.38)	(2.11)	(2.77)
Proposed % 5+ Units	11.41	5.88	8.74	11.49	9.17	13.14
	(1.37)	(0.42)	(1.66)	(1.77)	(1.42)	(2.33)
Proposed Height (feet)	29.0	28.4	28.9	30.1	29.7	29.7
	(0.5)	(1.0)	(0.6)	(0.6)	(0.5)	(0.7)
Lot Size	133,165	272,961	46,481	41,738	43,079	47,674
	(71,258)	(8,416)	(5,422)	(4,066)	(3,743)	(5,031)
Panel B: Neighborhood Characteristics						
Population Density	10,467	5,372	7,703	8,311	7,834	8,596
	(354)	(48)	(281)	(282)	(256)	(361)
% Poor	12.2	10.9	10.4	11.0	10.6	11.1
	(0.3)	(0.1)	(0.3)	(0.3)	(0.2)	(0.3)
% Non-Hispanic White	42.5	41.2	41.6	41.5	40.4	42.9
	(0.9)	(0.2)	(1.2)	(0.9)	(0.9)	(1.1)
% Hispanic	34.3	37.6	35.0	34.9	35.5	33.9
	(0.9)	(0.2)	(1.1)	(0.9)	(0.9)	(1.1)
% Renter	44.6	31.9	36.1	38.0	36.3	38.9
	(0.8)	(0.2)	(0.9)	(0.9)	(0.8)	(1.0)
% Vacant	8.9	9.6	8.5	8.7	8.3	9.0
	(0.2)	(0.1)	(0.3)	(0.2)	(0.2)	(0.3)
% College Graduate	0.5	0.4	0.5	0.5	0.5	0.5
	(0.0)	(0.0)	(0.0)	(0.0)	(0.0)	(0.0)

Notes: This table reports the means and standard errors of the mean (in parentheses) of property and neighborhood variables for six populations. Columns 1 and 2 report these statistics for properties that do not switch approval status while in sample, respectively always approved and never approved. Columns 3 and 4 report these statistics for properties whose approval status switches while in sample, respectively showing the pre-switch and post-switch means. Columns 5 and 6 report weighted means and standard errors for the identifying population. The weights in Columns 5 and 6 are $w_{it} = (\mathrm{RTI}_{it} - \bar{\mathrm{RTI}}_i - \bar{\mathrm{RTI}}_t + \bar{\mathrm{RTI}})^2$, where the bars denote property-level, year-level, and grand mean.

Table A10: Dollar-Value Approval Premia

	Cross-Section		Repeat-Sales DID	
	Mean	Median	Mean	Median
	(1)	(2)	(3)	(4)
Panel A: Univariate Specification				
Per Lot	905,417***	480,436***	701,644***	347,151***
	(55,506)	(21,521)	(45,037)	(16,564)
Per Square Foot	62.0***	36.4***	43.7***	25.0***
	(2.6)	(1.3)	(2.2)	(1.2)
Panel B: Joint Specification (Full RTI)				
Per Lot	936,876***	497,129***	769,932***	380,937***
	(56,145)	(21,731)	(47,292)	(17,049)
Per Square Foot	64.2***	37.7***	48.0***	27.4***
	(2.6)	(1.3)	(2.3)	(1.2)

Notes: This table reports dollar-value mean and median approval premia. To express premia in dollars, we apply the premium estimated in Equations 11 or 12 to the log price, exponentiate, and report statistics for properties that were preapproved. Bootstrap standard errors incorporate re-estimating the approval premium in each replication, resampling at the parcel level. $^{*} = p < 0.10$, $^{**} = p < 0.05$, $^{***} = p < 0.01$.

Table A11: Inter-Rater Reliability Between Humans and Large Language Model

	(1) Full Sample	(2) High Confidence	(3) Not High Confidence
Panel A: Vacant Land			
Human-LLM Comparisons			
Pooled Coefficient	0.78	0.89	0.27
	(0.01)	(0.02)	(0.03)
Coefficient Range	[0.78, 0.89]	[0.87, 0.92]	[0.19, 0.71]
Num. Regressions	4	4	4
Num. Ratings	3,721	2,256	1,465
Human-Human Comparisons			
Pooled Coefficient	0.71	0.80	0.29
	(0.02)	(0.02)	(0.07)
Coefficient Range	[0.63, 0.78]	[0.72, 0.88]	[0.24, 0.48]
Num. Regressions	6	6	6
Num. Ratings	5,946	4,756	1,190
Panel B: Nonvacant Properties			
Human-LLM Comparisons			
Pooled Coefficient	0.85	0.84	0.92
	(0.08)	(0.09)	(0.05)
Coefficient Range	[0.68, 0.91]	[0.82, 0.91]	[0.17, 1.00]
Num. Regressions	5	5	5
Num. Ratings	2,372	2,074	298
Human-Human Comparisons			
Pooled Coefficient	0.86	0.88	0.70
	(0.02)	(0.02)	(0.09)
Coefficient Range	[0.80, 0.95]	[0.82, 0.96]	[0.16, 0.96]
Num. Regressions	8	8	8
Num. Ratings	1,316	1,150	166

Notes: This table reports inter-rater reliability coefficients from pairwise regressions. Both "Human-LLM" and "Human-Human" rows report mean coefficients from pooled regressions where all rater pairs are stacked together. Standard errors (in parentheses) are clustered by property description. Coefficient ranges come from individual regressions for each rater pair. For Human-LLM: 4 regressions (Human 1 on LLM, Human 2 on LLM, etc.). For Human-Human: 12 directed pairwise regressions (Human 1 on Human 2, Human 2 on Human 1, etc.). High confidence indicates both raters in the pair reported high confidence (3). Not High Confidence indicates at least one rater in the pair did not report high confidence.

Table A12: Instrumental-Variables Estimates of the Approval Premium, Vacant Lots

	(1)	(2)	(3)	(4)	(5)	(6)
	Cross Section			Repeat Sale		
	Keyword	LLM	Human	Keyword	LLM	Human
1. OLS	0.322***	0.308***	0.310***	0.409***	0.487***	0.598***
	(0.088)	(0.084)	(0.088)	(0.105)	(0.165)	(0.120)
2. First Stage	0.552***	0.682***		0.542***	0.605***	
	(0.030)	(0.023)		(0.098)	(0.157)	
3. IV	0.583***	0.451***		0.755***	0.805***	
	(0.160)	(0.124)		(0.183)	(0.159)	
Num. Listings	2,347	2,347	2,347	1,265	1,265	1,265
Num. Parcels	1,620	1,620	1,620	537	537	537

Notes: This table estimates approval premia using various specifications and measures of preapproval. All estimates come from the human-validation subsample of vacant lots. Columns 1–3 estimate the cross-sectional regression (Equation 11), whereas Columns 4–6 estimate the repeat-sales regression (Equation 12). Columns 1 and 4 use the keyword-based measure of approval, Columns 2 and 5 use the LLM-based measure, and Columns 3 and 6 use the hand-coded human measure. Ordinary least squares (OLS) and instrumental-variables (IV) coefficients are shown as estimated, without exponentiation. All first-stage regressions use the human rating as the outcome and instrument using the listed preapproval measure (keyword or LLM), and the coefficients are not exponentiated. Standard errors, reported in parentheses, are clustered at the parcel level. Weights are used in all regressions to account for sampling procedures for human validation. $^{*} = p < 0.10$, $^{**} = p < 0.05$, $^{***} = p < 0.01$.

Table A13: Instrumental-Variables Estimates of the Approval Premium, Nonvacant Lots

	(1)	(2)	(3)	(4)	(5)	(6)
	Cross Section			Repeat Sale		
	Keyword	LLM	Human	Keyword	LLM	Human
1. OLS	0.509***	0.329***	0.215**	0.281***	0.202***	0.202***
	(0.171)	(0.118)	(0.105)	(0.061)	(0.043)	(0.048)
2. First Stage	0.625***	0.809***		0.605***	0.800***	
	(0.087)	(0.054)		(0.120)	(0.091)	
3. IV	0.814**	0.406***		0.465***	0.253***	
	(0.327)	(0.153)		(0.122)	(0.068)	
Num. Listings	1852	1821	1852	471	463	471
Num. Parcels	1596	1569	1596	215	211	215

Notes: This table estimates approval premia using various specifications and measures of preapproval. All estimates come from the human-validation subsample of nonvacant lots. Columns 1–3 estimate the cross-sectional regression (Equation 11), whereas Columns 4–6 estimate the repeat-sales regression (Equation 12). Columns 1 and 4 use the keyword-based measure of approval, Columns 2 and 5 use the LLM-based measure, and Columns 3 and 6 use the hand-coded human measure. Ordinary least squares (OLS) and instrumental-variables (IV) coefficients are shown as estimated, without exponentiation. All first-stage regressions use the human rating as the outcome and instrument using the listed preapproval measure (keyword or LLM), and the coefficients are not exponentiated. Standard errors, reported in parentheses, are clustered at the parcel level. Weights are used in all regressions to account for sampling procedures for human validation. $^{*} = p < 0.10$, $^{**} = p < 0.05$, $^{***} = p < 0.01$.

Table A14: Sensitivity of Approval Premium to Property-Level Controls

	Baseline		Year–Month FE		Lot Size		Listing Length		All Controls	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
RTI (Narrow)	44.2***		44.5***		44.7***		42.3***		43.1***	
	(2.8)		(2.8)		(2.7)		(2.7)		(2.7)	
RTI (Complete)		34.0***		34.0***		33.8***		32.4***		32.2***
		(2.1)		(2.1)		(2.1)		(2.1)		(2.1)
RTI (In Process)		50.8***		51.0***		51.2***		48.7***		49.5***
		(2.9)		(2.9)		(2.9)		(2.8)		(2.8)

Notes: This table reports estimated approval premia from Equation 12 (repeat-sales difference-in-differences). All coefficients are exponentiated, showing $100[\exp(\beta) - 1]$. See the main text for definitions of the measure of permit status. Columns report estimates under varying sets of controls. Observations with missing values for control variables are included using binary indicators for missingness (coefficients not reported). Standard errors, reported in parentheses, are calculated using the delta method and reflect clustering by property. $^{*} = p < 0.10$, $^{**} = p < 0.05$, $^{***} = p < 0.01$.

Table A15: Heterogeneity in Approval Premium by Developer Specialization

	(1)	(2)	(3)	(4)
Fully Approved	0.449***	0.494***	0.425***	0.472***
	(0.043)	(0.043)	(0.049)	(0.049)
Fully Approved × Specialist Share	-0.119*	-0.080	-0.073	-0.038
	(0.071)	(0.070)	(0.079)	(0.077)
In Progress		0.256***		0.234***
		(0.026)		(0.029)
In Progress × Specialist Share		0.129*		0.156*
		(0.071)		(0.081)
Property F.E.	X	X	X	X
Year F.E.	X	X		
Zipcode–Year F.E.			X	X

Notes: This table reports coefficient estimates from a variant of Equation 12 in which permit status is interacted with our measure of developer specialization, the share of other properties associated with the same tax mailing address that were preapproved before sale. Coefficients are not exponentiated, and we present results using our narrow and joint measures of preapproval, along with various sets of fixed effects. Standard errors, reported in parentheses, are clustered at the parcel level. $^{*} = p < 0.10$, $^{**} = p < 0.05$, $^{***} = p < 0.01$.

Table A16: Approval Premium in Nonvacant Sample:
Controlling for Existing-Structure Characteristics

	(1) Baseline	(2) + Lot Area	(3) + Floorspace
Panel A: Cross Section			
Complete RTI	-0.6	1.3	1.6**
	(1.1)	(0.9)	(0.8)
RTI in Progress	-0.4	0.5	3.3**
	(2.0)	(1.8)	(1.6)
Panel B: Repeat Sales			
Complete RTI	10.4***	10.3***	8.0***
	(0.7)	(0.7)	(0.7)
RTI in Progress	4.9***	4.8***	5.4***
	(1.4)	(1.4)	(1.4)

Notes: This table reports estimated approval premia in the nonvacant sample with varied controls for the existing structure. All cells report exponentiated coefficients ($100 \cdot [\exp(\beta) - 1]$). Panel A reports estimates of Equation 11, whereas Panel B reports estimates of Equation 12. Column 1 includes only the fixed effects, Column 2 controls for (log) lot area, and Column 3 adds a control for (log) floorspace. Standard errors, reported in parentheses, are clustered at the parcel level. $^{*} = p < 0.10$, $^{**} = p < 0.05$, $^{***} = p < 0.01$.

Table A17: Approval Premium with Imputed Permit Controls

	(1)	(2)	(3)	(4)
Specification:	Cross Section		Repeat Sales	
Controls:	No	Yes	No	Yes
RTI	0.547***	0.545***	0.377***	0.377***
	(0.018)	(0.018)	(0.020)	(0.020)
Floorspace		0.010		0.000
		(0.008)		(0.008)
Units		0.003		-0.001
		(0.009)		(0.007)
Valuation		-0.009		0.001
		(0.008)		(0.007)
Height		-0.015		-0.002
		(0.013)		(0.008)
Lot Area		0.329***		0.090***
		(0.006)		(0.009)
Approval Premium	72.8***	72.5***	45.8***	45.8***
	(3.1)	(3.1)	(2.9)	(2.9)
Observations	63,845	63,845	45,543	45,543
Unique Lots	31,682	31,682	13,376	13,376

Notes: This table reports coefficient estimates for the hedonic value of preapproval (RTI) as well as other project characteristics as reported in the permit data. Columns 1 and 2 report estimates of Equation 11, whereas Columns 3 and 4 report estimates of Equation 12. Columns 2 and 4 use a multiple-imputation procedure (see Appendix C) to control for potential permit attributes when permits have not yet been filed. The approval-premium row reports exponentiated coefficients ($100[\exp(\beta) - 1]$) on preapproval. All regressions are estimated on the vacant-lot sample. Standard errors, reported in parentheses, are clustered at the parcel level. $^{*} = p < 0.10$, $^{**} = p < 0.05$, $^{***} = p < 0.01$.

Table A18: Conditional Differences in Permit Outcomes:
Preapproved Lots Versus Similar But Not Preapproved

	(1)	(2)	(3)	(4)
	Floorspace	Units	Valuation	Height
RTI	0.681***	0.594***	0.814***	0.124***
	(0.029)	(0.027)	(0.039)	(0.009)
Lot Area	0.050***	-0.103***	-0.026***	-0.068***
	(0.003)	(0.003)	(0.004)	(0.001)
Exponentiated Coefficient	97.6***	81.1***	125.8***	13.2***
	(5.8)	(4.9)	(8.9)	(1.0)
Observations	645,230	643,851	647,705	644,574
Unique Lots	31,687	31,686	31,688	31,687

Notes: This table reports conditional differences in four permitting outcomes between preapproved (RTI) and not-preapproved lots. The four outcomes (all in logarithms) are floorspace, unit count, reported valuation of construction cost, and building height. Exponentiated coefficients ($100[\exp(\beta)-1]$) are also reported. Standard errors are clustered at the parcel level. $^{*} = p < 0.10$, $^{**} = p < 0.05$, $^{***} = p < 0.01$.

Table A19: Approval Premium by Preapproved Structure Type, Vacant Lots

	(1) Narrow	(2) Broad	(3) Joint
Multifamily	44.4***	28.8***	
	(4.4)	(3.2)	
All Other	41.7***	30.2***	
	(3.0)	(1.8)	
Test, Equality	0.3 [0.58]	0.1 [0.73]	
Complete, Multifamily			53.2***
			(4.7)
Complete, All Other			46.6***
			(3.1)
Test, Equality			1.6 [0.20]
In Progress, Multifamily			43.2***
			(4.4)
In Progress, All Other			29.0***
			(2.2)
Test, Equality			9.0*** [0.00]

Notes: This table reports estimated approval premia for vacant lots, split by the preapproved structure type (multifamily residential or other) and the status of the preapproval (complete or in-progress). All estimates reflect exponentiated coefficients, estimated using our repeat-sales specification (Equation 12). "Narrow" in Column 1 refers to defining (binary) preapproval status as complete versus not-complete, whereas "Broad" in Column 2 defines preapproval status as complete or in progress versus neither complete nor in progress. Column 3 includes both complete and in-progress reapprovals jointly in the same regression. Standard errors (in parentheses) reflect clustering at the parcel level and are calculated via the delta method. F-statistics and p-values (the latter in square brackets) are also reported for tests of equality that contrast multifamily and all-other approval premia.

Table A20: Permitting Versus the Housing Cost Wedge

	(1)	(2)	(3)	(4)	(5)
	Permitting Cost		Housing Cost Wedge		Permitting Share
	Dollars Per Sqft. Land	Share of Constr. Cost	Dollars Per Sqft. Land	Share of Constr. Cost	Share of Cost Wedge
City Average	4.7	36.3	40.4	112.8	32.1
	(0.2)	(1.4)			(1.3)
Wedge Quintile					
Q1	0.1	1.7	2.4	11.9	14.6
	(0.1)	(1.4)			(12.1)
Q2	0.5	13.6	11.2	42.0	32.3
	(0.1)	(1.7)			(4.0)
Q3	4.5	29.2	34.8	101.2	28.9
	(0.4)	(2.5)			(2.4)
Q4	6.7	42.3	53.8	147.7	28.6
	(0.5)	(3.6)			(2.5)
Q5	9.5	84.6	93.2	243.9	34.7
	(0.5)	(4.8)			(2.0)

Notes: This table reports estimates of permitting costs, the aggregate housing cost wedge, and the permitting share of the aggregate wedge. All of these objects are reported as a citywide average and by quintile of the wedge, assigning all properties by zipcode according to its median wedge. Permitting costs as a share of construction cost are computed as $[\theta/(1+\theta)] \cdot \omega$, where θ is estimated from the land market and ω is the mean total housing cost wedge, either citywide or by quintile. The permitting share of the total housing cost wedge is $\theta/(1+\theta)$. Standard errors are computed using a nonparametric bootstrap, reflecting clustering by property, and do not account for estimation uncertainty in ω.

B Theory

B.1 Equilibrium Definition

Definition 1. *In the limit as the number of parcels $J \to \infty$, a steady-state equilibrium is a vector of endogenous objects $[\mathbf{g}, \ell, \mathbf{r}, \mathbf{q}^*, \mathbf{p}_1, \mathbf{p}_0]$, containing the following for each parcel:*

(i) a probability mass function over levels of housing capital $g_j(k; \mathbf{r}) = \Pr(h_j = k \mid \mathbf{r})$, forming $\mathbf{g}(\mathbf{r}) = [g_1(\mathbf{r}), \ldots, g_J(\mathbf{r})]$, with masses at $h_j = 0$ and a unique h_j^*, and otherwise $g_j(k; \mathbf{r}) = 0$;

(ii) a probability mass function over parcel shares, $\ell_j(s; \mathbf{r}) = \Pr(s_j = s \mid \mathbf{r})$, forming $\ell(\mathbf{r}) = [\ell_1(\mathbf{r}), \ldots, \ell_J(\mathbf{r})]$, also with masses only at $s_j = 0$ and $s_j = s_j^*$;

(iii) a unique rental price of housing capital r_j, forming $\mathbf{r} = [r_1, \ldots, r_J]$;

(iv) a unique permitting hazard q_j^*, forming $\mathbf{q}^* = [q_1^*, \ldots, q_J^*]$;

(v) a unique price $p_{1,j}(\bar{h}_j)$ of approved land, forming $\mathbf{p}_1 = [p_{1,j}(\bar{h}_j), \ldots, p_{1,J}(\bar{h}_j)]$;

(vi) a unique price $p_{0,j}(\bar{h}_j)$ of raw land, forming $\mathbf{p}_0 = [p_{0,j}(\bar{h}_j), \ldots, p_{0,J}(\bar{h}_j)]$;

such that, given primitives $\{\alpha, \beta, c_j, \delta, D\}$ and government policy $\bar{\mathbf{h}} = [\bar{h}_1, \ldots, \bar{h}_J]$, the following conditions hold:

(i) Household maximization: Parcel choices $\mathbf{s}(\mathbf{r})$ satisfy Equation (1);

(ii) Developer maximization: Capital investment choices h_j^* satisfy Equation (2) and permitting effort q_j^* satisfies Equation (5);

(iii) Housing market clearing: The distribution of household choice shares equals the distribution of capital shares: $\mathbf{s}(\mathbf{r}) = \mathbf{g}(\mathbf{r})/G(\mathbf{r})$, where the aggregate stock is $G(\mathbf{r}) = \sum_j \mathrm{E}[g_j(\mathbf{r})]$;

(iv) Land market clearing: In the long run, inflow (demolition) equals outflow (permitting): $\delta \Pr(h_j = h_j^* \mid \mathbf{r}) = q_j^* \Pr(h_j = 0 \mid \mathbf{r})$ for all j.

B.2 Proofs

Lemma 1. *Equilibrium prices of raw and approved land are given respectively by $p_{0,j}(\bar{h}_j) = \frac{\beta q_j^*}{1-\beta(1-q_j^*-c(q_j^*))} \int_0^{\bar{h}_j} \lambda_j(s)\, ds$ and $p_{1,j}(\bar{h}_j) = \int_0^{\bar{h}_j} \lambda_j(s)\, ds$.*

Proof. We first establish the equilibrium price $p_{1,j}(\bar{h}_j)$ of approved land. After approval, the developer maximizes $\pi_j(\bar{h}_j) = \max_{h \le \bar{h}_j} \{R_j(h) - k_j(h)\}$. By the Envelope Theorem, the marginal value of relaxing the quantity constraint is the shadow price $\lambda_j(\bar{h}_j)$. The value function can therefore be expressed as the integral of the shadow price:

$$\pi_j(\bar{h}_j) = \pi_j(0) + \int_0^{\bar{h}_j} \lambda_j(s)\, \mathrm{d}s.$$

Since the profit from a parcel with a quantity limit of zero is zero ($\pi_j(0) = 0$), and developers compete away profits in the approved land market such that price equals value ($p_{1,j} = \pi_j$), we obtain $p_{1,j}(\bar{h}_j) = \int_0^{\bar{h}_j} \lambda_j(s)\,\mathrm{d}s$. Integrability of $\lambda_j(\cdot)$ follows from the regularity of the revenue and cost functions R_j and k_j.

Next, we determine the price of raw land. The fundamental value of raw land, $v_{0,j}$, is determined by the solution to the Bellman equation in Equation (4):

$$v_{0,j}(\bar{h}_j) = \frac{\beta q_j^*}{1 - \beta(1 - q_j^* - c_j(q_j^*))}\,\pi_j(\bar{h}_j).$$

Imposing the zero-profit condition for raw land ($p_{0,j} = v_{0,j}$) and substituting the expression for $\pi_j(\bar{h}_j)$ derived above yields the expression for $p_{0,j}(\bar{h}_j)$. □

Proposition 1. *Relative to its counterfactual raw price, approved land trades at a premium of*

$$\theta_j = \frac{p_{1,j}(\bar{h}_j) - p_{0,j}(\bar{h}_j)}{p_{0,j}(\bar{h}_j)} = \frac{1-\beta}{\beta q_j^*} + \frac{c_j(q_j^*)}{q_j^*},$$

and we express the dollar-value premium as $\tau_j(\bar{h}_j) = p_{1,j}(\bar{h}_j) - p_{0,j}(\bar{h}_j)$.

Proof. By definition, the proportional approval premium is $\theta_j = \frac{p_{1,j}(\bar{h}_j)}{p_{0,j}(\bar{h}_j)} - 1$. Using the expressions for equilibrium prices derived in Lemma 1, the ratio of approved to raw land prices is:

$$\frac{p_{1,j}(\bar{h}_j)}{p_{0,j}(\bar{h}_j)} = \frac{\int_0^{\bar{h}_j} \lambda_j(s)\,\mathrm{d}s}{\left(\frac{\beta q_j^*}{1-\beta(1-q_j^*-c_j(q_j^*))}\right)\int_0^{\bar{h}_j} \lambda_j(s)\,\mathrm{d}s} = \frac{1 - \beta(1 - q_j^* - c_j(q_j^*))}{\beta q_j^*}.$$

Expanding the numerator yields:

$$\frac{p_{1,j}(\bar{h}_j)}{p_{0,j}(\bar{h}_j)} = \frac{1 - \beta + \beta q_j^* + \beta c_j(q_j^*)}{\beta q_j^*} = \frac{1-\beta}{\beta q_j^*} + 1 + \frac{c_j(q_j^*)}{q_j^*}.$$

Subtracting 1 from this ratio yields the expression for θ_j. □

Proposition 2. *Reducing time-to-build through preapproval increases the present value of rental income by the following proportion:*

$$\frac{\sum_{t=0}^{\infty} \beta^t \left(F_{1,j}(t) - F_{0,j}(t)\right)}{\sum_{t=0}^{\infty} \beta^t F_{0,j}(t)} = \frac{1-\beta}{\beta q_j^*}.$$

Proof. The present value of the rental stream for a parcel j is given by $V_j = \sum_{t=0}^{\infty} \beta^t \mathrm{E}[I_j(t)] r_j$, where $I_j(t)$ is an indicator variable equal to 1 if the building is completed by period t and 0 otherwise. Since $\mathrm{E}[I_j(t)] = F_j(t)$, we can write the value as $V_j = r_j \sum_{t=0}^{\infty} \beta^t F_j(t)$.

For preapproved land ($j = 1$), $F_{1,j}(t) = 1$ for all $t \geq 0$. The value is $V_{1,j} = r_j/(1-\beta)$. For raw land ($j = 0$), the cumulative probability of permit arrival is $F_{0,j}(t) = 1 - (1 - q_j^*)^t$. The sum is:

$$\sum_{t=0}^{\infty} \beta^t [1 - (1 - q_j^*)^t] = \frac{1}{1-\beta} - \frac{1}{1-\beta(1-q_j^*)} = \frac{\beta q_j^*}{(1-\beta)(1-\beta(1-q_j^*))}.$$

Multiplying by r_j gives $V_{0,j}$. The LHS of the proposition is equivalent to $(V_{1,j} - V_{0,j})/V_{0,j}$. Substituting the expressions derived above:

$$\frac{V_{1,j}}{V_{0,j}} - 1 = \frac{\frac{1}{1-\beta}}{\frac{\beta q_j^*}{(1-\beta)(1-\beta(1-q_j^*))}} - 1 = \frac{1-\beta(1-q_j^*)}{\beta q_j^*} - 1 = \frac{1-\beta}{\beta q_j^*}.$$

□

Proposition 3. *The housing cost wedge decomposes into permitting costs and raw land value:*

$$R_j(\bar{h}_j) - k_j(\bar{h}_j) = \tau_j(\bar{h}_j) + p_{0,j}(\bar{h}_j) = p_{1,j}(\bar{h}_j),$$

such that the permitting share of the wedge has two equivalent formulations:

$$\frac{\tau_j(\bar{h}_j)}{R_j(\bar{h}_j) - k_j(\bar{h}_j)} = \frac{\theta_j(\bar{h}_j)}{1 + \theta_j(\bar{h}_j)}.$$

Proof. From the developer's optimization problem for approved land in Equation 2, the total profit realized at the quantity limit $\bar{h}_j$ is $\pi_j(\bar{h}_j) = R_j(\bar{h}_j) - k_j(\bar{h}_j)$. In a perfectly competitive land market, this profit is capitalized into the value of a permitted parcel. By Lemma 1, the equilibrium price of approved land is the integral of the shadow prices of the land-use constraint, $p_{1,j}(\bar{h}_j) = \int_0^{\bar{h}_j} \lambda_j(s)ds$, which by the fundamental theorem of calculus equals the total profit $\pi_j(\bar{h}_j)$. This establishes the first identity: $R_j(\bar{h}_j) - k_j(\bar{h}_j) = p_{1,j}(\bar{h}_j)$.

From the definition of the dollar-value approval premium in Proposition 1, we have $\tau_j(\bar{h}_j) = p_{1,j}(\bar{h}_j) - p_{0,j}(\bar{h}_j)$. Rearranging this gives $p_{1,j}(\bar{h}_j) = \tau_j(\bar{h}_j) + p_{0,j}(\bar{h}_j)$, completing the three-way decomposition.

To derive the second result, let $W_j = R_j(\bar{h}_j) - k_j(\bar{h}_j)$ denote the housing cost wedge. Since $W_j = p_{1,j}$, the permitting share of the wedge is $\tau_j/p_{1,j}$. By the definition of the relative approval premium, $p_{1,j} = p_{0,j}(1 + \theta_j)$. Substituting $\tau_j = p_{1,j} - p_{0,j}$ into the share formula yields:

$$\frac{\tau_j}{p_{1,j}} = \frac{p_{1,j} - p_{0,j}}{p_{1,j}} = 1 - \frac{p_{0,j}}{p_{0,j}(1+\theta_j)} = 1 - \frac{1}{1+\theta_j} = \frac{\theta_j}{1+\theta_j}.$$

This demonstrates that the permitting share of the regulatory wedge can be recovered entirely from the land-market premium θ_j. □

C Other Materials

C.1 Supplementary Data Sources

Property Tax Assessments. To identify the role of specialized real estate intermediaries, we utilize annual property tax assessment rolls for Los Angeles County from 1995 to 2024 obtained via Cotality. Following Ganduri et al. (2023), we identify unique property owners by their tax mailing addresses rather than legal entity names. This approach allows us to "see through" the common practice of holding individual properties in distinct Limited Liability Companies (LLCs) that share a common principal. We define a unique owner ID by the combination of the mailing street address, city, and state. We link each property listing in our sample to its tax mailing address recorded in the assessment roll for the three years prior to the listing date ($t-1$, $t-2$, and $t-3$). For each unique owner, we calculate the share of their other listings in the county that were sold with preapproved (RTI) permits to characterize their degree of specialization in permitting. We compute the final shares as the maxima across the three horizon years.

Zoning Classification. Data on land-use zoning for Los Angeles County are from the Southern California Association of Governments (SCAG) 2019 land-use dataset.[28] The underlying data follow a Level IV Modified Anderson Land Use Classification system (Anderson, 1976), which provides detailed classess of built-up areas and natural land cover. For empirical analysis, we aggregate these classes into eight mutually exclusive and completely exhaustive categories (see Appendix Table A21).

American Community Survey. We use tract-level data from the 2019–2023 American Community Survey (ACS) 5-year estimates. To control for time-invariant neighborhood characteristics, we use a vector of ten covariates: population density, the share of renters, the poverty rate, racial and ethnic composition, and educational attainment shares. The tract data are linked to the listings using a spatial join on the geographic coordinates.

Permitting Fees. We estimate fiscal predevelopment costs by imputing building permit fees according to the official documents.[29] We apply a 1.9x multiplier to base fees to account for the plan check fee. We express these fees as a share of counterfactual raw-land value, which we recover by deflating preapproved land prices by our estimated 50-percent approval premium.

C.2 Time-to-Build Across Cities

This appendix provides additional details on the data standardization process used in our cross-city permit analysis.

We use publicly-available building permit data from 11 cities and the City of Los Angeles. The data were generally retrieved from municipal open-data websites. For Los Angeles, we augment data from the City of Los Angeles with data from the County's other constituent cities, which we purchased from a third-party data vendor (Shovels).

[28] These data are publicly available at https://hub.scag.ca.gov/datasets/ea9fda878c1947d2afac5142fd5cb658_0/.

[29] https://planning.lacity.gov/odocument/62df4599-b5b3-4ce9-880f-ec4de6b1a1c2/Appendix_2.4_-_Summary_of_Case_Filing_and_Building_Permitting_Fees_(Adopted).pdf.

Table A21: Aggregation of Modified Anderson Land Use Codes

Aggregated Category	Zoning Class	Included Descriptions
1. Single-Fam. Res.	1100, 1110, 1112, 1113	Single-family detached units ranging from low to medium density (estates to modern subdivisions).
2. Multifam. Res.	1120–1125, 1140	Attached units, duplexes, apartments (low to high-rise), and older neighborhoods with mixed residential types.
3. Commercial	1200–1247	General offices, retail stores, shopping centers, medical facilities, and public service buildings (e.g., libraries).
4. Industrial	1300–1340	Light and heavy manufacturing, film/TV studios, wholesaling, warehousing, and mineral extraction.
5. Open Space	1800–1890, 1900, 3400, 8888	Parks, golf courses, urban vacant lots, beaches, and land protected for environmental or slope constraints.
6. Specific Plan	7777	[Refers to Specific Plan districts in Los Angeles]
7. Rural/Agriculture	1150, 2000	Farmsteads, ranches, and land used for the production of food, fiber, or livestock.
8. Mixed-Use	1500, 1600, 1610, 1620	Sites combining commercial and industrial uses, or residential-oriented and commercial-oriented mixed-use buildings.

Notes: This table describes our mapping of the 2019 Southern California Association of Government (SCAG) zoning codes into the aggregated categories used for analysis in Appendix Figure A14. SCAG follows a Modified Anderson Land Use Classification available at `https://scag-spm-documentation.readthedocs.io/en/latest/scag_lu_codes_description/`.

In cleaning the data, we made the following sample restrictions. First, we retain only permits for new multifamily buildings that have valid dates of filing, issuance, and completion (i.e., are thus completed). Second, we limit to permit applications submitted between 2000 and 2025 that proposed fewer than 200 units and fewer than 10 years from submission to completion. Third, we exclude expired or voided permits from the sample whenever possible.

The key variables for our cross-city analysis are:

- *Units:* We filter the sample to include only permits for new buildings with at least two units. Some cities' data include multiple variables that could plausibly capture the number of units (e.g., `units` and `proposedunits`). In these cases, we define units as the maximum value across the candidate variables.
- *Neighborhood Population Density:* We spatially merged permits to 2019 Census tracts using the parcel latitude and longitude coordinates included in permit data. When coordinates were not available, we geocoded permits from their addresses and then merged them to the 2019 tracts. We then attach tract-level population density collected in the 2019 American Community Survey.

- *Time-to-Build:* We define time-to-build as the number of calendar days between the dates of permit submission and completion (issuance of the permanent certificate of occupancy). We always present this variable in years.

Appendix Table A22 reviews coverage and sample characteristics by city with the sample restrictions. Motivating our standardization efforts, it is challenging to see what could learn from unadjusted summary statistics about time-to-build in cross-city comparisons. Across cities, the median project varies by an order of magnitude in both its unit count and the density of the surrounding neighborhood.

Table A22: Permit Coverage and Sample Characteristics by City

(1) City	(2) Years	(3) Number of Permits	(4) Unit Count (Med.)	(5) Tract Density (Med.)
Austin, TX	2000–2025	7,606	6	4,022
Cincinnati, OH	2010–2024	248	8	5,766
Denver, CO	2010–2024	727	24	4,063
Fort Worth, TX	2001–2025	4,541	4	2,998
LA County, CA	2000–2024	9,634	2	17,943
New York, NY	2000–2021	37,429	2	35,674
Portland, OR	2000–2020	1,551	8	5,053
Raleigh, NC	2000–2019	1,154	22	2,750
San Diego, CA	2003–2022	581	21	8,767
San Francisco, CA	2000–2021	1,249	4	25,848
Seattle, WA	2000–2024	6,669	3	10,976
Tampa, FL	2007–2025	263	6	4,309

We use a regression to organize the data by predicting time-to-build for a standardized project. For each city c, we estimate a Poisson regression of time-to-build (TTB_{ic}) on unit count (Units_{ic}), local population density (PopDen_{ic}), and year fixed effects. Our baseline specification is

$$\mathrm{E}\left[\text{TTB}_{ic} \mid X_{ic}\right] = \exp\left(\beta_{0c} + \beta_{1c}\log(\text{Units}_{ic}) + \beta_{2c}\log(\text{PopDen}_{ic}) + \sum_t \gamma_{tc}\mathbf{1}\{\text{Year}_{ic} = t\}\right), \quad (17)$$

with standard errors clustered at the permit level.

We then use the model to construct, for each city c, the predicted time to build for a reference project by evaluating the fitted mean at $\log(\text{Units}_{ic}) = \log(30)$, $\text{Year}_{ic} = 2019$, and $\log(\text{PopDen}_{ic}) = \log(6000)$, where 6,000 is the mean tract-level population density for projects with total units within 10 units of 30, applied for between 2000 and 2025 and with time to build less than 10 years.

C.3 Estimation Details for Time-to-Build Map (Panel B, Figure 1)

We model the time from permit application filing to certificate of occupancy, denoted T_{iz} for project i in zipcode z, using the following Poisson Pseudo-Maximum Likelihood (PPML) regression:

$$E\left[T_{iz} \mid \text{Units}_i, \alpha_z\right] = \exp\left[\alpha_z + \beta_1\log(\text{Units}_i) + \beta_2 1(\text{Missing Units})_i\right],$$

where α_z is a zipcode fixed effect, Units_i is the proposed dwelling unit count of the project and In this model, the key explanatory variables are the natural log of the project's square footage and $1(\text{Missing Units})_i$ is an indicator for whether unit count are missing.

From the estimated model, we generate a baseline prediction for a standardized 30-unit apartment building in each zipcode. This predicted time from permit-to-completion, denoted $\widehat{T}_z$, is calculated as:

$$\widehat{T}_z = \exp\left[\hat{\alpha}_z + \hat{\beta}_1 \log(30)\right].$$

Finally, we adjust the estimates in Panel B to account for the time to file a permit after land acquisition. As shown in Panel A, data from vacant land sales reveals a stable, linear relationship between the time-to-file and the subsequent time-to-complete a project. We use this empirical relationship to proportionally scale up our baseline prediction, $\widehat{T}_z$. This final adjusted prediction is what is mapped in Panel B of Figure 2. To ensure reliability, we only map predictions for zipcodes with at least five completed projects in our sample.

Appendix Figure A15 performs an analogous exercise for the probability with which a 30-unit apartment building is ever completed, given permit submission or issuance. We find little spatial variation in the completion rate.

C.4 Multiple Imputation of Permitted Use

This section describes the multiple-imputation procedure (Multivariate Imputation by Chained Equations, MICE) used to evaluate whether the approval premium is driven by differences in the intensity of permitted land use. Section 5 motivates this strategy and discusses necessary assumptions for it to yield valid estimates of the approval premium.

The imputation model estimates values for the entire set of project-specific characteristics (X_j), specifically building height, construction valuation, unit count, and total floorspace. Given the right-skewness of these variables, we perform the imputation on their logarithms. The imputation is conditioned on the preapproval indicator (RTI_j) and a full set of zipcode and year fixed effects. For the k-th imputation variable,

$$X_{j,k} = \rho_k \text{RTI}_j + \eta_{k,z(j)} + \delta_{k,y(j)} + \nu_{j,k},$$

where $z(j)$ is the zipcode and $y(j)$ is the listing year for property j.

Estimation and Pooling. We implement the procedure using the following steps:

1. *Imputation:* We generate $M = 10$ imputed datasets using a Gibbs sampler (chained equations). Each missing variable is predicted iteratively using the other variables in the system until the distribution converges.

2. *Estimation and Sample:* We estimate our primary cross-sectional and repeat-listing models separately on each of the M datasets. We do this ensuring that the inclusion of controls does not change the underlying sample. That is, Columns 1 and 3 of Appendix Table A17 reflect the exact same sample of observations for which the MICE procedure successfully converges.

3. *Inference:* We pool the coefficients and standard errors across all M imputations following Rubin (1987), so as to ensure that the standard errors correctly reflect both standard sampling uncertainty and the additional uncertainty introduced by imputation.

C.5 Imputation Procedure for Construction Cost

To estimate the gap between sale prices and the physical cost of construction, we impute construction cost for each property using detailed characteristics. Our procedure involves (1) defining our sample from CoreLogic MLS data, (2) mapping MLS features to a set of standardized inputs, (3) imputing construction quality from MLS text fields, and (4) applying a component-based cost function calibrated from RS Means 2024 cost data.

Data and Sample Definition. We restrict our sample to arm's-length transactions by keeping only records where the MLS indicates it is a standard sale, which also excludes REO, short sales, and foreclosures. We further restrict the sample to single-family structures, due to greater challenges of inferring building characteristics of multifamily structures in the MLS.

The calibration uses values from the 2024 edition of R.S. Means "Square Foot Costs" book, adjusted using their national historical cost index as appropriate. The calibration has five components, following the setup of the R.S. Means books:

1. *Main Structure Costs*. Provides the base cost per square foot for the main structure and basements, disaggregated by quality, number of stories, wall type, and floorspace in square feet. This cost is linearly interpolated according to floorspace to account for economies of scale in construction.

2. *Variable Cost Modifiers*. Provides cost adjustments to base values, on a per-square-foot basis, for non-standard roof and heating systems, disaggregated by quality and number of stories. This is also interpolated according to floorspace.

3. *Fixed Costs*. Provides fixed, per-unit costs for amenities (e.g., full bath, fireplace, pool), disaggregated by quality.

4. *Linear-Foot Costs*. Provides per-linear-foot costs for fencing materials.

5. *Historical Cost Index*. We adjust construction costs to the year of a given transaction.

Cost Model. We define the total construction cost C_{it} for a house i sold in year t as its 2024-vintage cost, $C_{i,2024}$, adjusted by the national historical cost index (P_t) and the 2024 location factor for Los Angeles (F) :

$$C_{it} = F \cdot \frac{P_t}{P_{2024}} \cdot C_{i,2024}.$$

The 2024 cost $C_{i,2024}$ is a component-based function that sums four main parts, as defined above:

$$C_{i,2024} = (\text{Main Cost}_i + \text{Variable-Cost Modifiers}_i) \times \text{Sqft}_i + \text{Fixed Cost}_i + \text{Linear-Foot Cost}_i \times P_i,$$

where Sqft_i is floorspace P_i is the lot perimeter.

Construction of Input Variables. We reformat the raw MLS data into the 19 input variables required by our imputation model. The key processing steps are as follows:

- *Quality*: We classify each property into one of the three R.S. Means tiers (luxury, average, economy) using a combination of keywords in the MLS remarks and selected MLS variables:

1. *Luxury.* A property is classified as "luxury" if it meets one of two conditions: (1) it contains one or more "tier 1" keywords (e.g., high-end brand names like Sub-Zero, Viking, Pella; specific materials and finishes like quartzite or hand-hewn); or (2) it contains two or more "tier 2" keywords (e.g., "luxury", "custom", "gourmet", "marble").
2. *Economy.* A property not flagged as "luxury" is subsequently flagged as "economy" if the MLS remarks contain one or more "economy" keywords (e.g., "starter," "wall heater") or has MLS data fields indicating a lower-quality build (e.g., records indicate the structure lacks a cooling system, or the floor-type variable contains "laminate," "vinyl" or "carpet").
3. *Average.* Properties not classified as "luxury" or "economy" default to "average."

- *Wall*: We map MLS wall construction materials into four categories using keywords: "brick," "masonry," "wood," and "stucco." We apply a cost-based hierarchy, such that "brick" keywords overwrite "masonry," which in turn overwrites "wood" or "stucco." Given the context of Los Angeles, "stucco" serves as the default for ambiguous values of the wall-construction-material field.

- *Roof Type*: We map the MLS roof-type field into five categories, again using keywords: "slate," "cedar shake," "clay tile," "composition roll," and "asphalt shingles." The "asphalt shingles" category includes common synonyms like "composition" and "shingle." We do not make any R.S. Means cost adjustments for roof type when this field is ambiguous (e.g., "Other", "See Remarks").

- *Stories*: We cap the number of stories at 3 to match the bounds of the R.S. Means cost data.

- *Baths*: We parse the number of full and half baths from the MLS fields, treating fractional values as half baths and the lowest whole number as the number of full baths. Counts of full baths and half baths are separately capped at a maximum of ten.

- *Fireplaces*: We impute the count by taking the maximum value from several MLS fields that report the presence of any fireplace, the number of fireplace, and the types of fireplaces. We transform the last of these into a count from the number of comma and semicolon delimiters.

- *Lot Dimensions*: For properties missing lot width and length but not area in square feet, we impute dimensions. We first calculate the "aspect ratio" of the median lot (r = length/width). We then impute width as $\sqrt{\text{area}/r}$ and length as area/width. For all properties with reported or imputed length and width, we can impute the perimeter as $2 \times (\text{length} + \text{width})$.

- *Basement*: We classify properties as have no basement, a unfinished basement, or an finished basement, by hierarchically combining values from several MLS basement variables. The hierarchy is to prioritize any field indicating higher-cost option (finished, then unfininished, then none). We also use keywords "finished basement" and "basement" in the MLS remarks.

- *Other Variables*: The building's heat type, porch type, and garage type are mapped directly from their corresponding MLS variables, as the existence of a pool and a fence.

C.6 Measuring Permit Status

This appendix section explains how we measure permit status using keywords and the large language model.

Keyword. We use the following keywords: "RTI", "ready to issue", "ready-to-issue", "entitled", "entitlement", "permits pulled", "fully permit", "fully approve", "ready to build", "approved plans". The keyword "RTI" is case sensitive, and we also require that there is a space before this keyword only to eliminate issues with "RTI" appearing within a capitalized word. For all other keywords, we use a substring search without requiring a boundary space.

Large Language Model. We prompted the LLM as follows (OpenAI, GPT-4.1, 2025-04-14 release, in batch mode). First, the following prompt classified the vacant listings.

```
You are an expert real estate permitting analyst. Your task is to analyze a
    property description and output a JSON object.

**Instructions on Weighing Conflicting Signals:**
- **Definitive Status Overrides General Language**: A definitive statement about
    permit status (e.g., "in plan check," "RTI," "Approved plans") is the
    strongest signal and should be prioritized over general, boilerplate language.
- **Boilerplate Disclaimers**: Standard disclaimers like "buyer to verify" or "
    sold as is" should only force a Category 1 if there is NO definitive
    information about permit progress.
- **Hard Negatives**: Explicit negative facts like "expired plans" always force a
     Category 1 classification.

**Categories:**
- **1: Not Ready / Unclear**: The description lacks any definitive permit status,
     contains hard negatives ("expired plans"), or relies solely on describing
    potential.
- **2: Ready to Issue**: The description explicitly states all necessary plans
    and permits are fully approved, paid for, and issued (e.g., "RTI," "Approved
    plans").
- **3: In Progress**: The description shows clear evidence of active engagement
    with a planning department, such as submitted plans or applications that are
    awaiting corrections or approval (e.g., "in plan check," "awaiting approval").

**Instructions:**
Provide your response as a JSON object with the following keys:
- **reasoning**: A concise, step-by-step reasoning (1-2 sentences).
- **confidence**: A confidence rating (Low, Medium, or High).
- **category**: A single integer for the category (1, 2, or 3).
Your response must be a valid JSON object and nothing else.
```

Second, the following prompt is used for the nonvacant listings:

```
Your task is to determine if a real estate listing includes approved plans for a
    **future, unbuilt construction project.** Analyze with extreme precision and
    respond ONLY with a valid JSON object.
```

```
## Core Logic & Rules
1. **Default to Category 1.** If there is any ambiguity, or if the project
   described is already built or is only a hypothetical potential, the category
   is **1**.
2. **Distinguish Past vs. Future.** You must ignore all mentions of permits for
   **existing structures** or past renovations (e.g., "permitted bonus room," "
   remodeled with permits"). These are **not** future projects and belong in
   Category 1.
3. **Reject 'Potential'.** Any language describing what "can be," "could be," or
   what zoning "allows" is a sign of potential, not reality. This must be
   classified as Category 1.
4. **Require Explicit Proof.** Only classify as Category 2 or 3 if the listing
   provides definitive proof of an active or approved application for an **
   unbuilt project**, using the keywords defined below.

## RTI Status Categories
-  **1: Default / Not RTI:** The default classification. Use for:
   -  Potential, possibilities, or opportunities (e.g., "great for developers").
   -  Existing or completed work, even if permitted.
   -  Descriptions of zoning allowances (e.g., "R3 zoned lot").
-  **2: Approved / RTI:**
   -  **Condition:** Plans for a **future, unbuilt project** are fully approved
      and ready for construction.
   -  **Keywords:** "RTI," "ready to build," "permits in hand," "fully entitled
      ," "approved plans."
-  **3: In Progress:**
   -  **Condition:** A formal application for a **future, unbuilt project** has
      been submitted but is not yet approved.
   -  **Keywords:** "in plan check," "submitted to city," "awaiting approval."

## Value Proposition Scale
-  **1: Primarily Current Use Value:** Pitch is based on the property as it is
   today.
-  **2: Hybrid Value:** Pitch includes both current use and a future project as
   valuable paths.
-  **3: Primarily Development Value:** Pitch is based on the value of the land
   and/or entitlements for a future project.

## JSON Output Schema
-  **value_proposition_level**: (Int 1-3) Classification from the Value
   Proposition Scale.
-  **category**: (Int 1-3) Classification from the RTI Status Categories.
-  **confidence**: (String "Low", "Medium", "High") Your confidence in the
   assigned `category`.
-  **rti_content_type**: (Array of Strings | null) List of project components.
   Must be `null` if category is 1.
```

```
-   **rti_content_summary**: (String | null) Objective summary of the **future
    project**. Must be ‘null‘ if category is 1.

The prompt’s reference to permits being ‘‘paid for and issued’’ is intended to
    capture the pre-issuance RTI state specific to Los Angeles, in which all plan
    approvals are secured and plan check fees paid, such that the buyer may pull
    permits and begin construction immediately upon closing without further
    regulatory review. This differs from the final issuance step, at which point
    development impact fees become due. Listings that indicate remaining fee
    payment or issuance steps beyond plan check are intended to be classified as
    Category 3 (in progress) rather than Category 2. The agreement rate between
    LLM and human classifications suggests the classification was applied
    consistently (see Appendix Figure \ref{tab:confusion_matrix}).
```

The vacant-sample prompt’s reference to permits being “paid for and issued” is intended to capture the pre-issuance RTI state specific to Los Angeles, in which all plan approvals are secured and plan check fees paid, such that the buyer may pull permits and begin construction immediately upon closing without further regulatory review. This differs from the final issuance step, at which point development impact fees become due. Listings that indicate remaining fee payment or issuance steps beyond plan check are intended to be classified as Category 3 (in progress) rather than Category 2. The agreement rate between LLM and human classifications suggests the classification was applied consistently (see Appendix Table A6).

C.7 Geographic Heterogeneity in Approval Premium

We characterize spatial heterogeneity in the approval premium using an adaptive geographically weighted regression. For each zipcode z, we center a localized regression at the median latitude and longitude of sales within that zipcode. We use an adaptive bandwidth selection procedure, including in the regression sample for each zipcode the nearest fifth of sales in the county. We assign weights to these observations using a bi-square kernel:

$$w_{iz} = \begin{cases} \left(1 - \frac{\text{dist}_{iz}^2}{B_z^2}\right)^2 & \text{if dist}_{iz}^2 < B_z^2 \\ 0 & \text{otherwise} \end{cases}$$

where dist_{iz}^2 is the squared Euclidean distance between sale i and the sales-based median coordinate of zipcode z, and B_z is the maximum distance among observations in the nearest fifth.

Using these weights, we estimate a series of localized regressions to recover spatially-varying approval premia. For each zipcode z, we estimate:

$$\log p_{it} = \alpha_i + \alpha_t + \beta_z \text{RTI}_{it} + \sum_k \gamma_{zk} (X_{itk} \cdot \text{RTI}_{it}) + u_{it} \tag{18}$$

where X_{itk} is a vector of ten Census tract-level covariates (including population density, poverty rates, and land prices) that have been weighted-demeaned for each local sample. By demeaning the covariates locally, the coefficient $\hat{\beta}_z$ identifies the average approval premium for zipcode z.

We then use the delta method to transform these local coefficients into the three outcomes mapped in Figure 9 and Appendix Figure A18:

1. *Contribution to Wedge:* $100 \cdot \frac{\exp(\hat{\beta}_z)-1}{\exp(\hat{\beta}_z)} \cdot \bar{\omega}_z$, where $\bar{\omega}_z$ is the local mean housing cost wedge measured as a markup over construction cost (i.e., $(R-k)/k$).
2. *Dollar Cost:* $\tilde{p}_z \cdot [\exp(\hat{\beta}_z) - 1]$, where $\tilde{p}_z$ is the local weighted median price per square foot for non-preapproved land.
3. *Percent Premium:* $100 \cdot [\exp(\hat{\beta}_z) - 1]$.

To ensure reliability, we only report estimates for zipcodes when the standard error of the premium falls within a predefined reasonable range for each outcome. Finally, after estimation for map display only, we trim all negative point estimates to zero.

References for Appendices

Anderson, James Richard, "A Land Use and Land Cover Classification System for Use with Remote Sensor Data," Technical Report 964, US Government Printing Office 1976.

Conley, Timothy G, "GMM Estimation with Cross Sectional Dependence," *Journal of Econometrics*, 1999, *92* (1), 1–45.

Ganduri, Rohan, Steven Chong Xiao, and Serena Wenjing Xiao, "Tracing the Source of Liquidity for Distressed Housing Markets," *Real Estate Economics*, 2023, *51* (2), 408–440.

Müller, Ulrich K and Mark W Watson, "Spatial Correlation Robust Inference in Linear Regression and Panel Models," *Journal of Business & Economic Statistics*, 2023, *41* (4), 1050–1064.

Rubin, DB, *Multiple Imputation for Nonresponse in Surveys*, J Wiley and Sons, 1987.