Evicted from Opportunity: Labor Market Effects of Housing Loss in High-Productivity Cities*

Brian Asquith† Kate Pennington‡ Charly Porcher§

March 27, 2026

Abstract

How important is stable housing near high-paying jobs for long-run career growth? We study San Francisco Ellis Act withdrawals, which remove rent-controlled buildings from the rental market, and compare displaced tenants with tenants in nearby rent-controlled buildings. Six years after eviction, displaced workers earn about $13,000 less (20 percent of baseline earnings) and live in lower-value housing and neighborhoods. Losses are largest for younger workers and remain large even for movers of 5–25 km, who transition to smaller, lower-paying firms and face longer commutes. To account for these patterns, we develop an equilibrium model of frictional housing and job search and matching, costly commuting, localized search, and learning, calibrated to the Bay Area. Simulated central-city evictions qualitatively reproduce earnings losses, especially for young workers. Matching the empirical magnitudes likely requires the city center to offer an earnings-growth premium through stronger matching, faster learning, and faster transitions to more productive firms.

JEL Codes: J31, R23, R31, J61

Keywords: Eviction, displacement, urban labor markets, housing stability, commuting, agglomeration

*The views expressed in this paper are those of the authors and do not represent those of the U.S. Census Bureau. The Census Bureau has ensured appropriate access and use of confidential data and has reviewed these results for disclosure avoidance protection (Project 7529092: CBDRB-FY26-CES010-003). We thank seminar participants at Georgetown University for helpful comments.

†W.E. Upjohn Institute for Employment Research. Contact: asquith@upjohn.org.

‡Center for Economic Studies, U.S. Census Bureau. Contact: kate.m.pennington@census.gov.

§McDonough School of Business, Georgetown University. Contact: charly.porcher@georgetown.edu.

1 Introduction

Economic opportunity is increasingly concentrated in dense, high-productivity labor markets. A growing body of evidence documents that workers accumulate human capital faster in big cities (Glaeser and Maré, 2001; Baum-Snow and Pavan, 2012; De La Roca and Puga, 2017), that urban density generates productivity spillovers (Ciccone and Hall, 1996; Moretti, 2004; Duranton and Puga, 2004), and that the matching of workers to firms improves with market thickness (Bayer et al., 2008). Yet accessing these opportunities requires securing and retaining housing in expensive, supply-constrained markets. When housing tenure is disrupted, workers may lose not only their homes but also their foothold in the labor markets that generate these returns.

We study the labor market consequences of losing housing in a high-productivity city. We exploit mass evictions triggered by California's Ellis Act, which allows landlords to permanently withdraw buildings from the rental market. Because the Ellis Act requires that *all* units in a building be vacated simultaneously, landlords cannot selectively target individual tenants, substantially reducing tenant-level selection. We identify approximately 1,300 Ellis Act withdrawal events in San Francisco between 1998 and 2019, affecting roughly 3,500 tenants in our linked sample. Using a stacked event-study design, we compare each displaced tenant to residents of other rent-controlled buildings within 200 meters—neighbors who face the same micro-neighborhood conditions and the same institutional housing regime, but whose buildings were not withdrawn. We combine this reduced-form design with a spatial equilibrium model that disciplines how housing loss translates into job access, commuting, and career growth within the city.

We address four sets of questions. First, what are the *average* effects of displacement on earnings and residential outcomes over one to six years? Second, *who* is most affected? We examine heterogeneity by age, baseline income, and distance moved from the city center. Third, through what margins do these effects arise? We investigate whether earnings losses reflect transitions to lower-paying firms, changes in establishment characteristics, or increases in commute distance. Fourth, can these patterns be explained by a spatial equilibrium model in which housing frictions shape labor market access within the city?

The empirical analysis exploits a new longitudinal dataset that links households' addresses, building characteristics, federal income records, and employer–employee matches from the Longitudinal Employer-Household Dynamics (LEHD) program. This linked panel allows us to track individuals' earnings, employers, housing quality, and neighborhood characteristics before and after an Ellis Act displacement event.

We find large, persistent effects across the main outcome domains. Six years after evic-

tion, displaced individuals earn approximately $13,000 less per year than their matched neighbors, a gap equivalent to roughly 20 percent of pre-eviction earnings. This earnings divergence emerges gradually: there is no discontinuous drop at the time of eviction, but rather a steadily widening gap over the following years. In contrast, residential outcomes deteriorate immediately. Evicted tenants move to buildings with assessed values per square foot that are roughly 0.6 log points lower than their pre-eviction housing, and to census tracts where median home values are approximately 20 percent lower, indicating persistent downgrading of both their own housing and neighborhood.

We document striking heterogeneity in these effects. Younger workers (under 30 at the time of eviction) experience earnings losses that are proportionally much larger relative to their baseline income, consistent with displacement interrupting the steep, early-career portion of the urban earnings profile documented by De La Roca and Puga (2017). Workers with below-median baseline income experience similar absolute earnings declines as higher baseline earners but from a much lower base, resulting in proportionally larger losses. In both cases, the residential downgrading is at least as large as for their respective comparison groups.

Earnings losses are large even among those who leave San Francisco but remain in the Bay Area. Those who find new housing within 5 km do not experience any detectable earnings effects, while those who move 5–25 km—far enough to lose convenient access to San Francisco's labor market, but not far enough to lose all access—see earnings gaps approaching $40,000. This pattern points to an important role for hyper-local labor market access, even within the same commuting zone.

To shed light on mechanisms, we examine employer characteristics and commuting patterns. Workers who move 5–25 km gradually transition to smaller establishments and lower-wage firms outside the city center, with the effects concentrated among younger workers observed 4–6 years post-eviction. The same groups that experience the largest earnings losses also see the largest increases in commute distance. Taken together, these patterns are consistent with displacement degrading access to the dense, high-wage labor market in San Francisco's urban core, either by increasing the friction of commuting or by reducing exposure to the agglomeration spillovers that drive urban wage growth.

In the second part of the paper, we develop a spatial equilibrium model to interpret these findings and to ask whether housing loss can generate persistent earnings effects by degrading access to the urban job ladder. The model separates where workers live from where they work inside a monocentric city, embeds search frictions in both labor and housing markets, and allows commuting to affect both utility and match productivity. Workers differ over the life cycle, accumulate human capital dynamically, and move up a job ladder through employed

and unemployed search. Firms differ in productivity and are more concentrated near the urban core, so a forced residential move can reduce earnings both by worsening current access to high-wage firms and by lowering future exposure to the locations and matches that generate urban wage growth.

We calibrate the model to the spatial structure of the San Francisco Bay Area using moments on population density, rents, job density, the distribution of high-wage firms, commuting, and mobility. The calibrated economy reproduces the broad geography of the region, with higher rents, higher wages, and a greater concentration of productive firms near the center. When we simulate eviction from central locations, the model matches the empirical findings of immediate residential downgrading, longer commutes, and a gradual decline in earnings as displaced workers sort over time toward more local, less productive, slower-learning firms. However, it understates the size of the long-run earnings losses, suggesting that the value of central-city access may operate through even stronger learning or faster progression to high-productivity firms than in the baseline calibration.

Related literature. The reduced-form evidence and the model connect the paper to several literatures. At a high level, the paper sits at the intersection of dynamic agglomeration, spatial search, commuting, and housing-constrained spatial equilibrium. In the spirit of De La Roca and Puga (2017), the model emphasizes that productive cities matter for earnings growth over the life cycle, not only for static wage levels. Like Manning and Petrongolo (2017) and recent frictional mobility models such as Schmutz and Sidibé (2018), it treats labor markets as spatially local, so residential location shapes job access and job-to-job mobility. Like Monte et al. (2018), it separates residence from workplace and gives commuting an explicit role. Like Hsieh and Moretti (2019), it interprets housing constraints as barriers to accessing high-productivity labor markets. Relative to dynamic urban models such as Martellini (2022), the distinctive contribution here is to bring these ingredients into a within-city framework with an explicit housing market and a forced-displacement shock.

First, we add to work on agglomeration and the dynamics of earnings. A robust finding in urban economics is that wages are higher in denser areas (Ciccone and Hall, 1996; Moretti, 2004; Combes et al., 2012), and that workers' earnings grow faster in larger cities, with part of the growth persisting even after they leave (Glaeser and Maré, 2001; Baum-Snow and Pavan, 2012; De La Roca and Puga, 2017). A central challenge in this literature is distinguishing agglomeration-driven productivity, learning from selection, and sorting of workers across places (e.g., Diamond, 2016; Combes et al., 2008; Gould, 2007; Nix, 2020). We contribute a complementary source of identification: a sharp, plausibly exogenous shock to residential access within an exceptionally productive city. We show that the returns to urban exposure

are not only about "having been" in a large city, but also about *continued access* to dense job networks: workers pushed away from the urban core experience a persistent break in their earnings trajectories relative to otherwise similar neighbors who remain.

Second, we contribute to the literature on housing constraints, mobility, and spatial sorting. In spatial equilibrium, productivity advantages are capitalized into housing costs, so housing supply constraints and regulation shape who can access (and remain in) high-wage labor markets (Rosen, 1979; Roback, 1982; Saiz, 2010; Glaeser et al., 2005; Gyourko et al., 2013). A large body of work argues that these constraints can reduce aggregate output by limiting migration toward productive places (Ganong and Shoag, 2017; Hsieh and Moretti, 2019). Related evidence shows that rent regulation and tenant protections can reshape neighborhood composition and housing market trajectories (Autor et al., 2014; Diamond et al., 2019; Asquith, 2019; Pennington, 2021), and that losing protected housing can harm tenants' economic outcomes (Cerqueiro et al., 2025). We study a complementary margin that is harder to capture in cross-city migration designs: even conditional on having already sorted into a high-opportunity neighborhood, households can lose access through displacement. We document the importance of continued housing access for *incumbent* residents – quantifying the earnings costs of being forced out of a top-wage location rather than the gains from moving into one (Chetty et al., 2016; Chetty and Hendren, 2018; Bergman et al., 2024).

Third, we contribute to the literature on spatial frictions in labor markets. Labor markets are remarkably local: workers trade off wages against commuting distance and job accessibility (Manning and Petrongolo, 2017; Le Barbanchon et al., 2021), and classic "spatial mismatch" mechanisms link residential location to employment and wage outcomes through access to jobs and networks (Kain, 1968; Gobillon et al., 2007; Marinescu and Rathelot, 2018). Recent research also highlights commuting as a key margin through which local shocks propagate across neighborhoods and nearby labor markets (Monte et al., 2018). We contribute by documenting a setting in which residential location shifts sharply while job separations are not mechanically forced, allowing us to isolate how increased distance to jobs, reduced access to dense labor market interactions, and disrupted job-to-job mobility translate into persistent earnings losses. This provides micro-level evidence that "where you live" within a superstar city can causally shape earnings dynamics and motivates our focus on eviction-driven displacement.

We also contribute to the literature on eviction and displacement. Because Ellis Act evictions are building-wide and not triggered by tenant nonpayment, they provide a rare opportunity to separate the consequences of losing access to a location from the financial distress that often accompanies eviction filings (Collinson et al., 2024); consistent with this interpretation, we find no earnings pre-trends before the eviction event. Our findings also

inform the growing work on equilibrium housing policy (Abramson, 2021; Abramson and Landvoigt, 2025; Waldinger, 2021) and complement evidence from large displacement shocks. After Hurricane Katrina, New Orleans households experienced substantial displacement yet only short-lived earnings losses; longer-run income gains were largest among those who left the city, consistent with relocation to stronger labor markets and post-storm wage increases in New Orleans (Deryugina et al., 2018). Likewise, Nakamura et al. (2021) show that a mobility shock can raise long-run earnings even when it induces moves away from a high-income location, but the gains accrue primarily to children through increased education and better skill matching, while parents bear most of the costs. Taken together, these studies suggest that forced moves can raise or lower earnings depending on whether they reallocate workers toward or away from places that reward their skill. In our setting, displacement from San Francisco—near the top of the skill-reward distribution—reduces access to those returns and leads to persistent earnings losses.

Finally, our results also provide an informative contrast with the findings of the job displacement literature, where the Ellis Act can be seen as a housing analog to a mass layoff event. While earnings losses from job loss are sharp (Jacobson et al., 1993; Lachowska et al., 2020), workers tend to recover half their initial earnings losses after six years in non-recessionary environments (Stevens, 1997; Couch and Placzek, 2010; Davis and von Wachter, 2011). However, even workers who voluntarily quit (e.g., if a new commute is too long) can experience persistent earnings losses if the jobless spell is sustained (Fallick et al., 2025). Unemployed workers tend to have shorter unemployment durations and stronger earnings recoveries in large, dense labor markets (Andersson et al., 2018; Moretti and Yi, 2025). Our findings complement this literature by underscoring how sensitive workers are to residential location as part of their ability to achieve the optimal employee-employer match: workers displaced under the Ellis Act who move out of San Francisco, even if they otherwise manage to stay in the Bay Area, experience steadily growing earnings declines with no observed recovery even six years *ex post*.

2 Institutional Context

2.1 Rent Control in San Francisco

San Francisco's Rent Stabilization Ordinance, enacted in 1979, applies to most residential buildings with two or more units built before June 13, 1979—representing nearly 70 percent

of rented units in the city.[1] When a new tenant moves in, the landlord sets the initial rent freely (vacancy decontrol), a provision mandated statewide by the Costa-Hawkins Act of 1995. After move-in, annual rent increases are capped at 60 percent of the Bay Area Consumer Price Index, with a ceiling of 7 percent; in practice, this averaged approximately 1.7 percent per year during our study period (2000–2019). Tenants may only be evicted for specific "just cause" reasons, such as nonpayment, nuisance, or owner move-in, and landlords must follow formal eviction procedures, including filing a notice with the Rent Board.

This combination of rent stabilization and just-cause eviction protections means that long-tenure tenants in rent-controlled buildings often pay rents substantially below market rates. Losing a rent-controlled unit thus entails both the disruption of moving and the loss of a valuable implicit housing subsidy.

2.2 The Ellis Act

The Ellis Act is a California statute enacted in 1985 in response to *Nash v. City of Santa Monica* (1984), in which the California Supreme Court upheld a city's right to prevent a landlord from withdrawing a rent-controlled building from the rental market. The legislature responded by passing the Act to establish the right of property owners to "go out of the rental business," subject to substantial tenant protections.

An Ellis Act withdrawal begins when the building owner files a notice of intent with the San Francisco Rent Board. Tenants generally receive 120 days' notice, while elderly (age 62+) and disabled tenants may receive up to one year.[2] The owner must pay relocation benefits of approximately $5,000–$7,000 per tenant (varying over time with inflation adjustments; see Appendix A). Crucially, the Act requires that *all* accommodations in the building be withdrawn simultaneously—the owner cannot selectively evict individual tenants.

Figure 1 summarizes the post-withdrawal restrictions. If the owner re-rents within two years, they are liable for damages and must offer the units back to displaced tenants. For five years, any re-rental must be at the prior lawful rent plus allowed annual adjustments—effectively preventing rent resets. San Francisco extends a right of return for up to ten years: if units are re-offered, the owner must notify the Rent Board and offer them to former tenants. If the building is demolished and rebuilt within five years, the replacement units are treated as rent-controlled.

These restrictions make Ellis withdrawals costly to reverse. In practice, the vast majority

[1] Newer buildings, most single-family dwellings, dormitories, nonprofit cooperatives, and units whose rents are regulated by another government program (e.g., public housing, LIHTC) are excluded. A small number of buildings have earned an exemption after undergoing "substantial rehabilitation."

[2] Until January 1, 2000, all tenants only received 60 days to vacate. See Appendix A for more information.

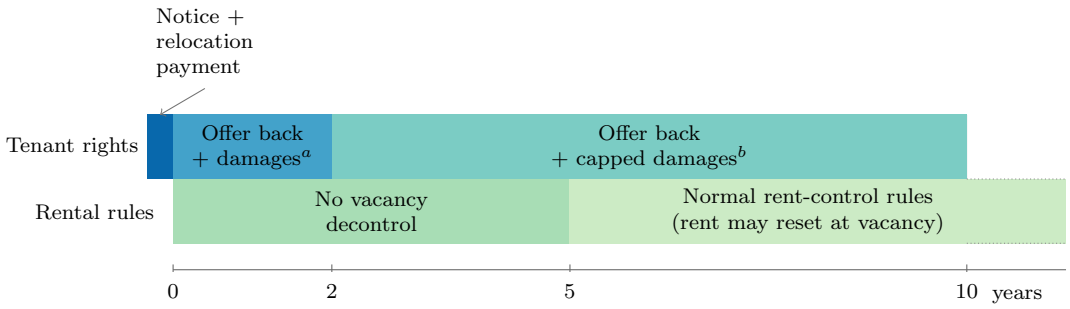


Figure 1: Timeline of Ellis Act provisions (San Francisco)

Notes: Time is measured in years since the withdrawal date. The pre-zero segment reflects the statutory notice period (120 days; up to one year for qualifying elderly/disabled tenants) and SF relocation payments. If a withdrawn unit is re-rented within ten years, SF requires that it first be offered to displaced tenants who timely requested notification. Enforcement differs by horizon. Within two years ([a]): noncompliant landlords face tenant lawsuits with possible punitive damages, and the City may also sue. After two but within ten years ([b]): displaced tenants may recover up to six months' contract rent. For any new tenancy begun within five years of withdrawal, rents are capped at the prior lawful rent plus allowable annual adjustments (no vacancy decontrol); standard SF rent-control rules apply thereafter. Sources: Cal. Gov. Code §7060.2(b)–(c); SF Admin. Code §37.9A(c)–(d).

of withdrawn buildings are converted to condominiums or to tenancy-in-common arrangements, the most common marketable alternative in San Francisco where condo conversions are rationed.[3] The restrictive re-rental rules and the pattern of conversion to ownership suggest that Ellis filings are motivated by the desire to exit the rental business entirely, rather than to remove specific tenants.

2.3 Ellis Act Events in San Francisco

Figure 2 maps all Ellis Act withdrawal events we observe between 1998 and 2019. Three features of this map are relevant for our empirical strategy. First, there are enough events for credible estimation: we identify approximately 1,300 distinct events affecting an estimated several thousand units. Second, the events are spread across the city but concentrated near the urban core, precisely where housing values and labor market access are highest. Third, because Ellis withdrawals remain a relatively low-probability event for any given building, they are unlikely to be fully anticipated by tenants or priced into location decisions.

[3]Only about 10 percent of withdrawn buildings are eventually converted through the city's condo conversion lottery.

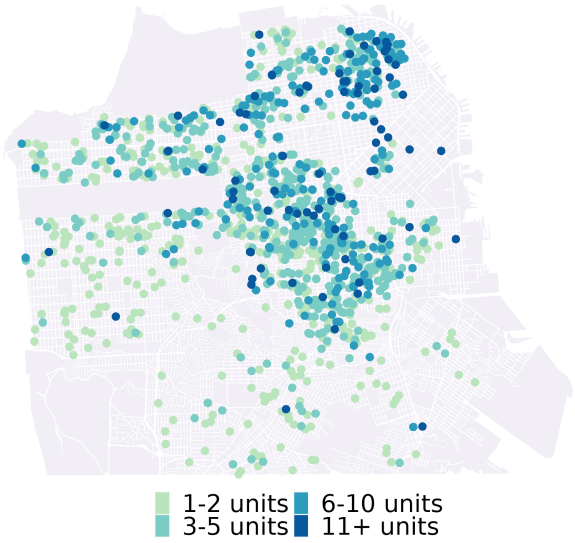


Figure 2: Ellis Act Building Withdrawals in San Francisco, 1998–2019

Notes: Each dot represents an Ellis Act building withdrawal event. Events are concentrated near the city center but occur throughout the city. Source: San Francisco Rent Board eviction filings.

3 Data

We construct an individual-level panel combining administrative records on addresses, employment, earnings, and housing characteristics.

Ellis Act evictions and treated buildings. We begin with the universe of eviction filings from 1998–2019 provided by the City and County of San Francisco (City and County of San Francisco, 2024). These records include the stated reason for eviction, filing date, property address, and parcel number. We restrict the sample to evictions whose reason is given as "Ellis Act" and consolidate filings for units on the same parcel into a single event ID, moving from 3,727 filings to 1,311 events. We estimate the number of affected units by counting unique address–unit combinations within each Ellis event. To reduce the chance of selection into an Ellis filing based on tenant characteristics, we limit our analysis to Ellis filings that targeted buildings with 3 or more units. In our main results, we further restrict to Ellis filings where 50% of residents moved out within one year and 75% moved out within three years. This increases confidence that a filing resulted in an eviction, which is not always the case if tenants and landlords can come to an agreement or there was an issue with the filing. These two restrictions leave us with 513 Ellis events.

To link these records to Census microdata, we spatially merge a shapefile of San Francisco land parcels with Census's internal property identifier, the MAFID, recorded in the Census

Master Address File (MAF). We then match Ellis records to the San Francisco MAF universe using standardized address strings and a deterministic, tiered linkage procedure that proceeds from exact matches on normalized addresses to less restrictive rules that ignore unit designators. This does not affect our research design, because we only need to identify residence in an evicted building rather than a specific unit. Matched records are removed at each stage to prevent duplicate assignments. Finally, we merge assigned MAFIDs to the parcel-MAFID crosswalk to attach parcel identifiers and Census geographic codes. Even with this strict requirement for address matching, we successfully match all Ellis events to at least one MAFID.

Control buildings. To construct the pool of potential control buildings, we combine property characteristics from the San Francisco Assessor's Historical Secured Property Tax Rolls and Black Knight property assessment files (2017–2019). We retain residential buildings with at least three units, built by 1978 so that they are also rent-controlled, and located within 200 meters of an Ellis Act event. The same building may serve as the control for multiple Ellis events.

Individual address histories. We identify treated and control individuals by intersecting treated and control MAFIDs with an original dataset of individual address histories. These histories draw on administrative sources that report residential location at a point in time, including 1040 and 1098 tax records, USPS change-of-address filings, and participation records from HUD. For each person, uniquely identified by a Census-assigned Protected Identification Key (PIK), we observe a sequence of timestamped MAFIDs which we use to infer the individual's residence at the time of an Ellis filing.

We identify treated individuals by merging Ellis events with address histories on MAFID and determining whether the filing date falls within an individual's inferred residence spell at the evicted MAFID. Because address histories may be interval-censored or incomplete, we apply a tiered set of deterministic overlap rules. Individuals are classified as treated when the filing date falls within their residence window (i.e., they are observed at the evicted MAFID both before and after the filing) or when the filing date is bracketed by adjacent address records (i.e., they are observed at the evicted MAFID before the filing and observed elsewhere afterward, or observed elsewhere before the filing and at the evicted MAFID afterward). We exclude individuals whose records indicate that they moved in after the filing or left beforehand. If an individual is linked to multiple Ellis events, we assign a single event using a fixed priority rule that favors stronger temporal overlap and closer timing to the filing date. Finally, we include spouses observed in the year prior to the filing.

Controls are defined analogously as individuals living in a control building at the time of the Ellis filing. To ensure a clean comparison, we exclude all potential controls who have ever

lived at an evicted address. As shown in Figure 2, Ellis events are spatially concentrated, so individuals may serve as controls for multiple events if they lived within 200 meters of more than one filing. We leverage this structure in a stacked difference-in-differences design, described in Section 4.2.

Demographics and Family Relationships. Demographics and family relationships come from the Census Databank, a cleaned repository built from 1040 and 1098 tax records, Census Numident birth and death records, and demographic information from the 2000 and 2010 Decennial Censuses and the American Community Survey (U.S. Census Bureau, 2024). Family relationships are identified through joint filing with a spouse and by claiming dependents. We merge these data to our sample on PIK for time-invariant characteristics (e.g., race and birth year) and on PIK–year for time-varying relationships.

Employment. We combine earnings and employment information from the Census Databank and the Longitudinal Employer–Household Dynamics (LEHD) infrastructure files. The Databank includes W-2 earnings, along with the tax identification number (TIN) of the highest-paying W-2 employer in a given year. LEHD is constructed primarily from state unemployment insurance (UI) wage records and provides quarterly earnings for each worker–employer match, covering nearly all private-sector employers and most state and local government employers subject to UI reporting requirements. Together, these sources capture wage and salary employment reported on a W-2 or through UI-covered employers and therefore exclude self-employment and other non-UI-covered earnings. Our primary annual income measure is the maximum of annual W-2 earnings and the sum of quarterly LEHD earnings. We merge these measures to our sample on PIK–year.

Housing characteristics. We attach MAFID-matched property assessment data from Black Knight. For each MAFID in the treated and control address histories, we pull location and housing characteristics for the year 2019 including assessed value, square footage, number of bedrooms, and year built. We merge these measures to our sample on MAFID.

Neighborhood characteristics. We include tract-level neighborhood characteristics including 2012 median household income, home values, and rents from the Longitudinal Tract Database (LTDB) Logan et al. (2014) and 2013 third-grade math scores from the Opportunity Insights data library Opportunity Insights (2024).

Sample construction. Our estimation sample includes individuals aged 18–64 at the time of the Ellis filing. We observe outcomes for up to four years before and six years after the event. The final sample consists of approximately 3,500 treated individuals and 403,800 controls across all events.

4 Empirical Strategy

4.1 Identification

Our goal is to estimate the dynamic causal effect of displacement on earnings and residential outcomes. The key advantage of the Ellis Act setting is that building-wide withdrawals prevent landlords from selectively targeting individual tenants. However, landlords choose which building to withdraw and when, so we cannot simply compare Ellis tenants to all other San Francisco renters.

For each Ellis filing, we compare tenants in the treated building to tenants living within 200 meters in other rent-controlled buildings—a narrow geographic band that ensures treated and control individuals face the same micro-neighborhood conditions, the same local labor market, and the same institutional housing regime. Our core identifying assumption is a local parallel trends condition: absent the Ellis withdrawal, treated tenants and their nearby rent-controlled neighbors would have followed the same trajectory in earnings and other outcomes.

This assumption would be violated if landlords systematically time Ellis filings in response to building-level shocks that predict tenants' future outcomes—for example, if treated buildings housed tenants who were already on a deteriorating earnings trajectory. We assess this assumption in two ways. First, we provide a detailed balance table that treated and control individuals are comparable on a wide range of baseline characteristics. Second, we verify the absence of differential pre-trends in both the raw data and in our event-study estimates.

Table 1 compares treated and control individuals in the year prior to the Ellis event. We report normalized differences in the third column, following the conventional benchmark that values below 0.1 in absolute value indicate negligible imbalance (Imbens and Wooldridge, 2009; Imbens and Rubin, 2015).[4] No propensity-score matching or reweighting is applied; these are raw comparisons.

The two groups are well balanced along most dimensions. Demographics are reasonably similar, though treated individuals are somewhat older (43 vs. 39 years) and slightly more likely to be married or have children.[5] Labor market outcomes are closely aligned: the groups have similar median incomes ($39,600 vs. $41,500), similar rates of employment

[4] A normalized difference of 0.1 means the groups differ by 0.1 pooled standard deviations (a small effect size).

[5] If anything, the fact that our treated sample is somewhat older and more likely to be married reinforces our identification assumption that landlords are not using the Ellis Act to target individual tenant households. Ellising older tenants is costlier: landlords have to give tenants over 62 higher relocation payments and longer notice periods. Further, since 2000, relocation payments are allotted on a per tenant basis, meaning a married couple is more expensive to Ellis than a singleton.

Table 1: Balance Between Treated and Control Individuals at $t = -1$

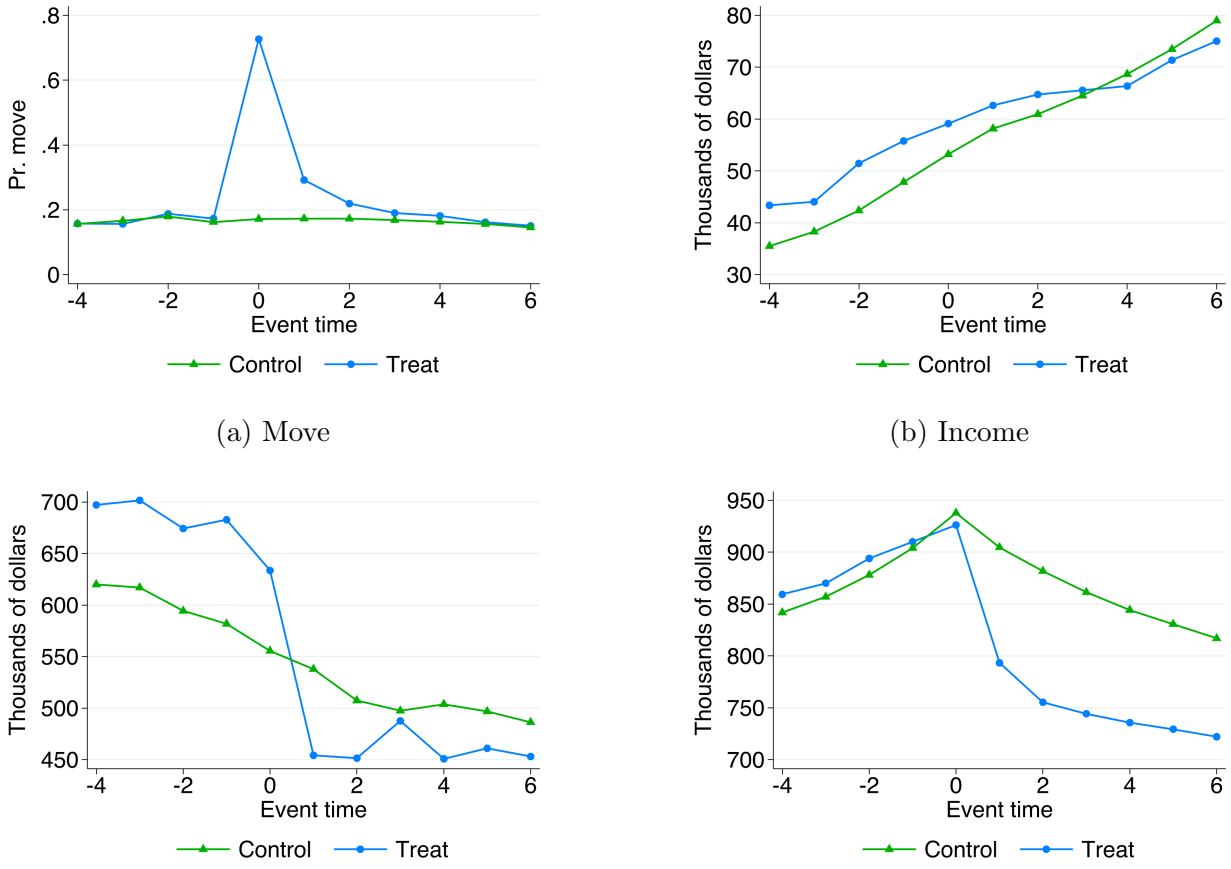
Group	Variable	Treat	Control	Norm. diff.
Demographics	Female	0.468	0.485	-0.034
	Black	0.033	0.033	-0.003
	White	0.595	0.516	0.159
	Age	43.1	39.24	0.252
	Married	0.25	0.194	0.135
	N children	0.184	0.123	0.118
Labor Market	Median income	39,600	41,500	-0.023
	Tech or finance industry	0.127	0.123	0.011
	Median firm average wage	72,300	77,400	-0.072
	Median estab. size (employees)	70	66.3	0.013
	Median commute distance (km)	5.33	5.48	-0.062
Housing	Tenure in current unit (years)	5.72	5.57	0.029
	Log assessed value/sqft	5.80	5.58	0.103
Neighborhood	Median tract rent	1,600	1,560	0.106
	Median tract home values	907,600	901,600	0.031
	Tract unemployment rate	0.048	0.050	-0.082
	Tract third grade math scores	0.149	0.155	-0.021
	N individuals	3,500	403,800	

This table compares treated and control individuals in the year prior to the Ellis event. Normalized differences below 0.1 in absolute value indicate negligible imbalance. No propensity score matching or reweighting is applied. Income, wages, and establishment size are reported at the median to reduce the influence of outliers.

in San Francisco's high-paying technology and finance industries, comparable median firm wages, establishment sizes, and commute distances. Housing tenure in the current unit is nearly identical (5.7 vs. 5.6 years), further reassuring us that landlords are not targeting tenants with longer tenures who pay submarket rents, and assessed building values are similar. Neighborhood characteristics—tract-level rents, home values, unemployment rates, and school quality—are comparable almost by construction, given the narrow 200-meter control radius. The control group is more than 100 times larger than the treatment group, reflecting the large number of rent-controlled buildings within each 200-meter catchment area.

Figure 3 displays raw outcome means for treated and control groups around the Ellis event. Panel (a) confirms that the Ellis filing successfully identifies a displacement shock: approximately 75 percent of treated individuals move during the event year, compared to a much lower baseline rate in the control group. The slightly elevated moving rate at event time +1 likely reflects elderly and disabled tenants who receive a full year's notice, as well as variation in when the filing falls within a calendar year.

Panel (b) shows that the two groups follow roughly parallel income trajectories before



(a) Move

(b) Income

(c) Assessed Value per Square Foot

(d) Tract Median Home Value

Figure 3: Raw Outcome Trends for Treated and Control Groups Around Ellis Events

Notes: Each panel plots raw group means for treated (Ellis-displaced) and control (nearby rent-controlled tenants) individuals by event time. Panel (a) shows the probability of moving in each year. Panel (b) shows annual income. Panel (c) shows the log assessed value per square foot of the individual's residence. Panel (d) shows the log median home value of the census tract of residence. Event time 0 is the year of the Ellis filing.

the event, after which the treated group's income growth decelerates while the control group continues on an upward trend. Panels (c) and (d) reveal immediate drops in housing and neighborhood quality for the treated group. The control group also trends downward in both measures, likely reflecting life-cycle patterns of residential mobility away from the most central neighborhoods. This motivates our event-study design, which differences out these common trends.

4.2 Stacked Event-Study Design

The staggered timing of Ellis events motivates a stacked event-study design. For each Ellis event d, we pair treated tenants with their local controls and stack these event-specific

datasets. We estimate:

$$Y_{itd} = \sum_{k=-\kappa_a}^{\kappa_b} (\beta_k + \delta_k \cdot T_{id}) \cdot \mathbb{1}\{E_{td} = k\} + X_{it}\gamma + \theta_{id} + \mu_t + \varepsilon_{itd}, \tag{1}$$

where Y_{itd} is the outcome for individual i in calendar year t exposed to Ellis event d; E_{td} denotes event time (years since the filing); T_{id} is an indicator for whether individual i was treated in event d; X_{it} are time-varying controls (age polynomials); θ_{id} are individual-by-event fixed effects; and μ_t are calendar year fixed effects. The coefficients of interest are $\{\delta_k\}$, which capture the treatment effect at each event-time horizon. We normalize $\delta_{-2} = 0$.

Standard errors are clustered at the event level. To ensure that each event contributes equally to the estimates regardless of the number of available controls, we weight control observations inversely by the number of controls per event, so that the effective number of controls per treated individual is equalized across events. This stacked design avoids the well-documented biases of two-way fixed effects estimators with staggered treatment timing and heterogeneous treatment effects (de Chaisemartin and D'Haultfœuille, 2020; Sun and Abraham, 2021; Borusyak et al., 2024).

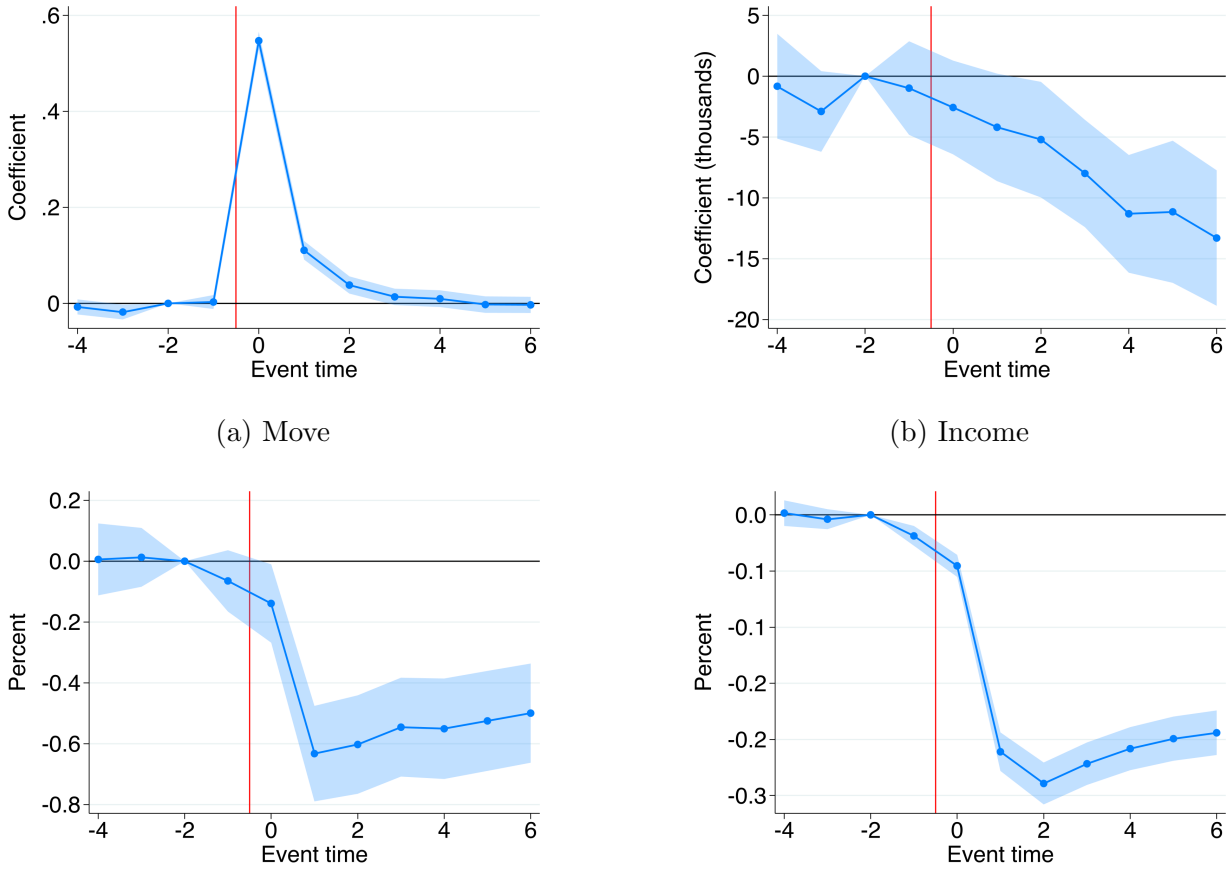
5 Results

5.1 Average Effects of Displacement

Figure 4 presents the main event-study estimates. Importantly, there is no differential pre-trend in any of the four panels, supporting the parallel trends assumption. Panel (a) confirms a sharp, immediate increase in the probability of moving: displaced tenants are approximately 60 percentage points more likely to move in the event year, and their probability of moving remains elevated until three years after the eviction.

Panel (b) presents the central finding: a gradual, widening earnings divergence. The treatment effect on annual income is small and statistically insignificant in the first year or two, then grows steadily, reaching approximately −$13,000 by year six. The slow emergence of the earnings gap—rather than an immediate drop—is notable. It suggests that displacement does not cause an abrupt earnings loss (as job displacement shocks tend to do, see e.g. Jacobson et al. (1993); Davis and von Wachter (2011); Lachowska et al. (2020)) but rather puts workers on a lower-growth trajectory, consistent with the disruption of the career-building process that high-productivity cities facilitate.

Panels (c) and (d) show a contrasting temporal pattern for residential outcomes. The assessed value per square foot of the treated group's housing drops by roughly 0.6 log points



(a) Move

(b) Income

(c) Log Assessed Value per Square Foot

(d) Log Tract Median Home Value

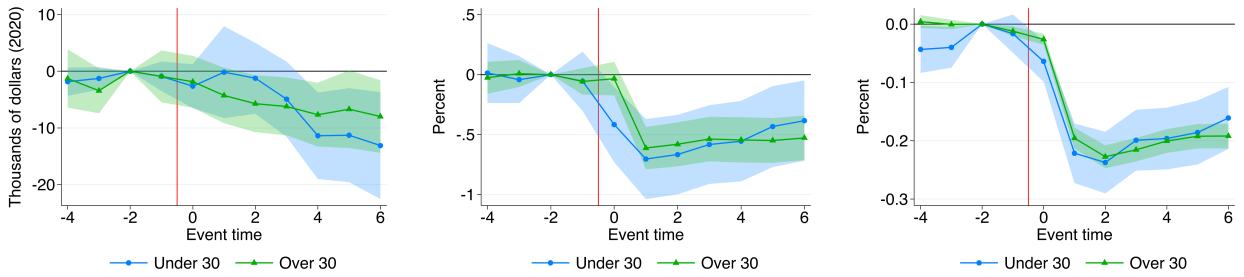
Figure 4: Event-Study Estimates: Average Effects of Ellis Act Displacement

Notes: Each panel plots estimated δ_k coefficients from equation (1), with 95% confidence intervals based on standard errors clustered at the Ellis event level. The omitted category is event time -2. Panel (a): probability of having moved from the pre-event address. Panel (b): annual income in dollars. Panel (c): log assessed value per square foot of current residence. Panel (d): log median home value in the census tract of residence.

immediately upon displacement, with modest recovery over the following six years. Tract-level median home values fall by approximately 0.2 log points (about 20 percent). These residential effects are large and immediate, in contrast to the gradual deterioration of earnings, suggesting that while households are quickly pushed to lower-quality housing and neighborhoods, the labor market consequences take years to fully materialize.

5.2 Heterogeneity by Age and Baseline Income

Figure 5 splits the sample by age at the time of the Ellis event, using a cutoff of 30 years. Although average pre-eviction earnings are lower for those under 30 (approximately $42,000 vs. $52,000 for older workers), the decline in income shown in Panel (a) is at least as large for



(a) Income (b) Log Value per Sq. Ft. (c) Log Tract Home Value

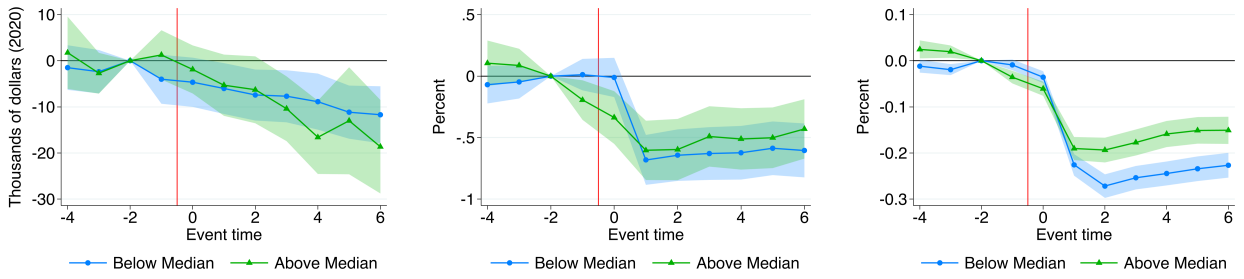
Figure 5: Event-Study Estimates by Age Group (Under 30 vs. 30 and Over)

Notes: Each panel plots separate event-study coefficients for individuals under 30 and 30 or older at the time of the Ellis filing. Confidence intervals are at the 95% level, with standard errors clustered at the event level.

younger workers, resulting in a proportionally much more severe impact. By six years post-eviction, the earnings gap for younger workers represents approximately 30 percent of their pre-eviction income, compared to about 15 percent for older workers. This age heterogeneity is consistent with displacement interrupting the steep, early-career phase of urban earnings growth, where learning on the job is an important part of human capital acquisition for younger workers (Nix, 2020). In the model of De La Roca and Puga (2017), young workers benefit disproportionately from learning in dense urban environments, accumulating portable human capital that generates returns throughout their careers. Being pushed away from the urban core during this critical period may permanently reduce the stock of experience-based human capital, with compounding effects over time.

In contrast, housing and neighborhood outcomes (Panels b and c) show relatively similar downgrading across age groups, suggesting that the differential earnings effects are not driven by younger workers moving to disproportionately worse neighborhoods. Rather, it appears that the same residential displacement has larger labor market consequences for workers at an earlier career stage.

Figure 6 presents heterogeneity by baseline income. Workers starting below the median (approximately $40,000) experience absolute earnings declines similar in magnitude to those above the median. These losses are much larger in proportion to their lower baseline income, consistent with lower-income households having less liquidity to cushion the displacement shock and maintain access to high-wage employment. If anything, the residential downgrading is somewhat more pronounced for the below-median group (Panels b and c), suggesting that budget constraints force lower-income evictees into neighborhoods and buildings of lower quality.



(a) Income (b) Log Value per Sq. Ft. (c) Log Tract Home Value

Figure 6: Event-Study Estimates by Baseline Income (Below vs. Above Median)

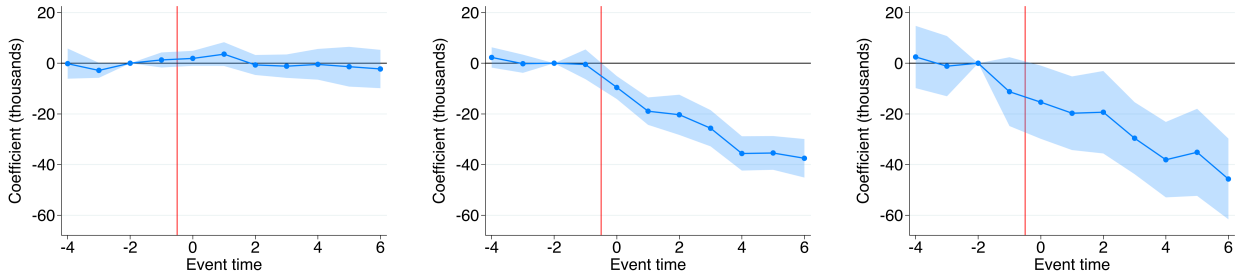
Notes: Each panel plots separate event-study coefficients for individuals with baseline income below and above the sample median (approximately $40,000). Confidence intervals are at the 95% level, with standard errors clustered at the event level.

5.3 Earnings Effects by Distance Moved

Next, we consider heterogeneity by distance moved from their pre-eviction address at event time +3. Because distance moved is itself an outcome of displacement, this analysis is descriptive rather than causal, but it reveals informative patterns about the role of location in the earnings effects. In these event studies, treated people who have moved a given distance are compared to all controls, regardless of how far a control may have moved.

Figure 7 presents the results. Individuals who remain within 5 km of their pre-eviction address (Panel a) do not experience economically or statistically significant earnings effects. In sharp contrast, those who move 5–25 km (Panel b)—far enough to leave San Francisco but not the broader Bay Area—experience earnings gaps approaching $40,000 by year six. Those who leave entirely (Panel c, > 25 km) also show substantial earnings losses.

The finding that the intermediate-distance group experiences large earnings effects is particularly informative. These individuals moved far enough to lose convenient access to San Francisco's dense, high-wage labor market, yet remained close enough that they have not fully "reset" in a new local economy. This pattern suggests an important role for hyper-local labor market access—even within the same commuting zone—consistent with the spatial job search literature documenting that labor markets are far more local than commuting zone boundaries imply (Manning and Petrongolo, 2017; Marinescu and Rathelot, 2018).



(a) Moved 0–5 km (b) Moved 5–25 km (c) Moved > 25 km

Figure 7: Event-Study Estimates for Income by Distance Moved at Year 3

Notes: Each panel plots event-study coefficients on income for a subsample defined by distance from the pre-eviction address at event time 3. Panel (a): individuals within 5 km ("stayers"). Panel (b): 5–25 km (left San Francisco but remained in the Bay Area). Panel (c): more than 25 km. Distance moved is an endogenous outcome; these splits are descriptive.

5.4 Firm Characteristics and Commuting

To investigate the mechanisms behind the earnings divergence, we examine whether displaced workers transition to different types of employers.

Table 2 interacts the treatment effect with age group, event-time horizon, and distance-moved category. Two findings stand out. First, the groups with the largest earnings losses—younger workers who move 5–25 km—also exhibit the largest declines in establishment size and firm average wages. In years 4–6, young workers in this group work at establishments with 35 fewer employees and at firms paying salaries that are $17,000 lower on average. Second, commute distance increases substantially for the groups that move farther, and especially for younger workers. Young workers who move 5–25 km see commute distances rise by 6–8 km, while those who remain within 5 km experience much smaller changes.

These patterns are consistent with two complementary mechanisms. The increase in commute distance may directly reduce labor supply or degrade the quality of the worker-firm match, as documented in the commuting literature (Gutiérrez-i-Puigarnau and van Ommeren, 2010; Le Barbanchon et al., 2021). At the same time, the transition to smaller, lower-wage establishments suggests that displacement reduces access to the types of employers concentrated in San Francisco's dense urban core—large, high-wage firms in technology and professional services—pushing workers toward a different segment of the employer distribution. Both channels affirm that housing location is key to labor market access.

Table 2: Establishment Size, Firm Wages, and Commute Distance by Age and Distance Moved

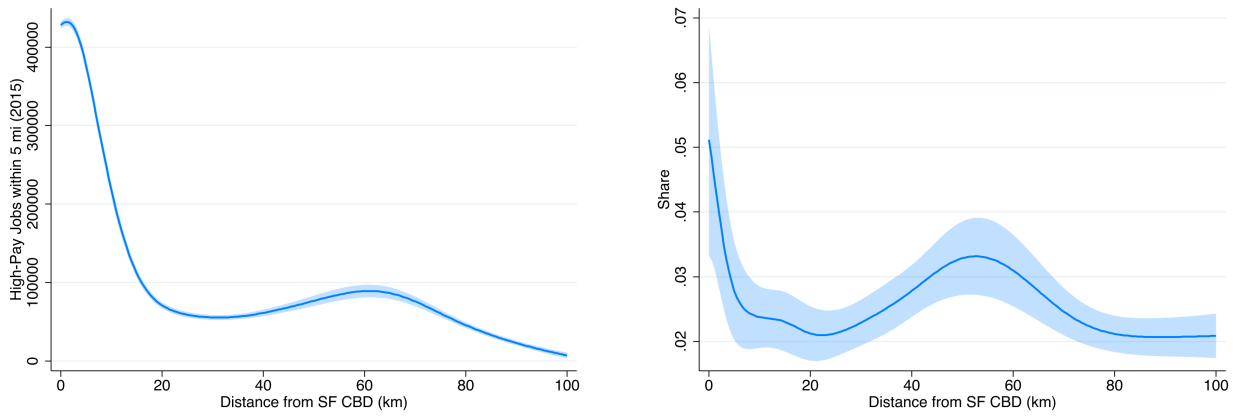
	Under 30			Age 30+		
	Estab. Emp.	Firm Wage ($k)	Commute (km)	Estab. Emp.	Firm Wage ($k)	Commute (km)
Years 0–3: 0–5 km	2.00	−4.80	−0.47	1.06	−0.45	−0.63
	(9.79)	(3.93)	(1.82)	(4.14)	(1.48)	(0.62)
Years 0–3: 5–25 km	−9.56	−8.25	8.04	−1.98	1.73	2.26
	(16.66)	(5.93)	(2.82)	(6.96)	(2.49)	(1.08)
Years 0–3: >25 km	1.80	−3.75	11.84	3.83	0.20	9.72
	(16.47)	(5.16)	(2.64)	(8.13)	(2.96)	(1.40)
Years 4–6: 0–5 km	1.77	−3.74	3.23	0.03	0.45	−0.88
	(12.57)	(4.71)	(2.09)	(5.37)	(1.99)	(0.78)
Years 4–6: 5–25 km	−35.10	−16.86	6.60	−2.15	2.62	2.90
	(21.50)	(6.72)	(3.56)	(9.06)	(3.19)	(1.32)
Years 4–6: >25 km	13.17	0.25	4.76	14.40	−2.04	9.55
	(20.37)	(6.83)	(3.41)	(10.94)	(3.94)	(1.78)
N	409k	407k	409k	1.746m	1.730m	1.745m
Adj. R^2	0.419	0.533	0.261	0.634	0.615	0.462

Each column reports treatment effects from a regression of the indicated outcome on interactions of treatment with event-time period (years 0–3 or 4–6) and distance-moved category (0–5 km, 5–25 km, or >25 km from pre-eviction address at event time 3). The two panels report separate regressions for individuals under 30 and 30+ at baseline. Standard errors (in parentheses) are clustered at the Ellis event level. Establishment employment is measured in levels; firm wages and commute distance are in thousands of dollars and kilometers, respectively.

5.5 Toward an Equilibrium Model of the Bay Area

The empirical findings point to a tight link between housing location and labor market outcomes, but they also raise questions that reduced-form evidence alone cannot answer. Why does moving 15 or 20 kilometers from the center—a modest distance—translate into such large and persistent earnings losses? And if jobs near the CBD are so valuable, why do so many evicted tenants relocate away from the center rather than searching nearby?

Figure 8 offers a first clue: spatial gradients in job quality are remarkably steep. Panel (a) plots the number of high-paying jobs (those paying more than $\$3,333$ per month) within 5 miles of a census tract against distance to the CBD. Moving only 20 km from the center reduces the count of surrounding high-paying jobs by a factor of four. Panel (b) shows that the share of establishments with more than 100 employees is more than twice as high in the city center as 10 km away. These gradients can explain why displacement degrades employer quality, but they cannot by themselves explain why workers who relocate do not simply



(a) High-paying jobs within 5 miles

(b) Establishments with $\geq$ 100 employees

Figure 8: Firm Characteristics by Distance to San Francisco's CBD

Notes: Each panel represents the results of a kernel regression of local firm and establishment characteristics as a function of distance to San Francisco's CBD. Panel (a) reports the number of jobs with earnings greater than $3,333 per month in own and neighboring tracts whose centroids fall within a radius of 5 miles from own tract centroid. Source: Opportunity Insights, constructed using LODES - WAC 2015 from the Census Bureau. Panel (b) reports the share of establishments per ZIP code with more than 100 employees. Source: ZIP Code County Patterns 2019. The regression uses a linear bandwidth of 3 km near the center reaching 15 km at 100 km. The shaded areas represent the 95% confidence intervals from bootstrapped standard errors.

commute back to the center and why their relative earnings continue to deteriorate.

To help disentangle the forces at work, we turn in the next section to a modeling exercise of the Bay Area.

6 Model

We build a spatial equilibrium model of the Bay Area featuring frictional housing and labor markets, costly commuting, localized job search, and on-the-job learning, calibrated to match the observed distribution of firms and jobs across locations. In this first baseline model, there are no explicit density-based agglomeration externalities in learning or matching. The only reason the center is special is that it exogenously hosts a greater density of high-wage employers and that these employers are difficult to reach because of localized search and commuting frictions.

We show that this framework generates persistent earnings losses from displacement and accounts for a meaningful share of their magnitude—housing frictions slow re-sorting, commuting costs limit access to distant jobs, and localized search confines displaced workers to a thinner set of nearby vacancies. Yet the model under-predicts the empirical effects, suggesting that proximity to the center confers benefits beyond mere access to the existing

stock of good jobs. We confirm that our calibrated model under-predicts the wage growth in the center, suggesting that agglomeration externalities play a key role that the next iteration of the model should address.

6.1 Overview

We model a disk-shaped metropolitan area with a fixed exogenous population of size L. Time is continuous. Workers differ in age $a \in \mathcal{A} = \{1, \ldots, A\}$, human capital $h \in \mathcal{H} = \{h_1, \ldots, h_H\}$, and employment status $e \in \{E, U\}$. Workers transition to age $a+1$ at hazard $\psi(a) \geq 0$ and exit the economy at hazard $\bar{\psi}(a) > 0$. Every worker exiting the economy is replaced with an unemployed worker in the same location of the lowest age and lowest human capital. Search and matching frictions in the labor and housing markets imply that job offers and housing opportunities occur as Poisson shocks.

The metropolitan area is radially symmetric around a central business district (CBD), located at $y = 0$. Each housing unit accommodates one resident, so the population density equals housing density and integrates to the total population L. All plots located at the same distance x from the CBD are identical; we therefore index locations by their ring $x \in [0, \bar{D}]$ from the center, where $\bar{D}$ is the outer boundary of the metro area. We use y when referring to a ring as a workplace, x when a residence. People who live on ring x, work on ring y pay commuting time $\tau(x, y)$.

Production occurs at every ring $y \in [0, \bar{D}]$. There is a fixed stock $\bar{M}(y, z) > 0$ of potential jobs, $z \in \mathcal{Z} = \{1, \ldots, Z\}$. A potential job can be filled, actively vacant, or dormant. Hence, the masses of active vacancies and filled jobs of each quality z in each workplace y are endogenous. Employed workers transition to h' with hazard rate $q^E(h'|a, h, z)$; unemployed workers with rate $q^U(h'|a, h)$.

Housing units are owned by absentee landlords. A worker residing in a housing unit at ring x pays a rent $r(x)$ per unit of housing quality, which is measured in the same units as the numeraire consumption good. There are $\bar{H}(x)$ housing units per ring. The landlord supplies the chosen quality level by exerting effort proportional to the total rent minus the value of land. Hence, their net return per plot is $r(x)$, regardless of the tenant's choice of quality. We interpret $\bar{D}$ as the urban fringe and close the housing block with the boundary condition $r(\bar{D}) = r_A$, where r_A is the fringe rent.

In what follows, we describe the utility flows, production and wages, and search in the labor and housing markets, then formalize the worker's problem. The utility flow determines indirect flow utilities and housing demand in the metro area, while the dynamic problem governs the transitions across housing and labor markets and, consequently, the aggregate

evolution of the metro area. We conclude with a definition of the metro area's equilibrium.

6.2 Flow Utility

A worker occupies one housing unit at the chosen location and chooses consumption $c > 0$ of the numeraire good and housing quality $q > 0$ on that unit. For a worker of age a, residence x, workplace y, firm type z, and income $w(h, x, y, z) > 0$, the flow utility problem is

$$\begin{aligned} u^E(a, h, x, y, z) &= \max_{c,q} \{\alpha \log c + (1-\alpha) \log q\} + B(x) - \chi\tau(x, y) && (2) \\ &\text{s.t.} \quad c + r(x)q \leq w(h, x, y, z), && (3) \end{aligned}$$

where $\alpha \in (0, 1)$ is the expenditure share on the consumption good and $B(x)$ is the amenity term. The worker's flow problem has the Cobb–Douglas solution:

$$c^*(w, x) = \alpha w, \qquad q^*(w, x) = \frac{(1-\alpha)w}{r(x)}. \tag{4}$$

Hence richer households buy higher *quality* rather than more units. Substituting yields the indirect flow utility function

$$u^E(h, x, y, z) = \kappa_\alpha + \log w(h, x, y, z) - (1-\alpha) \log r(x) + B(x) - \chi\tau(x, y), \tag{5}$$

where $\kappa_\alpha = \alpha \log \alpha + (1-\alpha) \log(1-\alpha)$ is a constant. For an unemployed worker in state (a, h, x), income equals unemployment income $b(h) > 0$, so the indirect utility flow is

$$u^U(h, x) = \kappa_\alpha + \log b(h) - (1-\alpha) \log r(x) + B(x). \tag{6}$$

6.3 Production and Wages

A match between worker (h, x) and firm z at ring y yields output

$$Y(h, x, y, z) = A(z) \exp\left(\kappa_h h - \alpha_y y\right) \max\{\bar{n} - \tau(x, y), 0\}, \tag{7}$$

where $\max\{\bar{n} - \tau(x, y), 0\}$ is the amount of productive time left after commuting. This specification has two intended roles. First, long commutes lower the productive time available within a match. Second, firm type and workplace location shift productivity even after conditioning on worker characteristics. The first channel captures direct commuting losses. The second captures the fact that better firms are more concentrated in some parts of the city.

A landlord who owns one housing unit at x can supply quality q on that unit at resource cost $r(x)(q-1)$. Since the resident pays $r(x)q$, the landlord's net return from an occupied unit is exactly $r(x)$, independently of the chosen quality. This separates the intensive margin of housing quality from the extensive margin of occupied units: richer households spend more on housing because they buy higher quality, but they do not absorb multiple plots.

Wages are determined by a money-metric period-by-period Nash bargaining. The current-flow utility gain from employment at y relative to unemployment for a worker (h, x) is $\log w(h, x, y, z) - \log b(h) - \chi\tau(x, y)$. This flow gain is zero when the wage equals the worker's reservation wage in goods units,

$$w^R(h, x, y) = b(h) \exp(\chi\tau(x, y)) . \tag{8}$$

The bargaining rule splits *current-flow* surplus in money-metric units. The worker's money-metric surplus from a wage w is $w - w^R(h, x, y)$, and the firm's current-flow surplus is $Y(h, x, y, z) - w$. The wage solves

$$\max_{w \in [w^R(h,x,y),\, Y(h,x,y,z)]} \big(w - w^R(h, x, y)\big)^{\beta} \big(Y(h, x, y, z) - w\big)^{1-\beta}, \tag{9}$$

where $\beta \in (0, 1)$ is the worker's bargaining weight. The first-order condition gives the closed-form wage

$$w(h, x, y, z) = (1 - \beta)w^R(h, x, y) + \beta Y(h, x, y, z). \tag{10}$$

A match is statically feasible if and only if $Y(h, x, y, z) \geq w^R(h, x, y)$. If this fails, the candidate match never forms. Current profit in a filled match is therefore

$$\pi(h, x, y, z) = (1 - \beta)\big(Y(h, x, y, z) - w^R(h, x, y)\big). \tag{11}$$

This is positive exactly when the static viability condition holds. Because wages are state contingent and re-evaluated whenever worker or match characteristics change, realized productivity and human-capital changes are immediately reflected in pay. This wage rule abstracts from the effect of continuation-value differences on the current wage, which would arise under full dynamic Nash bargaining, but preserves a closed-form wage schedule under logarithmic worker utility.

6.4 Search and Matching in the Labor Market

6.4.1 Job-search kernels

Labor markets are a combination of a workplace y and a firm type z. Workers (employed or not) allocate their search differently depending on where they live and work.

Let $\omega^E(y' \mid x, y)$ denote the share of employed search from residence x and workplace y directed to workplace ring y', and let $\omega^U(y' \mid x)$ denote the corresponding share for unemployed search. We assume that search favors workplaces near the current residence and current workplace according to

$$\omega^E(y' \mid x, y) = \frac{\exp\big(-\eta_x \tau(x, y') - \eta_y^E \tau(y, y')\big)}{\int_{\tilde{y}} \exp\big(-\eta_x \tau(x, \tilde{y}) - \eta_y^E \tau(y, \tilde{y})\big) d\tilde{y}}. \tag{12}$$

Here $\eta_x \geq 0$ governs how job search decays with commuting cost, $\eta_y^E \geq 0$ governs persistence in the current workplace on the job ladder. For unemployed workers, $\omega^U(y' \mid x)$ follows the same expression but with $\eta_y^U = 0$.

6.4.2 Search mass, matching, and labor-market tightness

Matching is specific to each ring y and the job offer rates depend on the mass of vacancies $V(y)$ and aggregate intensity of search $S(y)$. Both $V(y)$ and $S(y)$ depend on the spatial distribution of workers. Let μ^E and μ^U denote the stationary cross-sectional distribution of employed and unemployed workers over states. Define employment at workplace ring y and firm type z as

$$L(y, z) = \int_x \sum_{a,h} \mu^E(a, h, x, y, z) dx. \tag{13}$$

Among the idle potential jobs $\bar{M}(y, z) - L(y, z)$, only a fraction $\zeta(y, z) \in [0, 1]$ are actively advertised. This fraction is determined by the vacancy activation problem defined below in Section 6.4.3. Then active vacancies are

$$V(y, z) = \zeta(y, z) \left[\bar{M}(y, z) - L(y, z)\right], \qquad V(y) = \sum_{z=1}^{Z} V(y, z). \tag{14}$$

Effective search directed toward workplace ring y is

$$S(y) = \int_{x,y'} \sum_{a,h,z} \lambda^E \omega^E(y \mid x, y') \mu^E(a, h, x, y', z) dx dy' + \int_x \sum_{a,h} \lambda^U \omega^U(y \mid x) \mu^U(a, h, x) dx, \tag{15}$$

where λ^E is the arrival rate of sending job applications for employed workers and λ^U is the rate for unemployed workers. At workplace ring y, the flow of contacts is

$$\mathcal{M}(y) = \chi_M S(y)^\gamma V(y)^{1-\gamma}, \qquad \gamma \in (0,1). \tag{16}$$

Hence the worker-side contact rate $\lambda(y)$ and the vacancy-side contact rates $\phi(y)$ are given by

$$\lambda(y) = \chi_M \left(\frac{V(y)}{S(y)}\right)^{1-\gamma}, \qquad \phi(y) = \chi_M \left(\frac{S(y)}{V(y)}\right)^{\gamma}, \tag{17}$$

with the usual interpretation that $\theta(y) = V(y)/S(y)$ is labor-market tightness at workplace ring y. Conditional on contacting a vacancy at y, the probability that the vacancy is of type z is

$$\nu(z \mid y) = V(y,z)/V(y). \tag{18}$$

6.4.3 Firm values and vacancy activation

In workplace ring y and firm type z, each filled match has value $J^F(a,h,x,y,z)$ and each active posting has value $J^V(y,z)$ to the firm. Let $J^0(y,z)$ denote the value of holding the corresponding job idle. Active postings incur a flow cost $c^V(y,z) > 0$.

Conditional on the firm receiving a job application at y, the probability that the contacting worker is an employed worker in state (a,h,x,y',z') is the share of search mass contributed by that worker state,

$$\pi^E(a,h,x,y',z';y) = \frac{\lambda^E \omega^E(y \mid x,y') \mu^E(a,h,x,y',z')}{S(y)}. \tag{19}$$

Let $p^E(a,h,x,y',z';y,z)$ denote the probability that a worker in state (a,h,x,y',z') accepts an offer from firm type z at ring y, which we derive in Section 6.6. Define analogously $\pi^U(a,h,x;y)$ the share of search mass contributed by unemployed workers in state (a,h,x) and $p^U(a,h,x;y,z)$ the probability they accept an offer. If that worker accepts, the vacancy becomes a filled match with firm value $J^F(a,h,x,y,z)$.

Vacancies are posted based on a smooth vacancy-activation rule. Each idle job draws an idiosyncratic extreme-value shock to the utility of being active versus dormant. The share of idle jobs that are actively advertised is then

$$\zeta(y,z) = \frac{\exp(J^V(y,z)/\sigma_V)}{1 + \exp(J^V(y,z)/\sigma_V)}, \qquad \sigma_V > 0. \tag{20}$$

6.5 Housing Market

Housing opportunities arrive at rate $\lambda^H > 0$. A worker currently living at x draws a candidate residence x' from the kernel $\Pi(x' \mid x)$ satisfying $\int_{x'} \Pi(x' \mid x) = 1$. A move from x to x' requires paying utility-equivalent moving cost $K(x, x') \geq 0$.

Let $\mathcal{R}(x) = \sum_{a,h} \mu^U(a, h, x) + \int_y \sum_{a,h,z} \mu^E(a, h, x, y, z)dy$ denote resident mass in ring x. Since each resident occupies exactly one housing unit, competitive housing-market clearing requires:

$$\mathcal{R}(x) = \bar{H}(x). \tag{21}$$

To close the housing market in levels, we use the urban-fringe boundary condition

$$r(\bar{D}) = r_A > 0. \tag{22}$$

This condition anchors the overall level of rents, while the cross-ring shape of $r(x)$ is pinned down endogenously by workers' location choices and ring-by-ring housing-market clearing.

6.6 Job and housing choices

Before deciding whether to accept job or housing offers, workers draw idiosyncratic preferences for these offers. This leads them to accept with probabilities that increase in the value of the offers, according to a simple logit rule.

Denote by $W(a, h, x, y, z)$ the value of employment to an employed worker, and by $U(a, h, x)$ that of an unemployed worker. For an employed worker in state (a, h, x, y, z) who receives job offer (y', z'), define the candidate-job value $\widetilde{W}(a, h, x, y', z'; y, z) = W(a, h, x, y', z')$ if $Y(h, x, y', z') \geq w^R(h, x, y')$, $-\infty$ otherwise. The probability of accepting the offer is

$$p^E(a, h, x, y', z'; y, z) = \frac{\exp(\widetilde{W}(a, h, x, y', z'; y, z)/\sigma_J)}{\exp(\widetilde{W}(a, h, x, y', z'; y, z)/\sigma_J) + \exp(W(a, h, x, y, z)/\sigma_J)}. \tag{23}$$

For an unemployed worker in state (a, h, x) offered a job (y, z), define the candidate-job value $\widetilde{W}(a, h, x, y, z) = W(a, h, x, y, z)$ if $Y(h, x, y, z) \geq w^R(h, x, y)$, $-\infty$ otherwise. The probability of acceptance is

$$p^U(a, h, x; y, z) = \frac{\exp(\widetilde{W}(a, h, x, y, z)/\sigma_J)}{\exp(\widetilde{W}(a, h, x, y, z)/\sigma_J) + \exp(U(a, h, x)/\sigma_J)}. \tag{24}$$

Housing choices are determined analogously. For an unemployed worker in state (a, h, x)

offered x', the probability of moving is

$$\xi^U(a,h,x;x') = \frac{\exp\big((U(a,h,x') - K(x,x'))/\sigma_H\big)}{\exp\big((U(a,h,x') - K(x,x'))/\sigma_H\big) + \exp(U(a,h,x)/\sigma_H)}. \tag{25}$$

For an employed worker (a,h,x,y,z), a housing offer x' creates an intermediate decision: after moving, the worker may either keep the current job or separate into unemployment. The post-move inclusive value of that job-retention decision is

$$R(a,h,x',y,z;x) = \sigma_J \log\left(\exp(W(a,h,x',y,z)/\sigma_J) + \exp(U(a,h,x')/\sigma_J)\right). \tag{26}$$

The probability of keeping the job after an accepted move is therefore

$$\kappa(a,h,x',y,z;x) = \frac{\exp\big(W(a,h,x',y,z)/\sigma_J\big)}{\exp\big(W(a,h,x',y,z)/\sigma_J\big) + \exp\big(U(a,h,x')/\sigma_J\big)}. \tag{27}$$

The probability of moving is therefore

$$\xi^E(a,h,x',y,z;x) = \frac{\exp\big((R(a,h,x',y,z;x) - K(x,x'))/\sigma_H\big)}{\exp\big((R(a,h,x',y,z;x) - K(x,x'))/\sigma_H\big) + \exp(W(a,h,x,y,z)/\sigma_H)}. \tag{28}$$

6.7 Worker and firm value functions

6.7.1 Worker Bellman equations

For an employed worker (a,h,x,y,z), let W denote the continuation value. Let a_+ denote the next age state, with exogenous transition hazard $\psi(a)$. The exogenous separation hazard is $\delta \geq 0$. The Bellman equation is

$$\begin{aligned}
\rho W = {} & u^E(a,h,x,y,z) + \psi(a)(W(a_+,h,x,y,z) - W) + \delta\big(U(a,h,x) - W\big) - \bar{\psi}(a)W \\
& + \sum_{h'} q^E(h'|a,h,z)\big(W(a,h',x,y,z) - W\big) \\
& + \lambda^H \int_{x'} \Pi(x' \mid x)\Big(\mathcal{S}\big(W, R(a,h,x',y,z;x) - K(x,x');\sigma_H\big) - W\Big)dx' \\
& + \lambda^E \int_{y'} \omega^E(y' \mid x,y)\lambda(y') \sum_{z'} \nu(z' \mid y')\Big(\mathcal{S}(W, \widetilde{W}(a,h,x,y',z';y,z);\sigma_J) - W\Big)dy',
\end{aligned} \tag{29}$$

where $\mathcal{S}(u, w; \sigma) = \sigma \log(\exp(u/\sigma) + \exp(w/\sigma))$. Let $U(a, h, x)$ denote the continuation value for an unemployed worker. The Bellman equation is

$$\begin{aligned}\rho U = {} & u^U(a, h, x) + \psi(a)\big(U(a_+, h, x) - U\big) + \sum_{h'} q^U(h'|a, h)\big(U(a, h', x) - U\big) - \bar{\psi}(a)U \\ & + \lambda^H \int_{x'} \Pi(x' \mid x)\big(\mathcal{S}(U, U(a, h, x') - K(x, x'); \sigma_H) - U\big)dx' \\ & + \lambda^U \int_{y'} \omega^U(y' \mid x)\, \lambda(y') \sum_{z'=1}^{Z} \nu(z' \mid y')\Big(\mathcal{S}\big(\widetilde{W}(a, h, x, y', z'), U; \sigma_J\big) - U\Big)dy'. \end{aligned} \tag{30}$$

The last term is the expected gain from job offers net of the value of staying unemployed.

6.7.2 Firm Bellman equations

The firm's value of a filled match in state (a, h, x, y, z) is

$$\begin{aligned}\rho J^F = {} & \pi(h, x, y, z) + \psi(a)\big(J^F(a_+, h, x, y, z) - J^F\big) + (\bar{\psi}(a) + \delta)\big(J^0(y, z) - J^F\big) \\ & + \sum_{h'} q^E(h'|a, h, z)\big(J^F(a, h', x, y, z) - J^F\big) \\ & + \lambda^H \int_{x'} \Pi(x' \mid x)\xi^E(a, h, x', y, z; x)\Big[\kappa(a, h, x', y, z; x)J^F(a, h, x', y, z) \\ & + (1 - \kappa(a, h, x', y, z; x))J^0(y, z) - J^F\Big]dx' \\ & + \lambda^E \int_{y'} \omega^E(y' \mid x, y)\, \lambda(y') \sum_{z'} \nu(z' \mid y')p^E(a, h, x, y', z'; y, z)\big(J^0(y, z) - J^F\big), \end{aligned} \tag{31}$$

where the final term captures worker quits to better outside jobs, and a destroyed match returns the job to the idle object $J^0(y, z)$, not directly to the active-posting object J^V.

The expected present value of keeping a job (y, z) actively advertised is therefore

$$\begin{aligned} & \rho J^V(y, z) \\ & = -c^V(y, z) + \phi(y) \int_{a,h,x} \Big(\pi^U(a, h, x; y)p^U(a, h, x; y, z) \\ & + \int_{y',z'} \pi^E(a, h, x, y', z'; y)p^E(a, h, x, y', z'; y, z)dy'dz'\Big)\big(J^F(a, h, x, y, z) - J^V(y, z)\big)dadhdx, \end{aligned} \tag{32}$$

The corresponding ex-ante idle-job value is the smooth envelope

$$J^0(y, z) = \sigma_V \log\big(1 + \exp(J^V(y, z)/\sigma_V)\big), \tag{33}$$

up to the usual Type-I extreme-value additive constant. As $\sigma_V \to 0$, $J^0(y,z)$ converges to $\max\{J^V(y,z), 0\}$. This specification has three advantages. It bounds vacancy mass by idle capacity, preserves a meaningful notion of market tightness, and avoids a difficult free-entry problem over the total number of firms. As $\sigma_V \to 0$, activation approaches a deterministic threshold at $J^V(y,z) = 0$.

Together, equations (31)–(33) determine the profitability of keeping a job advertised.

6.8 Stationary equilibrium

The model's dynamics are represented by a continuous-time generator $Q(\Xi)$ indexed by the equilibrium objects $\Xi = (r, U, W, J^F, J^V, \mu)$, where $\mu = (\mu^E, \mu^U)$. Each row of $Q(\Xi)$ collects transition hazards out of one worker state: human-capital transitions, aging, job applications followed by job acceptance probabilities, housing offers followed by housing acceptance probabilities, exogenous job separations, and exit followed by replacement through the entrant distribution.

A stationary cross-section is a probability vector μ such that

$$\mu Q(\Xi) = 0, \qquad \sum_{a,h,z} \left(\mu^U(a,h,x) + \int_{x,y} \mu^E(a,h,x,y,z) dx dy \right) = 1. \tag{34}$$

Definition 1 (Stationary equilibrium)**.** *A stationary equilibrium is a collection* $\Xi^* = (r^*, U^*, W^*, J^{F,*}, J^{V,*}, \mu^*)$ *such that i) within-period choices satisfy* (4)*; ii) wages satisfy* (10) *and infeasible candidate matches never form; iii) worker values satisfy* (29)–(30)*; iv) firm values satisfy* (31)–(33)*; v) vacancy activation satisfies* (20)*; vi) the stationary distribution satisfies* (34)*; vii) housing-market clearing holds ring by ring through* (21) *the rent schedule satisfies the urban-fringe boundary condition* (22)*; viii) aggregate objects are generated from* μ^* *by* (13)–(17)*.*

We now establish that our economy admits at least one stationary equilibrium, under typical regularity assumptions on the model's primitives.

Theorem 1 (Existence of stationary equilibrium)**.** *Under regularity conditions described in Appendix B, the economy admits at least one stationary equilibrium.*

Proof. See Appendix B. □

7 Empirical Application

7.1 Calibration

We calibrate the model to the San Francisco Bay Area, defined as the area within 100 km of the CBD.[6] Space is discretized into $I = 10$ concentric rings of width 10 km, centered at 5, 15, ..., 95 km. Workers are characterized by three age groups ($A = 3$; young, middle-aged, old), and two endogenous human-capital levels ($H = 2$). Firms are of two types ($Z = 2$; low- and high-productivity). Most calibration targets are obtained from publicly available data and smoothed using population-weighted local-linear kernel regressions with a Gaussian kernel and 10 km bandwidth, evaluated at each ring center. Parameter values are reported in Table 3 in Appendix C. Future iterations of the model will use GMM to target empirical moments.

Geography and housing supply. Ring-level population and land area are computed from census-tract centroids (Census Bureau Gazetteer) matched to the American Community Survey (ACS) 2015–2019 via IPUMS NHGIS. The housing-unit capacity $\bar{H}(x)$ in each ring is set equal to the observed population density times ring area, normalized so that $\sum_x \bar{H}(x) = 1$. The total population is $L = 7{,}579{,}110$, matching the 100 km area. Ring-level rents are taken from the Longitudinal Tract Data Base (LTDB) (two-bedroom median rent, 2012) and entered directly as targets; we set r_A to the average rent in the outermost ring.

Production and firm geography. Job density by workplace ring is constructed from LEHD LODES 2019 origin-destination files, aggregated from census-block level to 10 km rings using the Census block-coordinate crosswalk. The stock of potential jobs $\bar{M}(y, z)$ in each ring is set to 1.2 times observed job density (to allow for a 20% vacancy-slack margin), split between firm types using the LODES workplace earnings distribution: the share of high-type ($z = 2$) jobs in ring y is equated to the share of jobs with annual earnings above \$3,333/month ("high-pay" in LODES SE03 files). Baseline firm productivities are $A_1 = 4{,}800$ and $A_2 = 8{,}800$. The workplace-productivity gradient $\alpha_{A,z}$ is type-specific: $\alpha_{A,1} = 0.003$ and $\alpha_{A,2} = 0.020$, so that high-type firm productivity decays roughly seven times faster with distance from the CBD. The worker's bargaining weight is $\beta = 0.55$. Vacancy posting costs are uniform at $c^V = 192$. These production parameters are set by hand to match the observed wage and high-type-share gradients.

Labor-market flows and search. The matching function is Cobb–Douglas with elasticity $\gamma = 0.50$ and scale $\chi_M = 1.35$. Offer arrival rates are age-declining: $\lambda^U = (1.45, 1.15, 0.90)$ for unemployed workers and $\lambda^E = (0.42, 0.30, 0.20)$ for employed workers. Exogenous sepa-

[6] The CBD is located at 37.793°N, 122.408°W, corresponding to the Financial District. The model could be expanded to include areas farther than 100 km in a final outermost ring.

ration rates are also age-declining: $\delta = (0.090, 0.071, 0.058)$, with an additional penalty for high commuting cost ($\delta_\tau = 0.10$) and a discount for high-type firms ($\delta_z = 0.25$). These values are informed by JOLTS separation and hiring rates by age, applied uniformly across rings. Job-search kernels decay exponentially with commuting cost at rate $\eta_x = \eta_y^E = 1$.

Human capital and wage growth. Workers accumulate human capital on the job at age- and firm-type-specific rates. They transition from the low to the high human capital type with probabilities ranging from 0.025 (young, low-z) to 0.090 (young, high-z), declining with age. The return to human capital is firm-type-specific: $\kappa_h = (0.15, 0.30)$, so that human-capital gains are more valuable in high-type firms. Unemployment income is $b(h) \approx 12.5\%$ of a reference wage level, increasing mildly with age and skill. These parameters are set to approximate the age–earnings gradient observed in LODES wage-growth data between 2002 and 2019.

Commuting. Commuting costs between rings x and y follow a slightly convex specification, $\tau(x, y) = \tau_1 |d_x - d_y|^{\tau_2}$, with $\tau_1 = 0.0017$ and $\tau_2 = 1.10$. Commuting disutility enters both the flow utility cost (with age-specific weight $\chi_a = (0.30, 0.35, 0.40)$) and effective productive time. Commuting distance targets come from LODES 2019 origin-destination flows, where we compute average workplace–residence distance for each home ring using block-level coordinates.

Geographic mobility. Housing-market opportunities arrive at rate $\lambda^H = 0.17$. The housing-offer kernel $\Pi(x'|x)$ is proportional to $\bar{H}(x') \exp(-\eta_H |d_{x'} - d_x|)$ with $\eta_H = 0.095$, so that workers are more likely to receive offers from nearby, densely populated rings. Moving costs have a fixed component $K_0 = 0.20$ (in utility units) and a distance component $K_1 = 0.0025$ per km. The choice smoothing parameters are $\sigma_J = 0.55$ for job decisions and $\sigma_H = 0.50$ for housing moves. The resulting annual move rate of approximately 7.5% is compared to its target, constructed from ACS 2015–2019 same-MSA movers at the tract level (NHGIS), smoothed with the same kernel-regression specification.

Preferences. The consumption share is $\alpha = 0.72$ and the discount rate is $\rho = 0.05$. Amenities are modeled as a linear increasing function of distance: $B(x) = \gamma_A(d_x - 5)$ with $\gamma_A = 0.004$, so that peripheral locations offer a modest amenity premium (e.g., larger lots, less congestion) that partially offsets lower wages. Age- and attachment-specific amenity shifters are set to zero in the current calibration.

Model fit. Figure 9 plots the model's equilibrium profiles (solid lines) against the data targets (dots) for nine outcomes as a function of distance to the CBD. The model reproduces the key qualitative gradients: population and job density decline steeply with distance; the high-z job share falls from around 0.98 near the center to 0.26 in the outermost ring; wages decline by roughly 50% from center to periphery; rents are highest near the

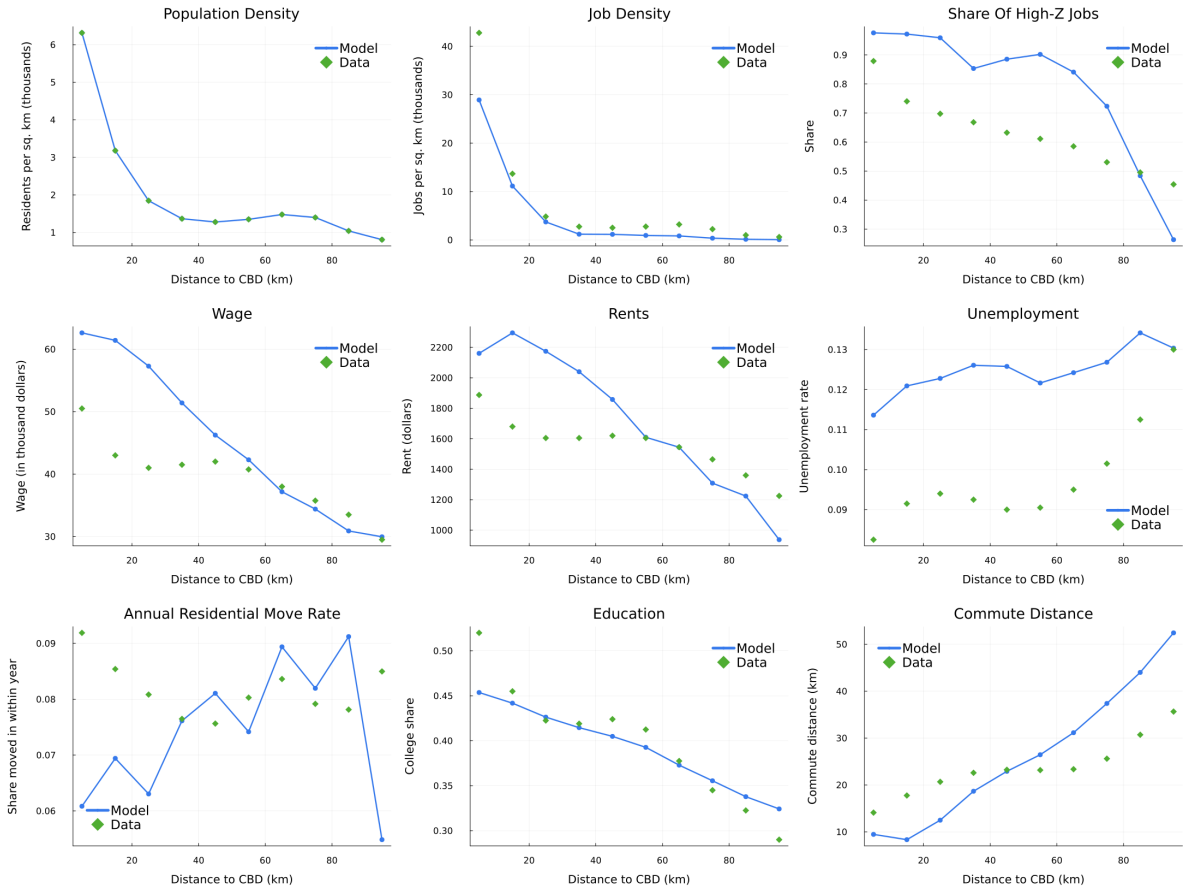


Figure 9: Model equilibrium profiles versus data targets, by distance to CBD

Notes: Each panel plots the model's stationary-equilibrium outcome (solid blue line) against the corresponding data target (green dots) as a function of distance from the San Francisco CBD. Data targets are constructed from LODES 2019, ACS/NHGIS 2015–2019, and LTDB sources, smoothed using population-weighted local-linear kernel regressions (bandwidth 10 km). The model is solved on a grid of $I = 10$ rings of width 10 km. See Appendix C for parameter values and data construction details.

CBD; and unemployment is slightly higher at the urban fringe. Commute distance rises monotonically from about 10 km to 50 km. The model generates a hump-shaped move-rate profile, qualitatively consistent with the data, though the fit at distances beyond 60 km is less precise. The college share is roughly flat across rings in the model, whereas the data show a mild decline—a limitation of the current two-skill specification. Overall, the model captures the broad spatial structure of the Bay Area economy while abstracting from finer geographic heterogeneity.

7.2 Simulating Evictions: Average Effects

We now use the calibrated model to simulate the effects of displacement from the city center in partial equilibrium. The exercise mirrors the empirical design of our Ellis Act event

studies: from the stationary equilibrium, we select a subset of agents residing in the most central ring and force them to vacate their housing unit at time t_0. These agents must accept the first housing offer they receive within two years of the eviction notice, reflecting the observed elevated move rates for two years after Ellis evictions. The simulated control group consists of agents who also reside in the city center at t_0 but are not subject to eviction. We then track both groups over time and compute differences in outcomes, analogous to the event-study coefficients estimated in Section 5.

Figure 10 presents the average effects of simulated evictions. In the model, evicted agents relocate on average 17 km from the city center, close to the 13 km average displacement distance observed empirically three years after an Ellis eviction. Rents paid by evicted agents drop by approximately 25 percent following the displacement, consistent with our estimates that evicted tenants move to areas with roughly 20 percent lower home values. Commuting distance increases by about 10 km, also in line with the empirical findings.

The earnings effects display the same gradual pattern documented in the data. Rather than an immediate drop, earnings of evicted agents decline progressively, reaching approximately \$6,000 below the control group after six years. This gradual deterioration arises because, even though displaced workers do not move far after the initial relocation, their earnings continue to fall as they transition to more local firms that are both less productive and offer slower human-capital accumulation. The model thus provides a structural rationale for the slow-building earnings losses observed in the Ellis Act event studies: displacement degrades workers' positions on the spatial job ladder, and the resulting career losses compound over time.

The magnitude of the simulated earnings decline, however, is roughly half the \$13,000 loss estimated empirically by year six. This gap suggests that the model, as currently calibrated, does not fully capture the forces that make the city center especially valuable for earnings growth. In particular, the baseline calibration features no direct advantage of central locations for human-capital accumulation beyond the higher concentration of high-productivity firms. Closing this gap likely requires introducing agglomeration forces that intensify earnings growth in the center—for example, a learning function $q^E(\cdot)$ that increases with local worker density $L(y)$, or a matching technology that generates higher contact rates $\lambda(y)$ in denser labor markets, in contrast to the current specification in which matching efficiency is location-invariant. We return to this point in Section 7.4 below.

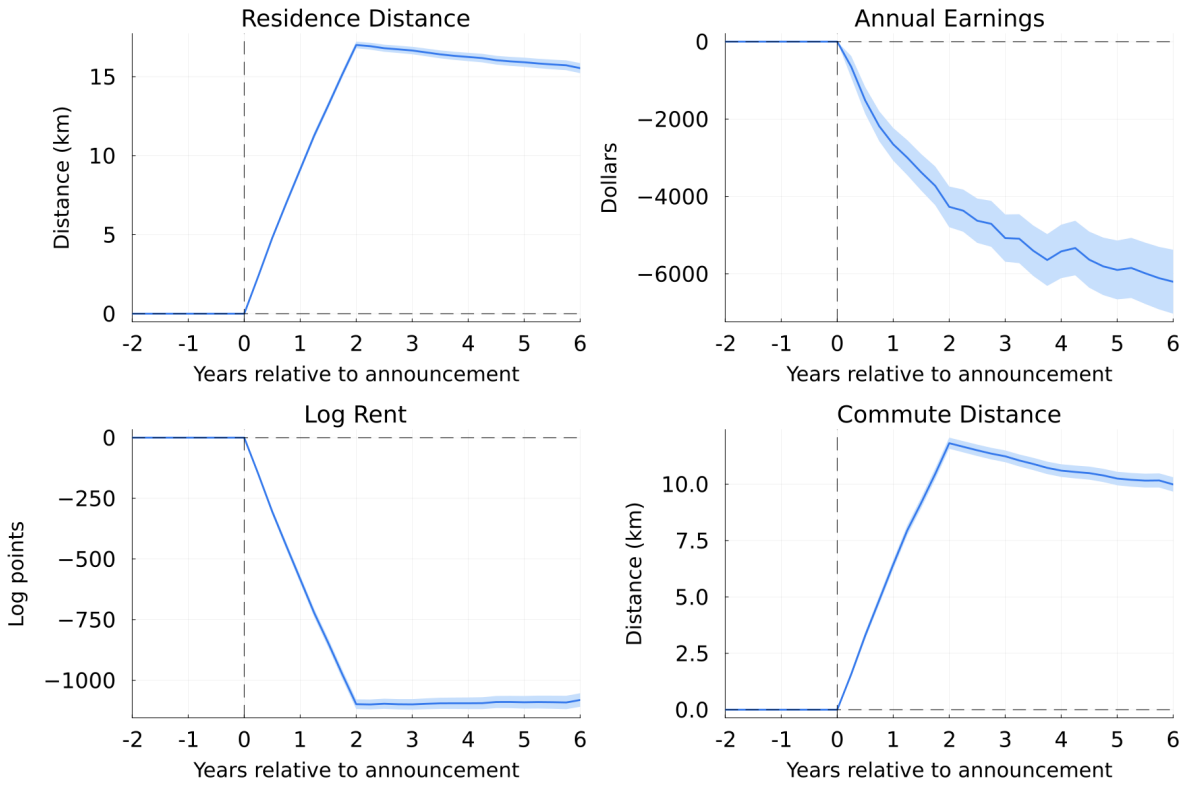


Figure 10: Simulated Evictions from the City Center: Average Effects

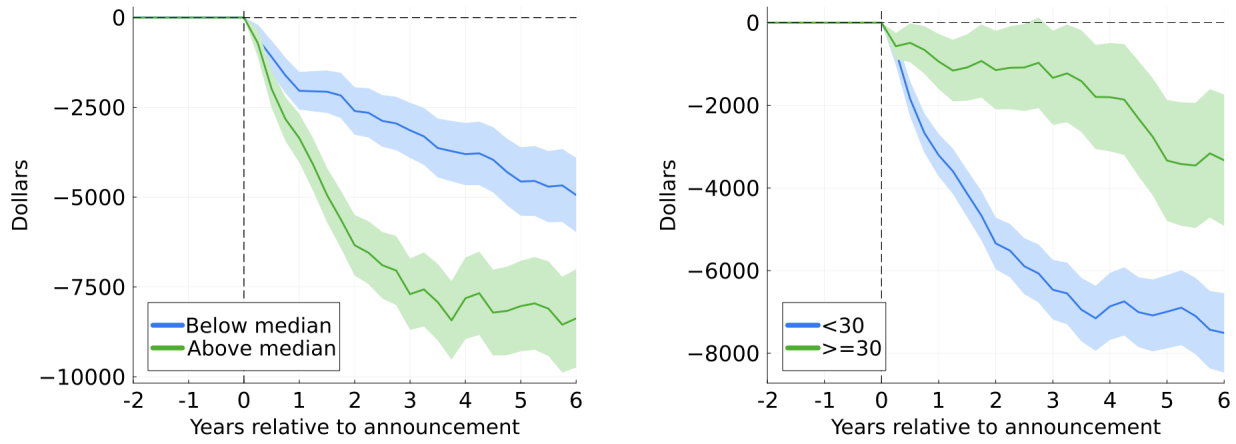
Notes: Each panel plots the results of the simulated event studies. From the stationary equilibrium of the model, 20,000 individuals are forced to leave their current housing unit and accept the first housing offer within two years. The simulated control group consists of agents in the city center at the time of the simulated eviction notice.

7.3 Simulating Evictions: Heterogeneity

The model is able to reproduce two key dimensions of heterogeneity documented in the empirical event studies. Figure 11 presents the simulated earnings effects separately by baseline income and by age.

Panel (a) shows that agents evicted when they have above-median income experience larger absolute earnings losses than below-median agents. This pattern is consistent with the structure of the model: higher-income workers are more likely to be employed at high-productivity firms concentrated near the center, and displacement pushes them toward peripheral locations where such firms are scarce. In relative terms, however, the earnings effects are proportionally larger for below-median income agents, mirroring the empirical finding that lower-income workers bear a disproportionate burden from displacement.

Panel (b) shows that younger agents are substantially more affected by eviction than older agents. In the model, younger workers benefit from faster human-capital accumulation and more frequent job-to-job transitions, both of which are more valuable when employed at



(a) By baseline income

(b) By age

Figure 11: Simulated Evictions from the City Center: Heterogeneity by Income and Age

Notes: Each panel plots the results of the simulated event studies for two subgroups. Panel (a) compares individuals evicted when their wage was below or above the median income in the economy. Panel (b) compares individuals evicted when they were more or less than 30 years old.

high-productivity firms near the center. Displacement interrupts this career-building process at a critical stage, generating larger and more persistent earnings losses for the young. This result confirms that the disruption of access to fast career progression in the city center is an important driver of the age heterogeneity observed in the Ellis Act event studies.

As with the average effects, the magnitude of the simulated earnings losses is smaller than the empirical estimates, reinforcing the conclusion that an additional amplification mechanism—such as density-dependent learning or matching—is needed to fully account for the observed effects of displacement.

7.4 Role of Commuting and Search Process

We conclude the simulation analysis by exploring how commuting costs and search frictions shape the magnitude of the earnings losses from displacement. Figure 12 compares the simulated earnings effects under three calibrations with increasing degrees of spatial friction.

Panel (b) reproduces the baseline earnings effects from Figure 10, with commuting disutility $\chi = 1$ and search decay parameter $\eta_x = 1$ in the search kernel defined in equation (12). In Panel (a), we decrease the commuting cost parameter to $\chi = 0.3$, making it less costly for displaced workers to commute back to the city center. We also decrease the search frictions to $\eta_x = 0.3$. This substantially reduces the effects of evictions: even if workers relocate away from the center, they can still search and commute to central jobs relatively easily. In Panel (c), we increase the commuting cost parameter to $\chi = 3$ and increase the search



(a) Low commuting cost ($\chi = 0.3$), Wide search ($\eta_x = 0.3$)

(b) Medium commuting cost ($\chi = 1$), Medium search ($\eta_x = 1$)

(c) High commuting cost ($\chi = 3$), Narrow search ($\eta_x = 3$)

Figure 12: Earnings Effects of Simulated Evictions: The Role of Commuting and Search

Notes: Each panel plots the earnings effects of the simulated event studies for three calibrations. Panel (a) considers a lower value of the commuting cost χ and a wider radius of search η_x, as defined in (12). Panel (b) is the baseline and replicates the earnings effects in Figure 10. Panel (c) considers both a higher value of the commuting cost χ and a narrower effective radius of search.

friction by setting $\eta_x = 3$, which narrows the effective geographic radius of job search. With both higher commuting costs and more localized search, earnings losses are amplified further, but only moderately, as displaced workers are increasingly confined to the lower-productivity firms available near their new, more peripheral residence.

These comparative statics confirm that commuting costs and search frictions are necessary ingredients for generating large earnings losses from city-center evictions: without spatial frictions, displaced workers could costlessly commute or search across the metropolitan area, and residential location would be irrelevant for labor market outcomes. However, even with substantially elevated frictions, the simulated earnings losses remain below the $13,000 average effect estimated in the Ellis Act event studies. This finding reinforces the conclusion that the model requires a stronger role for the city center in generating fast income growth—through agglomeration externalities in learning or matching—to fully replicate the empirical magnitude of displacement costs. Incorporating such forces is the focus of ongoing work.

8 Conclusion

This paper provides evidence that displacement from a high-productivity city center has persistent and substantial costs for workers' earnings trajectories. Using mass evictions under California's Ellis Act as a plausibly exogenous shock to residential location, we find that displaced tenants in San Francisco experience earnings losses averaging 20 percent of baseline over six years, accompanied by large and immediate declines in housing and

neighborhood quality. The earnings effects are largest for younger and lower-income workers and are concentrated among those who leave the urban core, even if they remain in the broader region—a pattern consistent with displacement degrading access to San Francisco's dense, high-wage labor market.

To interpret these findings, we develop and calibrate a spatial equilibrium model with frictional housing and labor markets, costly commuting, and on-the-job learning. Even without explicit agglomeration externalities, the model qualitatively replicates our main results: simulated evictions generate gradual earnings losses for workers who relocate 10–20 kilometers from the core, by removing them from immediate access to the most productive firms and disrupting their progression up the spatial job ladder. Fully matching the magnitude of the empirical effects requires a larger wage-growth premium from city-center proximity, which agglomeration forces—faster learning, better matching, or quicker transitions to more productive firms in denser markets—would naturally deliver. This is the clearest next step for developing the model.

Our findings underscore that housing stability is a form of labor market access. In high-productivity cities, disruptions to housing compound into long-run disruptions to earnings, especially at early career stages. The model clarifies the mechanism: displacement does not cause an immediate earnings shock but diverts workers from the career-building opportunities that dense urban labor markets uniquely provide. The heterogeneous effects further suggest that displacement widens inequality in access to urban opportunity, as younger and lower-income workers bear disproportionate costs.

More broadly, our results point to a deep interdependence between housing and labor markets in productive cities. Understanding the full welfare effects of housing instability—and designing policies to address it—requires accounting for the spatial job ladder that links where workers live to the evolution of their careers.

References

Abramson, Boaz, "The Welfare Effects of Eviction and Homelessness Policies," 2021. Working paper.

_ **and Tim Landvoigt**, "Curbing Rising Housing Costs: A Model-Based Policy Comparison," *Journal of Economic Perspectives*, 2025, *39* (3), 27–44.

Andersson, Fredrik, John C. Haltiwanger, Mark J. Kutzbach, Henry O. Pollakowski, and Daniel H. Weinberg, "Job Displacement and the Duration of Joblessness: The Role of Spatial Mismatch," *The Review of Economics and Statistics*, 2018, *100*, 203–218.

Asquith, Brian, "Do Rent Increases Reduce the Housing Supply Under Rent Control? Evidence from San Francisco," 2019. Upjohn Working Paper #19-296.

Autor, David H., Christopher J. Palmer, and Parag A. Pathak, "Housing Market Spillovers: Evidence from the End of Rent Control in Cambridge, Massachusetts," *Journal of Political Economy*, 2014, *122* (3), 661–717.

Baum-Snow, Nathaniel and Ronni Pavan, "Understanding the City Size Wage Gap," *Review of Economic Studies*, 2012, *79* (1), 88–127.

Bayer, Patrick, Stephen L. Ross, and Giorgio Topa, "Place of Work and Place of Residence: Informal Hiring Networks and Labor Market Outcomes," *Journal of Political Economy*, 2008, *116* (6), 1150–1196.

Bergman, Peter, Raj Chetty, Stefanie DeLuca, Nathaniel Hendren, Lawrence F. Katz, and Christopher Palmer, "Creating Moves to Opportunity: Experimental Evidence on Barriers to Neighborhood Choice," *American Economic Review*, 2024, *114* (5), 1281–1337.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess, "Revisiting Event-Study Designs: Robust and Efficient Estimation," *Review of Economic Studies*, 2024, *91* (6), 3253–3285.

Cerqueiro, Geraldo, Isaac Hacamo, and Pedro Raposo, "Priced-out: Rent Control, Wages, and Inequality," Technical Report 2025. Working paper.

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

_ , _ , **and Lawrence F. Katz**, "The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment," *American Economic Review*, 2016, *106* (4), 855–902.

Ciccone, Antonio and Robert E. Hall, "Productivity and the Density of Economic Activity," *American Economic Review*, 1996, *86* (1), 54–70.

City and County of San Francisco, "Assessor Historical Secured Property Tax Rolls," San Francisco Open Data Portal 2024. Dataset, accessed February 2, 2023.

Collinson, Robert, John Eric Humphries, Nicholas Mader, Davin Reed, Daniel Tannenbaum, and Winnie van Dijk, "Eviction and Poverty in American Cities," *Quarterly Journal of Economics*, 2024, *139* (1), 57–120.

Combes, Pierre-Philippe, Gilles Duranton, and Laurent Gobillon, "Spatial Wage Disparities: Sorting Matters!," *Journal of Urban Economics*, 2008, *63* (2), 723–742.

_ , _ , _ , **Diego Puga, and Sébastien Roux**, "The Productivity Advantages of Large Cities: Distinguishing Agglomeration from Firm Selection," *Econometrica*, 2012, *80* (6), 2543–2594.

Couch, Kenneth A. and Dana W. Placzek, "Earnings Losses of Displaced Workers Revisited," *American Economic Review*, 2010, *100*, 572–589.

Davis, Steven J. and Till M. von Wachter, "Recessions and the Costs of Job Loss," *Brookings Papers on Economic Activity*, 2011, *43* (2 (Fall)), 1–72.

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

De La Roca, Jorge and Diego Puga, "Learning by Working in Big Cities," *Review of Economic Studies*, 2017, *84* (1), 106–142.

Deryugina, Tatyana, Laura Kawano, and Steven Levitt, "The Economic Impact of Hurricane Katrina on Its Victims: Evidence from Individual Tax Returns," *American Economic Journal: Applied Economics*, April 2018, *10* (2), 202–33.

Diamond, Rebecca, "The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980-2000," 2016, *106*, 479–524.

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

Duranton, Gilles and Diego Puga, "Micro-Foundations of Urban Agglomeration Economies," in J. Vernon Henderson and Jacques-François Thisse, eds., *Handbook of Regional and Urban Economics*, Vol. 4, Elsevier, 2004, pp. 2063–2117.

Fallick, Bruce, John Haltiwanger, Erika McEntarfer, and Matthew Staiger, "Job Displacement and Earnings Losses: The Role of Joblessness," *American Economic Journal: Macroeconomics*, 2025, *17*, 177–205.

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

Glaeser, Edward L. and David C. Maré, "Cities and Skills," *Journal of Labor Economics*, 2001, *19* (2), 316–342.

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

Gobillon, Laurent, Harris Selod, and Yves Zenou, "The Mechanisms of Spatial Mismatch," *Urban Studies*, 2007, *44* (12), 2401–2427.

Gould, Eric D., "Cities, Workers, and Wages: A Structural Analysis of the Urban Wage Premium," *Review of Economic Studies*, 2007, *74* (2), 477–506.

Gutiérrez-i-Puigarnau, Eva and Jos N. van Ommeren, "Labour Supply and Commuting," *Journal of Urban Economics*, 2010, *68* (1), 82–89.

Gyourko, Joseph, Christopher Mayer, and Todd Sinai, "Superstar Cities," *American Economic Journal: Economic Policy*, 2013, *5* (4), 167–199.

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

Imbens, Guido W. and Donald B. Rubin, *Causal Inference for Statistics, Social, and Biomedical Sciences*, Cambridge University Press, 2015.

_ **and Jeffrey M. Wooldridge**, "Recent Developments in the Econometrics of Program Evaluation," *Journal of Economic Literature*, 2009, *47* (1), 5–86.

Jacobson, Louis S., Robert J. LaLonde, and Daniel G. Sullivan, "Earnings Losses of Displaced Workers," *American Economic Review*, 1993, (4), 685–709.

Kain, John F., "Housing Segregation, Negro Employment, and Metropolitan Decentralization," *Quarterly Journal of Economics*, 1968, *82* (2), 175–197.

Lachowska, Marta, Alexandre Mas, and Stephen A. Woodbury, "Sources of Displaced Workers' Long-Term Earnings Losses," *American Economic Review*, 2020, *110* (10), 3231–3266.

Le Barbanchon, Thomas, Roland Rathelot, and Alexandra Roulet, "Gender Differences in Job Search: Trading off Commute Against Wage," *Quarterly Journal of Economics*, 2021, *136* (1), 381–426.

Logan, John R., Zengwang Xu, and Brian Stults, "Interpolating U.S. Decennial Census Tract Data from 1970 to 2010: A Longitudinal Tract Database," 2014. Technical report, Brown University.

Manning, Alan and Barbara Petrongolo, "How Local Are Labor Markets? Evidence from a Spatial Job Search Model," *American Economic Review*, 2017, *107* (10), 2877–2907.

Marinescu, Ioana and Roland Rathelot, "Mismatch Unemployment and the Geography of Job Search," *American Economic Journal: Macroeconomics*, 2018, *10* (3), 42–70.

Martellini, Paolo, "Local Labor Markets and Aggregate Productivity," 2022. Working paper.

Monte, Ferdinando, Stephen J. Redding, and Esteban Rossi-Hansberg, "Commuting, Migration, and Local Employment Elasticities," *American Economic Review*, 2018, *108* (12), 3855–3890.

Moretti, Enrico, "Workers' Education, Spillovers, and Productivity: Evidence from Plant-Level Production Functions," *American Economic Review*, 2004, *94* (3), 656–690.

_ **and Moises Yi**, "Size Matters: Matching Externalities and the Advantages of Large Labor Markets," Technical Report 32250, National Bureau of Economic Research 2025.

Nakamura, Emi, Josef Sigurdsson, and Jon Steinsson, "The Gift of Moving: Intergenerational Consequences of a Mobility Shock," *The Review of Economic Studies*, 09 2021, *89* (3), 1557–1592.

Nix, Emily, "Learning spillovers in the firm," Working Paper 2020:14 2020.

Opportunity Insights, "Opportunity Insights Data Library," https://opportunityinsights.org/data/ 2024. Accessed: 2026-02-24.

Pennington, Kate, "Does Building New Housing Cause Displacement?: The Supply and Demand Effects of Construction in San Francisco," 2021.

Roback, Jennifer, "Wages, Rents, and the Quality of Life," *Journal of Political Economy*, 1982, *90* (6), 1257–1278.

Rosen, Sherwin, "Wage-Based Indexes of Urban Quality of Life," *Current Issues in Urban Economics*, 1979, pp. 74–104.

Saiz, Albert, "The Geographic Determinants of Housing Supply," *Quarterly Journal of Economics*, 2010, *125* (3), 1253–1296.

Schmutz, Benoît and Modibo Sidibé, "Frictional Labour Mobility," *The Review of Economic Studies*, 09 2018, *86* (4), 1779–1826.

Stevens, Ann Huff, "Persistent Effects of Job Displacement: The Importance of Multiple Job Losses," *Journal of Labor Economics*, 1997, *15*, 165–188.

Sun, Liyang and Sarah Abraham, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

U.S. Census Bureau, "Center for Administrative Records Research and Applications (CARRA) Databank," Restricted-access administrative data 2024. Confidential data accessed under approved Census Bureau project.

Waldinger, Daniel, "Targeting In-Kind Transfers through Market Design: A Revealed Preference Analysis of Public Housing Allocation," *American Economic Review*, 2021, *111* (8), 2660–2696.

A Policy Appendix: Rent Control and the Ellis Act

A.1 Origins and Purpose of the Ellis Act

The Ellis Act (California Government Code §§7060–7060.7) was enacted in 1985 following the California Supreme Court's decision in *Nash v. City of Santa Monica* (1984), which upheld a municipality's right to prevent a landlord from withdrawing rent-controlled property from the rental market. The legislature established the principle that property owners cannot be compelled to continue offering accommodations for rent, subject to certain conditions and protections for tenants.

A.2 Key Provisions

The Act requires that all units in a building be withdrawn simultaneously (Section 7060.7(3)). Tenants receive 120 days' notice (extended to one year for elderly (62+) and disabled tenants since January 2000 under SB 948). Prior to then, all tenants were given 60 days' notice. Landlords must pay relocation benefits, which have been expanded over time:

- **1986–1994:** Relocation payments based on unit size (studios: $1,500; one-bedroom: $1,750; two or more bedrooms: $2,500), plus $3,000 for elderly/disabled tenants.
- **1994–2004:** Following the *Channing* decision (1992), payments were restricted to low-income tenants, set at $4,500 per tenant starting in 2000. $3,000 surcharge for elderly/disabled tenants retained. Income cutoffs followed HUD's qualifying definition for "low-income" for San Francisco.
- **2004–present:** Ordinance 21-05 (effective retroactively from August 2004) extended payments to all tenants at $4,500 per person (maximum $13,500 per unit). $3,000 surcharge for elderly/disabled tenants retained. These amounts are adjusted annually for inflation.

A.3 Re-Rental Restrictions

If the building is offered for rent within two years of withdrawal, the landlord is liable for damages and must offer units back to displaced tenants. For five years after withdrawal, any re-rental must be at the prior lawful rent plus allowable annual adjustments. San Francisco extends a right of return for up to ten years. If the building is demolished and rebuilt within five years, replacement units are treated as rent-controlled.

A.4 San Francisco Rent Stabilization Ordinance

San Francisco's Rent Stabilization Ordinance (enacted June 13, 1979) covers buildings with two or more units built before its effective date. Annual rent increases are capped at 60 percent of the regional CPI (with a 7 percent ceiling), averaging approximately 1.7 percent during our study period. Evictions can only be for "just cause" and are defined by statute under the ordinance. The Costa-Hawkins Act (1995) mandates vacancy decontrol statewide, exempts new construction, and permanently bars rent control on single-family homes, condominiums, and cooperatives.

A.5 Condo Conversion Restrictions

Conversions to condominium have been strictly regulated in San Francisco since 1985, capped at 200 units per year for small buildings and allocated by lottery (suspended indefinitely in 2013). Beginning in 2005–2006, San Francisco penalized Ellis Act properties by reducing their lottery chances or banning them from the lottery outright for 10 years if an Ellis eviction had been performed. These restrictions effectively channel most post-Ellis conversions into tenancy-in-common arrangements.

B Existence of Equilibrium

Let

$$S_U = \mathcal{A} \times \mathcal{H} \times X, \qquad S_E = \mathcal{A} \times \mathcal{H} \times X \times Y \times \mathcal{Z}, \qquad S = S_U \sqcup S_E.$$

A stationary cross-section is represented by nonnegative continuous densities $\mu = (\mu^U, \mu^E)$, with $\mu^U \in C(S_U)$ and $\mu^E \in C(S_E)$, normalized so that

$$\sum_{a,h} \int_X \mu^U(a,h,x)\,dx + \sum_{a,h,z} \int_{X\times Y} \mu^E(a,h,x,y,z)\,dx\,dy = 1.$$

The implied resident density is

$$R_\mu(x) = \sum_{a,h} \mu^U(a,h,x) + \sum_{a,h,z} \int_Y \mu^E(a,h,x,y,z)\,dy.$$

Housing-market clearing is the pointwise identity $R_\mu(x) = \bar{H}(x)$, and we normalize the housing stock so that

$$\int_X \bar{H}(x)\,dx = 1.$$

The rent at the urban fringe is taken as given:

$$r(\bar{D}) = r_A > 0.$$

Assumption 1 (Regularity of primitives)**.** *The sets $\mathcal{A}$, $\mathcal{H}$, and $\mathcal{Z}$ are finite, and $X = Y = [0, \bar{D}]$ are compact intervals. The functions $\bar{H}$, $\bar{M}$, B, b, τ, K, and c^V, together with the densities of $\Pi(x' \mid x)$, $\omega^E(y' \mid x, y)$, and $\omega^U(y' \mid x)$, are bounded and Lipschitz on their domains. There exist strictly positive constants $\underline{H}$, $\underline{M}$, $\underline{b}$, $\underline{\Pi}$, $\bar{\Pi}$, $\underline{\omega}$, and $\bar{\omega}$ such that*

$$\bar{H}(x) \geq \underline{H}, \qquad \bar{M}(y, z) \geq \underline{M}, \qquad b(h) \geq \underline{b},$$

and

$$\underline{\Pi} \leq \Pi(x' \mid x) \leq \bar{\Pi}, \qquad \underline{\omega} \leq \omega^k(y' \mid \cdot) \leq \bar{\omega}, \qquad k \in \{E, U\},$$

for all admissible arguments. The smoothing parameters satisfy $\sigma_H > 0$, $\sigma_J > 0$, and $\sigma_V > 0$. The total hazard generated by aging, exit, separation, human-capital transitions, housing offers, and job contacts is bounded above by some $\bar{\Lambda} < \infty$, and the housing-offer hazard is bounded below by a strictly positive constant. Finally, there exists $\varepsilon_\pi > 0$ such that

$$\left|Y(h, x, y, z) - w^R(h, x, y)\right| \geq \varepsilon_\pi \qquad \text{for all } (h, x, y, z) \in \mathcal{H} \times X \times Y \times \mathcal{Z}.$$

Theorem 2 (Existence of stationary equilibrium)**.** *Under Assumption 1, the economy admits at least one stationary equilibrium.*

Proof. The equilibrium can be summarized by three aggregate objects: a rent schedule r, a vacancy-activation rule ζ, and a stationary worker density μ. Once such a triple is given, the aggregate variables that enter individual problems are already pinned down. The density μ determines resident density and search mass; the pair (ζ, μ) determines active vacancies, tightness, contact rates, and vacancy shares; and those aggregate objects determine the worker and firm Bellman equations. The Bellman solutions then imply an updated vacancy rule and an updated law of motion for workers. The proof uses these observations to construct a fixed-point problem on a compact set $\mathcal{K} = \mathcal{R} \times \mathcal{Z} \times \mathcal{M}$. The only remaining equilibrium restriction is housing-market clearing. To guarantee it, we define a rent correspondence that selects, from a compact set of admissible rent schedules satisfying $r(\bar{D}) = r_A$, the schedules that maximize the inner product with excess housing demand. If a fixed point of that correspondence failed to clear the housing market, a small tilt of rents toward rings with excess demand would improve the objective, which is impossible at a maximizer. The rest of the proof makes that argument precise.

Consider the compact convex sets

$$\mathcal{R} = \Big\{ r \in C(X) : \; r(\bar{D}) = r_A, \; \underline{r} \le r(x) \le \bar{r} \text{ for all } x \in X, \; \mathrm{Lip}(r) \le L_r \Big\},$$

$$\mathcal{Z} = \Big\{ \zeta \in C(Y \times \mathcal{Z}) : \; \underline{\zeta} \le \zeta(y,z) \le 1 - \underline{\zeta}, \; \mathrm{Lip}(\zeta) \le L_\zeta \Big\},$$

$$\mathcal{M} = \Big\{ (\mu^U, \mu^E) : \; \mu^U, \mu^E \ge 0, \; \|\mu^U\|_\infty + \|\mu^E\|_\infty \le M_\mu, \; \mathrm{Lip}(\mu^U), \mathrm{Lip}(\mu^E) \le L_\mu, \; \mu \text{ is normalized} \Big\}.$$

The constants in these definitions are chosen below from primitive bounds. Once they are fixed, each set is nonempty and convex. Moreover, since the underlying domains are compact and the admissible functions are uniformly bounded and Lipschitz, each set is compact in the sup norm. Therefore, $\mathcal{K}$ is compact and convex as well.

Fix any $(r, \zeta, \mu) \in \mathcal{K}$. The employment density

$$L_\mu(y,z) = \int_X \sum_{a,h} \mu^E(a,h,x,y,z)\, dx$$

is continuous, and so are active vacancies $V_{\zeta,\mu}(y,z) = \zeta(y,z)(\bar{M}(y,z) - L_\mu(y,z))$, total vacancies, search mass, tightness, contact rates, and vacancy shares. Because r is bounded away from zero and above on $\mathcal{R}$, the worker flow utilities and firm flow profits are uniformly bounded. The worker Bellman system is therefore well defined for every $(r, \zeta, \mu) \in \mathcal{K}$. Uniformizing at rate $\bar{\Lambda}$ from Assumption 1, the worker Bellman operator maps $C(S_U) \times C(S_E)$ into itself and has contraction modulus $\bar{\Lambda}/(\rho + \bar{\Lambda}) < 1$. The only nonlinear terms are inclusive values of the form

$$\mathcal{S}(u, v; \sigma) = \sigma \log\big(e^{u/\sigma} + e^{v/\sigma}\big),$$

and $\mathcal{S}$ is 1-Lipschitz in (u, v) under the sup norm. Hence there is a unique pair $(U_{r,\zeta,\mu}, W_{r,\zeta,\mu})$. Holding those worker values fixed, the firm Bellman system is a contraction on $C(S_E) \times C(Y \times \mathcal{Z})$, so there is a unique pair $(J^F_{r,\zeta,\mu}, J^V_{r,\zeta,\mu})$. Standard parametric-contraction arguments imply that all four value functions depend continuously on (r, ζ, μ). The same arguments on the space of Lipschitz functions yield primitive constants C_V and L_V such that, uniformly over $\mathcal{K}$,

$$\|U_{r,\zeta,\mu}\|_\infty + \|W_{r,\zeta,\mu}\|_\infty + \|J^F_{r,\zeta,\mu}\|_\infty + \|J^V_{r,\zeta,\mu}\|_\infty \le C_V,$$

and the corresponding Lipschitz seminorms are bounded by L_V.

These value functions determine the model-implied vacancy rule

$$\widehat{\zeta}_{r,\zeta,\mu}(y,z) = \frac{\exp\big(J^V_{r,\zeta,\mu}(y,z)/\sigma_V\big)}{1 + \exp\big(J^V_{r,\zeta,\mu}(y,z)/\sigma_V\big)}.$$

Because $J^V_{r,\zeta,\mu}$ is uniformly bounded, $\widehat{\zeta}_{r,\zeta,\mu}$ is bounded away from 0 and 1 by a primitive constant. Because $J^V_{r,\zeta,\mu}$ is uniformly Lipschitz, so is $\widehat{\zeta}_{r,\zeta,\mu}$. Choosing $\underline{\zeta}$ small enough and L_ζ large enough therefore ensures that

$$\widehat{\zeta}_{r,\zeta,\mu} \in \mathcal{Z} \qquad \text{for every } (r,\zeta,\mu) \in \mathcal{K}.$$

The same candidate triple and the associated value functions determine the one-step law of motion for workers. Let $Q_{r,\widehat{\zeta},\mu}$ be the continuous-time generator implied by the resulting acceptance probabilities and transition hazards, and define the uniformized kernel

$$P_{r,\widehat{\zeta},\mu} = I + \bar{\Lambda}^{-1} Q_{r,\widehat{\zeta},\mu}.$$

The distribution obtained after one uniformized step is

$$\widehat{\mu}_{r,\zeta,\mu} = \mu P_{r,\widehat{\zeta},\mu}.$$

Because the transition kernels admit bounded densities, this update can be written pointwise as

$$\widehat{\mu}_{r,\zeta,\mu}(s) = a_{r,\widehat{\zeta},\mu}(s)\mu(s) + \int_S g_{r,\widehat{\zeta},\mu}(s',s)\mu(s')\,ds',$$

where the no-jump weight satisfies $0 \le a_{r,\widehat{\zeta},\mu}(s) \le q_0 < 1$ uniformly, because the total jump hazard is bounded below by the strictly positive housing-offer hazard, and where $g_{r,\widehat{\zeta},\mu}$ is jointly continuous, uniformly bounded, and uniformly Lipschitz in the arrival state. It follows that

$$\|\widehat{\mu}_{r,\zeta,\mu}\|_\infty \le q_0\|\mu\|_\infty + C_0$$

and

$$\operatorname{Lip}(\widehat{\mu}_{r,\zeta,\mu}) \le q_0 \operatorname{Lip}(\mu) + C_1\|\mu\|_\infty + C_2$$

for primitive constants C_0, C_1, and C_2. Choosing M_μ and L_μ above the corresponding fixed points therefore makes $\mathcal{M}$ invariant under the density update. Moreover, because the housing-offer kernel has full support and housing acceptance probabilities are bounded away from zero by $\sigma_H > 0$, the boundedness of values, and the boundedness of moving costs, there exists a primitive constant $c_- > 0$ such that

$$R_{\widehat{\mu}_{r,\zeta,\mu}}(x) \ge c_- \qquad \text{for every } x \in X \text{ and every } (r,\zeta,\mu) \in \mathcal{K}.$$

This lower bound will be used to control supporting rent schedules.

The boundary condition $r(\bar{D}) = r_A$ pins down the level of rents, so it remains to show that

stationary supporting schedules stay away from arbitrarily large or arbitrarily small levels. At this point we use the candidate Lipschitz restriction built into $\mathcal{R}$. Fix $\eta_A = r_A/(4L_r)$ and let $N_A = [\bar{D} - \eta_A, \bar{D}]$. For every $r \in \mathcal{R}$ and every $x' \in N_A$,

$$|r(x') - r_A| \leq L_r |x' - \bar{D}| \leq r_A/4,$$

so rents on N_A remain in the fixed interval $[3r_A/4, 5r_A/4]$. Suppose now that (r, ζ, μ) satisfies $\mu = \widehat{\mu}_{r,\zeta,\mu}$. Then every ring carries resident mass at least $c_- > 0$. Because the non-rent part of residence utility and all continuation values are uniformly bounded, there exists $\bar{r}^0 < \infty$ such that, whenever $r(x) > \bar{r}^0$, every worker located at x strictly prefers with a uniform margin any offer into N_A to staying at x, while workers in N_A reject moves into x with probability uniformly bounded away from one. Since $\Pi(x' \mid x)$ and $\Pi(x \mid x')$ are bounded below on sets of positive measure and $|N_A| = \eta_A > 0$, this would generate a strictly positive net outflow from ring x, contradicting stationarity of resident mass at x. The same argument, with the direction of the comparison reversed, yields a constant $\underline{r}^0 > 0$ such that no stationary supporting schedule can satisfy $r(x) < \underline{r}^0$. Hence every stationary supporting rent schedule in $\mathcal{R}$ satisfies

$$\underline{r}^0 \leq r(x) \leq \bar{r}^0 \qquad \text{for all } x \in X.$$

Once these level bounds are available, the unemployed Bellman equation can be rearranged for any fixed reference state $(a_0, h_0) \in \mathcal{A} \times \mathcal{H}$ as

$$(1 - \alpha) \log r(x) = \kappa_\alpha + \log b(h_0) + B(x) + \Psi_{a_0,h_0}(x; U, W) - \rho U(a_0, h_0, x),$$

where $\Psi_{a_0,h_0}(x; U, W)$ collects the continuation terms. The right-hand side is uniformly Lipschitz because U, W, and the primitives are, so every stationary supporting rent schedule has $\log r$ uniformly Lipschitz. Since such schedules are also bounded away from zero and above, the rent schedule itself is uniformly Lipschitz. Thus there exists a primitive constant L_r^0 such that $\mathrm{Lip}(r) \leq L_r^0$ for every stationary supporting rent schedule. Choosing

$$\underline{r} < \underline{r}^0, \qquad \bar{r} > \bar{r}^0, \qquad L_r > L_r^0,$$

makes $\mathcal{R}$ a compact set that contains every stationary supporting rent schedule satisfying $r(\bar{D}) = r_A$ in its interior.

The aggregate update can now be defined on $\mathcal{K}$. For each $(r, \zeta, \mu) \in \mathcal{K}$, let $\widehat{\zeta}_{r,\zeta,\mu} \in \mathcal{Z}$ and $\widehat{\mu}_{r,\zeta,\mu} \in \mathcal{M}$ be the objects just constructed, and define excess housing demand generated by

the updated cross-section as

$$e_{r,\zeta,\mu}(x) = R_{\widehat{\mu}_{r,\zeta,\mu}}(x) - \bar{H}(x).$$

Because $\widehat{\mu}_{r,\zeta,\mu}$ is normalized and $\int_X \bar{H}(x)\,dx = 1$, the function $e_{r,\zeta,\mu}$ is continuous and satisfies

$$\int_X e_{r,\zeta,\mu}(x)\,dx = 0.$$

Define the rent correspondence by

$$\Phi(r,\zeta,\mu) = \arg\max_{q\in\mathcal{R}} \int_X q(x) e_{r,\zeta,\mu}(x)\,dx.$$

The set $\mathcal{R}$ is compact and convex, and the objective is continuous and affine in q. Berge's maximum theorem therefore implies that Φ is nonempty, convex-valued, compact-valued, and upper hemicontinuous. The full aggregate correspondence is

$$\mathcal{T}(r,\zeta,\mu) = \Phi(r,\zeta,\mu) \times \{\widehat{\zeta}_{r,\zeta,\mu}\} \times \{\widehat{\mu}_{r,\zeta,\mu}\}.$$

By construction, $\mathcal{T}$ maps $\mathcal{K}$ into itself, has nonempty convex compact values, and is upper hemicontinuous. Fan–Glicksberg's theorem therefore yields a fixed point $(r^*,\zeta^*,\mu^*) \in \mathcal{K}$ such that

$$(r^*,\zeta^*,\mu^*) \in \mathcal{T}(r^*,\zeta^*,\mu^*).$$

In particular,

$$\zeta^* = \widehat{\zeta}_{r^*,\zeta^*,\mu^*}, \qquad \mu^* = \widehat{\mu}_{r^*,\zeta^*,\mu^*}.$$

The second equality is equivalent to $\mu^* Q_{r^*,\zeta^*,\mu^*} = 0$, so μ^* is stationary for the worker dynamics generated by (r^*,ζ^*,μ^*).

It remains to verify housing-market clearing. Let $e^*(x) = R_{\mu^*}(x) - \bar{H}(x)$. Suppose $e^* \not\equiv 0$. Since $\int_X e^*(x)\,dx = 0$, the function e^* takes both positive and negative values. Because r^* is a stationary supporting rent schedule, it lies in the interior of $\mathcal{R}$. There therefore exists $\delta > 0$ small enough that

$$q_\delta(x) = r^*(x) + \delta\big(e^*(x) - e^*(\bar{D})\big)$$

still belongs to $\mathcal{R}$; subtracting $e^*(\bar{D})$ preserves the boundary condition $q_\delta(\bar{D}) = r_A$. But then

$$\int_X q_\delta(x) e^*(x)\,dx = \int_X r^*(x) e^*(x)\,dx + \delta \int_X \big(e^*(x)\big)^2\,dx,$$

because $\int_X e^*(x)\,dx = 0$. The second term is strictly positive, which contradicts the fact

that $r^* \in \Phi(r^*, \zeta^*, \mu^*)$. Therefore $e^* \equiv 0$, so

$$R_{\mu^*}(x) = \bar{H}(x) \qquad \text{for every } x \in X.$$

The fixed point has therefore delivered all equilibrium objects that matter for behavior. Let $(U^*, W^*, J^{F,*}, J^{V,*})$ denote the Bellman solution associated with (r^*, ζ^*, μ^*). By construction, these objects satisfy the worker and firm Bellman equations, the vacancy rule, the stationarity condition, the boundary condition $r^*(\bar{D}) = r_A$, and housing-market clearing. Hence

$$\left(r^*, U^*, W^*, J^{F,*}, J^{V,*}, \mu^*\right)$$

satisfies the definition of a stationary equilibrium. This proves that the economy admits at least one stationary equilibrium. □

C Details on Empirical Application

C.1 Calibration targets: data sources and construction

Population and land area. Census-tract centroids are obtained from the Census Bureau Gazetteer (2010 boundaries). Each tract is assigned to a 10 km ring based on Haversine distance from the CBD. Ring-level population comes from ACS 2015–2019 (NHGIS table) and land area from the Gazetteer. We compute population density as total ring population divided by total ring land area (in km^2).

Job density and high-pay share. We use LEHD LODES 2019 Workplace Area Characteristics (WAC) and Origin-Destination (OD) files for California. Each census block is assigned a ring using block-level coordinates from the LODES geographic crosswalk (`ca_xwalk.csv`). Job density at workplace ring y is total LODES jobs located in ring y divided by ring land area. The high-pay share is the fraction of jobs with monthly earnings above \$3,333 (SE03 category), computed from WAC SE01–SE03 files.

Commuting distance. Average commuting distance by home ring is computed from LODES 2019 OD files. For each home–work block pair, we compute Haversine distance between block centroids. We then average across all workers residing in each home ring, weighting by the OD flow count.

Rents and income. Ring-level median two-bedroom rent is taken from the Longitudinal Tract Data Base (LTDB, 2012 wave). Per-capita income is from the same source. Both are entered as ring-level averages weighted by tract population.

Unemployment and college share. Tract-level unemployment rates and college-graduate shares come from the LTDB (2008–2012 ACS estimates), aggregated to rings with population weights.

Move rate. The annual residential move rate is computed from ACS 2015–2019 (NHGIS) as the tract-level share of residents who report having moved from another location within the same MSA in the past year. We run a population-weighted local-linear kernel regression (Gaussian kernel, bandwidth 10 km) of this share on distance to the CBD.

Kernel-regression procedure. All distance–outcome profiles are estimated using the same procedure. Tract-level (or block-level, for LODES variables) observations are binned into 0.1 km cells to reduce computation. We then run a weighted local-linear kernel regression with a Gaussian kernel, bandwidth 10 km, evaluated at the ring centers (5, 15, . . . , 95 km). Standard errors are computed via 200 bootstrap replications at the tract level.

C.2 Parameter values

C.3 Discussion of calibration approach

The current calibration proceeds in two stages. In a first pass, parameters governing spatial profiles—housing supply $\bar{H}(x)$, firm-job capacity $\bar{M}(y, z)$, and calibration targets for density, rents, and high-pay shares—are set directly from the data as described above. In a second pass, preference, technology, and search parameters are adjusted by hand to bring the model's equilibrium profiles close to the data targets shown in Figure 9. Specifically, production parameters (A_z, $\alpha_{A,z}$, κ_{age}) are tuned to match the wage and high-type-share gradients; search parameters (λ^U, λ^E, χ_M) target aggregate flow rates (job-finding, separation, job-to-job); and housing parameters (λ^H, K_0) target the overall move rate. A formal method-of-moments estimation is planned for future iterations of the model.

Table 3: Calibrated parameter values

Parameter	**Symbol**	**Value**	**Source / Target**
Preferences			
Discount rate	ρ	0.05	Standard
Consumption share	α	0.72	CEX housing expenditure share
Bargaining weight	β	0.55	Hand-set
Job-choice smoothing	σ_J	0.55	Hand-set
Housing-choice smoothing	σ_H	0.50	Hand-set
Commuting disutility	χ_a	(0.30, 0.35, 0.40)	Commute-distance profile
Amenity gradient	γ_A	0.004	Rent gradient
Production			
Firm productivity	A_z	(4800, 8800)	Wage level
HC return	κ_h	(0.15, 0.30)	Wage growth
Workplace gradient (low z)	$\alpha_{A,1}$	0.003	High-z share profile
Workplace gradient (high z)	$\alpha_{A,2}$	0.020	High-z share profile
Commute-cost level	τ_1	0.0017	Commute-distance profile
Commute-cost convexity	τ_2	1.10	Commute-distance profile
Labor-market flows			
UE contact rate	λ_a^U	(1.45, 1.15, 0.90)	Job-finding rate
OTJ contact rate	λ_a^E	(0.42, 0.30, 0.20)	Job-to-job rate
Matching elasticity	γ	0.50	Standard
Matching scale	χ_M	1.35	Unemployment level
Separation rate	δ_a	(0.090, 0.071, 0.058)	JOLTS separation rates
Separation–commute penalty	δ_τ	0.10	Hand-set
Search decay	η	8.0	Commute-distance profile
Housing and mobility			
Housing arrival rate	λ^H	0.17	Annual move rate
Housing kernel decay	η_H	0.095	Move-distance distribution
Moving cost (fixed)	K_0	0.20	Move-rate level
Moving cost (distance)	K_1	0.0025	Move-distance distribution
Vacancy activation smoothing	σ_V	7000	Hand-set
Vacancy posting cost	c^V	192	Activation share
Vacancy slack target	—	0.18	JOLTS vacancy rate