Rent Caps, Spillovers, and Investment: A Dynamic Housing Game with Landlord Portfolios

S. Navid Hojaji*

May 2025

Early Draft

Abstract

We study how fragmented rent regulation redistributes markups, maintenance, and investment across space. We build a dynamic portfolio model of landlords who price new leases against downward-sloping residual demand, maintain a durable quality stock, and face incumbent rent-increase caps, eviction costs, and permit delays. Three results follow. (i) Tightening caps shifts markups to the unconstrained margin, widening the incumbent–new wedge. (ii) Heterogeneous rules reallocate listings and investment toward less restrictive jurisdictions, raising rents there (cross-border spillovers). (iii) With interior frictions, renovictions occur at a threshold that trades reset gains against regulatory costs. A transparent welfare breakdown attributes decentralization losses to investment, allocation, mobility, markups, and spillovers on identical state/action supports. Calibrated to five moments, the model reproduces patterns from San Francisco, Berlin, and Cambridge. Mobility and investment account for 55–65% of the welfare gap; harmonizing rules lowers it by 6–7%, while tightening caps raises it by about 3%. frictions

Keywords: rent regulation; spatial spillovers; landlord portfolios; maintenance investment; Markov-perfect equilibrium; welfare decomposition

*Heriot-Watt University, navhoj@hw.ac.uk

1 Introduction

Urban rents have risen sharply in many productive cities, renewing interest in rent-increase caps and eviction restrictions (Arnott, 1995; Diamond et al., 2019a; Kholodilin, 2024). These protections are often adopted unevenly across nearby jurisdictions. Classic spatial logic suggests that local frictions can spill across borders when supply is inelastic and mobility imperfect (Rosen, 1974; Roback, 1982; Saiz, 2010; Hsieh and Moretti, 2019), yet we lack a unified dynamic account of how fragmented protections reallocate market power and investment across space and over time.

This paper asks: when rent regulation differs across jurisdictions, how do landlords with portfolios reallocate markups, maintenance, and investment—and what are the implications for rents, quality, and welfare on both sides of the border? The question matters because many regimes insure incumbents while leaving new leases and neighboring districts relatively unconstrained.

We establish three central findings. First, tightening incumbent caps shifts pricing power to the unconstrained margin: new-lease rents rise relative to capped incumbent rents, widening the incumbent–new wedge. Second, regulatory heterogeneity tilts listings and maintenance toward laxer jurisdictions, raising rents there and generating cross-border spillovers. Third, when eviction and permitting are costly but not prohibitive, renovictions emerge at a threshold that trades the reset value of a new lease against regulatory costs. A welfare accounting attributes decentralization losses primarily to muted mobility and under-investment on the incumbent margin, with smaller contributions from markups and cross-border spillovers. In our calibration, mobility+investment account for roughly 55–65% of the gap between the decentralized equilibrium and a planner benchmark; tightening caps increases the gap by about 3%, mobility vouchers reduce it by about 6%, and light harmonization cuts it by about 7%. These patterns align with San Francisco's expansion, Berlin's freeze, and Cambridge's decontrol (Diamond et al., 2019a; Autor et al., 2014; Hahn et al., 2024).

Intuition. Consider two adjacent districts. Caps bind first on incumbents, compressing returns to maintaining those units and increasing the value of resetting to a new lease. Because tenant mobility is imperfect, demand on new listings is downward-sloping, so the bite of the policy shows up where it is not regulated: higher markups on new leases and on the laxer side of the border. Landlords with units in both districts shift listings and investment until portfolio returns equalize. With interior frictions, renovictions occur when the reset gain exceeds eviction and

permit costs. This is pass-through with a regulated margin; incidence depends on local demand curvature and the elasticity of portfolio substitution (cf. Miravete et al., 2018).

We build a dynamic environment in which landlords own portfolios across jurisdictions, new listings face downward-sloping residual demand microfounded by a standard dynamic discrete-choice block (Rust, 1987; Hotz and Miller, 1993), quality is durable and maintained at a convex cost, and rent increases, eviction, and permitting are regulated. We study stationary equilibria and derive comparative statics mapping policy primitives to rents, vacancy, and quality-adjusted supply. To interpret welfare, we implement an accounting that isolates five channels—investment, allocation, mobility, markups, and spillovers—by evaluating neutralized environments on identical state/action supports. Calibrated examples reproduce qualitative patterns from San Francisco, Berlin, Cambridge decontrol, and Scotland's freeze (Diamond et al., 2019a; Autor et al., 2014; Hahn et al., 2024; Hojaji et al., 2025).

Value added relative to existing work. Compared with spatial–equilibrium and hedonic foundations (Rosen, 1974; Roback, 1982; Saiz, 2010; Hsieh and Moretti, 2019), we add landlord portfolios and incumbent-only regulation in a durable-quality setting, which yields cross-border pass-through and portfolio reallocation that those frameworks leave implicit. Relative to dynamic IO and durable adjustment (Berry et al., 1995; Nevo, 2001; Bajari et al., 2007; Doraszelski and Pakes, 2007), we embed rent-cap, eviction, and permit frictions and characterize how markup reallocation to new leases arises when the incumbent margin is constrained. Relative to empirical rent-control episodes (Diamond et al., 2019a; Asquith, 2019; Autor et al., 2014, 2019; Hahn et al., 2024; Hojaji et al., 2025), we provide a single mechanism that organizes wedges, renovictions, and spillovers and a welfare breakdown that attributes changes to specific channels.

Our contribution is twofold. Substantively, we offer one mechanism that rationalizes wedges, renovictions, and cross-border spillovers and quantify their welfare importance. Methodologically, we provide a transparent and portable welfare attribution in a dynamic, general-sum setting by holding the environment fixed across counterfactuals and documenting identification.

Section 2 presents the model; Section 3 develops equilibrium, comparative statics, and the welfare accounting; Section 4 reports calibrated results and policy counterfactuals; Section 5 concludes.

2 Model setup

Notation. We use capitals for random variables and lower-case for their realisations (e.g., S_t vs. s_t); sets are calligraphic. For a statement A, $\mathbb{1}\{A\}$ denotes the indicator. Expectations are $\mathbb{E}[\cdot]$ (conditional on s_t unless noted), probabilities are $\mathbb{P}(\cdot)$, and $\log$ is the natural log. Time is discrete $t = 0, 1, \ldots$; tenants are indexed by i, units by $h \in \mathcal{H}$, and jurisdictions by $j \in \mathcal{J}$ with mapping $j(h)$. Incumbent rents are $r_{inc,h,t}$, new-lease rents $r_{h,t}$, quality $q_{h,t} \in \mathcal{Q}$, maintenance $m_{h,t} \geq 0$, and $\mathrm{dist}(h)$ is a location disutility proxy. The policy vector in j is $\mathcal{R}_j^{\mathrm{pol}} = \{\Delta\bar{r}_j, \phi_j^{\mathrm{evict}}, \psi_j^{\mathrm{permit}}\}$. We normalise the outside option to zero, take the common discount factor $\delta \in (0, 1)$, and model tenant taste shocks as Type-I EV with logit scale $1/\beta$ ($\beta > 0$). The aggregate state is $S_t = (\{s_{h,t}\}_{h\in\mathcal{H}}, z_t)$; unless stated, rents and qualities are in levels. A complete table of notation appears in Appendix L.

We develop a dynamic game-theoretic model of the urban rental housing market in which a finite set of landlords interact strategically over time in a city composed of multiple neighbourhoods. Each neighbourhood, indexed by $j \in \mathcal{J}$, differs in its regulatory environment and rental market characteristics. Landlords are forward-looking agents who own spatially distributed housing portfolios and face frictions in rent adjustment, tenant eviction, and investment decisions. Tenants are heterogeneous and mobile, and choose housing based on rent, location, and match-specific utility. The model integrates frictions from housing policy with elements of spatial industrial organisation and dynamic contracting.

Agents and timing. Let $\mathcal{L} = \{1, \ldots, L\}$ denote a finite set of landlords. Each landlord $\ell \in \mathcal{L}$ owns a portfolio $\mathcal{H}_\ell \subseteq \mathcal{H}$ of rental units across neighbourhoods. In each period t, landlords choose rents $r_{h,t}$, set maintenance or upgrade levels $m_{h,t}$, and decide whether to evict sitting tenants, subject to regulatory constraints. Each landlord aims to maximise the expected discounted value of net rental income.

Tenants are modelled as a continuum indexed by $i \in [0, 1]$. In each period, tenant i decides whether to remain in their current dwelling or search for a new unit. Utility is match-specific and depends on rent, location, housing quality, and idiosyncratic preferences. Mobility is costly due to search frictions, switching costs, and moving expenses.

The regulatory environment is represented by a non-strategic agent who sets neighbourhood-specific policies $\mathcal{R}_j^{\mathrm{pol}} = \{\Delta\bar{r}_j, \phi_j^{\mathrm{evict}}, \psi_j^{\mathrm{permit}}\}$, governing: (i) rent increase caps $\Delta\bar{r}_j$; (ii) eviction delays or costs ϕ_j^{evict}; and (iii) permitting lags for renovations ψ_j^{permit}.

Timing within period t. The sequence of events within each period is as follows: (i) landlords observe state variables and set rents $r_{h,t}$, investment/maintenance decisions $m_{h,t}$, and eviction attempts; (ii) tenants observe rents and housing attributes, then decide whether to stay or move; (iii) vacant units are matched with searching tenants; (iv) investments are realised with delays, and units depreciate if not maintained; and (v) regulatory constraints bind and state transitions occur.

2.1 State variables

Let $S_t = \big(\{s_{h,t}\}_{h\in\mathcal{H}},\, z_t\big) \in \mathcal{S}$, where each unit state is

$$s_{h,t} = \big(T_{h,t},\, r_{inc,h,t},\, q_{h,t},\, p_{h,t},\, j(h)\big),$$

with $T_{h,t} \in \{0,1\}$, $r_{inc,h,t} \in \mathcal{G}^r_{j(h)}$, $q_{h,t} \in \mathcal{Q}$, $p_{h,t} \in \{0,1,\dots,\psi^{\text{permit}}_{j(h)}\}$, and $j(h) \in \{A,B,\dots\}$. Idiosyncratic match values $\theta_{i,h,t}$ are integrated out and enter only via the matching kernel implied by the tenant DDC; they are not part of S_t. The industry summary $z_t \in \mathcal{Z}$ collects finite sufficient statistics (e.g., vacancy by quality bin and jurisdiction).

2.2 Tenant problem

Each tenant i solves a stationary dynamic discrete choice problem with Type-I EV idiosyncratic shocks $\{\varepsilon_{i,h,t}\}$. Let the per-period systematic utility be

$$u_i(h; s_t) = \theta_{i,h,t} + \eta_{j(h)} - \kappa\,\mathrm{dist}(h) + \xi(q_{h,t}),$$

and let incumbents face $r_{inc,h,t}$ while new leases pay $r_{h',t}$. The Bellman equation is

$$V_i(h, s_t) = \max\Big\{\, u_i(h; s_t) - r_{inc,h,t} + \delta\,\mathbb{E}[V_i(h, s_{t+1}) \mid s_t],\ \tilde{V}^{\text{move}}(s_t)\Big\}, \tag{1}$$

with inclusive value for moving (logit):

$$\tilde{V}^{\text{move}}(s_t) = -c^{\text{move}} + \frac{1}{\beta}\log\sum_{h'\in\mathcal{H}} \exp\Big(\beta\big[u_i(h'; s_t) - r_{h',t} + \delta\,\mathbb{E}[V_i(h', s_{t+1}) \mid s_t]\big]\Big), \tag{2}$$

where $\beta > 0$ is the logit scale. Endogenous turnover follows from the optimal stay/move choice probabilities.

2.3 Landlord problem

Each landlord ℓ maximises the expected discounted stream of net rental income by choosing rents $r_{h,t}$, maintenance decisions $m_{h,t}$, and eviction efforts $e_{h,t}$ to solve:

$$\max_{\{r_{h,t},m_{h,t},e_{h,t}\}} \mathbb{E}\left[\sum_{t=0}^{\infty}\delta^t \sum_{h\in\mathcal{H}_\ell}\left(\pi_{h,t} - C(m_{h,t}) - \phi_j^{\text{evict}}\cdot e_{h,t}\right)\right], \tag{3}$$

where $\pi_{h,t} = r_{h,t}\cdot T_{h,t}$ denotes rental income when the unit is occupied, and $C(m_{h,t})$ is the cost of maintenance or quality upgrades. The term $e_{h,t}\in\{0,1\}$ represents an eviction attempt, which succeeds only if permitted by regulation, and incurs cost ϕ_j^{evict}.

Quality and permit status. Let $p_{h,t}\in\{0,1,\dots,\psi_{j(h)}^{\text{permit}}\}$ be the remaining permit lag (0 means no pending permit). Then

$$q_{h,t+1} = (1-\delta_q)\,q_{h,t} \;+\; \chi\, m_{h,t}\,\mathbb{1}\{p_{h,t}=0\}, \tag{4}$$

$$p_{h,t+1} = \begin{cases} \psi_{j(h)}^{\text{permit}} & \text{if a permitted upgrade is initiated at } t,\\ \max\{0,\, p_{h,t}-1\} & \text{otherwise.}\end{cases} \tag{5}$$

Thus, investment only augments quality once the permit countdown hits zero.

Frictional constraints create trade-offs: removing a low-paying tenant may improve portfolio returns in the long run but entails regulatory costs and possible vacancy.

2.4 Equilibrium concept

We characterise a Markov Perfect Equilibrium (MPE) in which each landlord's strategy $\sigma_\ell(s_t)$ is a best response to other landlords' strategies and tenant behaviour. Tenants solve a stationary dynamic decision problem, forming expectations over future rents and matching outcomes.

The evolution of the state vector s_{t+1} is governed by the transition function

$$s_{t+1} = f(s_t,\sigma,\mu), \tag{6}$$

where μ summarises institutional policies such as eviction regulation and rent caps.

The MPE framework captures the endogenous dynamics of market structure, tenant turnover, rent distributions, and housing quality over time.

2.5 Baseline discrete model

To obtain clean existence and comparative-static results, we work with a finite, discretized version of the model that is behaviourally equivalent in spirit to the continuous formulation above but yields a discounted finite stochastic game.

Jurisdictions and landlords. There are two jurisdictions $j \in \{A, B\}$. Jurisdiction A has a binding rent-increase cap $\Delta\bar{r}A > 0$; jurisdiction B is effectively uncapped (implemented by a weak/loose cap that does not bind on the state space we consider).[1] A finite set of landlords $\mathcal{L} = \{1, \ldots, L\}$ each owns a finite portfolio $\mathcal{H}_\ell \subseteq \mathcal{H}$.

Finite action sets. For every unit h in jurisdiction $j(h)$ and period t, the landlord chooses:

$$r_{h,t} \in \mathcal{G}^r_{j(h)}, \qquad m_{h,t} \in \mathcal{M} := \{0, 1, \ldots, \bar{m}\}, \qquad e_{h,t} \in \{0, 1\},$$

where $\mathcal{G}^r_j$ is a finite grid of feasible contract rents in j , $m_{h,t}$ is maintenance/upgrade intensity, and $e_{h,t}$ is an eviction attempt.

Finite state space. The per-unit state is

$$s_{h,t} = \big(T_{h,t},\, r_{inc,h,t},\, q_{h,t},\, j(h)\big),$$

where $T_{h,t} \in \{0, 1\}$ is occupancy, $r_{inc,h,t} \in \mathcal{G}^r_{j(h)}$ is the incumbent's regulated rent, and $q_{h,t} \in \mathcal{Q}$ is quality on a finite grid $\mathcal{Q} = \{q^1, \ldots, q^{\bar{q}}\}$. The industry summary is a finite vector

$$z_t \in \mathcal{Z}$$

(e.g., vacancy counts by j and quality bin, or any finite sufficient statistic for rivals and demand).

The full state is $S_t = (\{s_{h,t}\}_{h\in\mathcal{H}},\, z_t) \in \mathcal{S}$ with $\mathcal{S}$ finite.

State transitions. (i) *Incumbent rent.* If no turnover,

$$r_{inc,h,t+1} = \min\{\, r_{inc,h,t} + \Delta\bar{r}j(h),\; r_{h,t}\,\}.$$

If turnover (voluntary exit or successful eviction), then $r_{inc,h,t+1} = r_{h,t}$.

[1] We keep all sets finite. "Uncapped" means the cap never binds on $\mathcal{S}$ rather than $\Delta\bar{r}B = \infty$.

(ii) *Quality dynamics.* Quality follows a discretized law of motion

$$q_{h,t+1} \;=\; \Gamma\big(q_{h,t},\, m_{h,t};\, j(h)\big) \in \mathcal{Q},$$

with Γ monotone in m and reflecting depreciation and permit frictions (delays can be encoded by the mapping Γ or by an auxiliary finite permit-status variable; both keep $\mathcal{S}$ finite).

(iii) Occupancy/matching. Given $(r_{h,t}, q_{h,t})$ and z_t, the tenant DDC implies a logit choice probability for each vacant unit and an optimal stay rate for incumbents. The resulting transition is a time-homogeneous Markov kernel

$$P(\cdot \mid S_t, A_t) \text{ on } \mathcal{S},$$

obtained by integrating the stay/move probabilities (from (1)) and eviction success $\rho_{j(h)}(e_{h,t})$ over the i.i.d. taste shocks.

Eviction success. An eviction attempt $e_{h,t} \in \{0,1\}$ succeeds with probability

$$\rho_{j(h)}(e_{h,t}) = \mathbb{1}\{e_{h,t}=1\} \cdot \lambda_{j(h)}(\phi^{\text{evict}}_{j(h)}),$$

where $\lambda_j(\cdot) \in [0,1]$ is decreasing and encapsulates legal frictions. The occupancy transition factors as

$$\mathbb{P}(T_{h,t+1}=0 \mid s_{h,t}, r_{h,t}, m_{h,t}, e_{h,t}, z_t) = \rho_{j(h)}(e_{h,t}) + \big(1-\rho_{j(h)}(e_{h,t})\big)\cdot \tilde{p}_{\text{vol}}(s_{h,t}, r_{h,t}, q_{h,t}; z_t),$$

where $\tilde{p}_{\text{vol}}$ is the voluntary-exit probability implied by the tenant DDC.

Per-period payoff. Per-unit profit is

$$\pi_{h,t} \;=\; r_{h,t}\, T_{h,t} \;-\; C(m_{h,t}) \;-\; \phi^{\text{evict}}_{j(h)}\, e_{h,t},$$

with $C : \mathcal{M} \to \mathbb{R}_+$ convex and $\phi^{\text{evict}}_j > 0$. Landlord ℓ's period payoff is $\sum_{h\in\mathcal{H}_\ell} \pi_{h,t}$ and the discounted objective is $\mathbb{E}\big[\sum_{t\geq 0}\delta^t \sum_{h\in\mathcal{H}_\ell}\pi_{h,t}\big]$, $0<\delta<1$.

Comment. Under this discretization, the environment is a discounted *finite* stochastic game with observable Markov state S_t, finite action sets, and bounded payoffs. We use this baseline in Section 3 to state propositions and (in Appendix A) to establish existence of a stationary Markov-perfect equilibrium; we restrict attention to the discretized model and do not claim continuous-state results.

Definition 1 (Stationary Markov-perfect equilibrium in the baseline game)**.** *Let $\mathcal{S}$ be finite and let $A_\ell(s)$ be finite for each ℓ and s. A stationary Markov strategy for landlord ℓ is a map $\sigma_\ell : \mathcal{S} \to \Delta(A_\ell(s))$. A profile $\sigma = (\sigma_\ell)_{\ell \in \mathcal{L}}$ is a stationary MPE if for every ℓ and $s \in \mathcal{S}$,*

$$\sigma_\ell(s) \in \operatorname{argmax}_{\alpha \in \Delta(A_\ell(s))} \; \mathbb{E}_{a \sim (\alpha, \sigma_{-\ell}(s))} \left[\pi_\ell(s,a) \; + \; \delta \sum_{s' \in \mathcal{S}} P(s' \mid s,a)\, V_\ell^\sigma(s') \right],$$

where V_ℓ^σ solves the Bellman fixed point

$$V_\ell^\sigma(s) = \mathbb{E}_{a \sim \sigma(s)} \left[\pi_\ell(s,a) + \delta \sum_{s' \in \mathcal{S}} P(s' \mid s,a)\, V_\ell^\sigma(s') \right].$$

2.6 Model Innovations and Extensions

This framework introduces several innovations that allow us to analyse regulatory policy in spatially segmented rental markets.

First, we incorporate spatial spillovers and arbitrage by allowing for regulatory heterogeneity across jurisdictions, such that $\mathcal{R}_j^{\text{pol}} \neq \mathcal{R}_{j'}^{\text{pol}}$. This induces landlord reallocation across space in response to policy differences. The model permits geographic mobility of capital, capturing spatial spillover effects in equilibrium.

Second, the model jointly endogenises eviction and investment decisions. This allows for the simulation of renovictions—evictions tied to maintenance or quality upgrading—as equilibrium outcomes, shaped by policy parameters and market expectations.

Third, we introduce landlord heterogeneity. Portfolios may differ in size, financial constraints, and regulatory exposure. These asymmetries permit the analysis of market segmentation and the differential incidence of regulation across landlord types.

Finally, the model supports the evaluation of counterfactual policy regimes. We simulate comparative statics and dynamic adjustment paths under alternative regulatory designs, including maintenance subsidies, eviction reform, and uniform rent caps.

Microfoundations and reduced form: mapping

We rationalize the linear residual demand $D_j(r,q) = A_j - B_j r + \gamma q$ as a local first-order approximation to a tenant DDC model with i.i.d. type-1 EV shocks. Let tenant i's utility from unit h in j be $u_i(h) = \theta_{i,h} + \eta_j + v(q_h) - r_h + \varepsilon_{i,h}$, $\varepsilon_{i,h} \sim$ i.i.d. EV, and let a standard logit matching rule

determine the mass of takers. Around a baseline $(\bar{r}_j, \bar{q}_j)$, the market share $s_j(r, q)$ is differentiable and obeys

$$s_j(r,q) \approx s_j(\bar{r}_j, \bar{q}_j) + \underbrace{\frac{\partial s_j}{\partial r}\Big|_{\bar{r}_j,\bar{q}_j}}_{-\beta_j}(r - \bar{r}_j) + \underbrace{\frac{\partial s_j}{\partial q}\Big|_{\bar{r}_j,\bar{q}_j}}_{\gamma_j}(q - \bar{q}_j).$$

Multiplying by market size gives a local inverse demand that is linear in (r, q); normalizing constants yields A_j and slopes $B_j = \beta_j > 0$, $\gamma = \gamma_j \geq 0$.

Full details in Appendix H.

Figure 1 summarizes how frictions move incentives, states, and welfare; dashed/dotted arcs mark the direct cap effect and static protection, respectively.

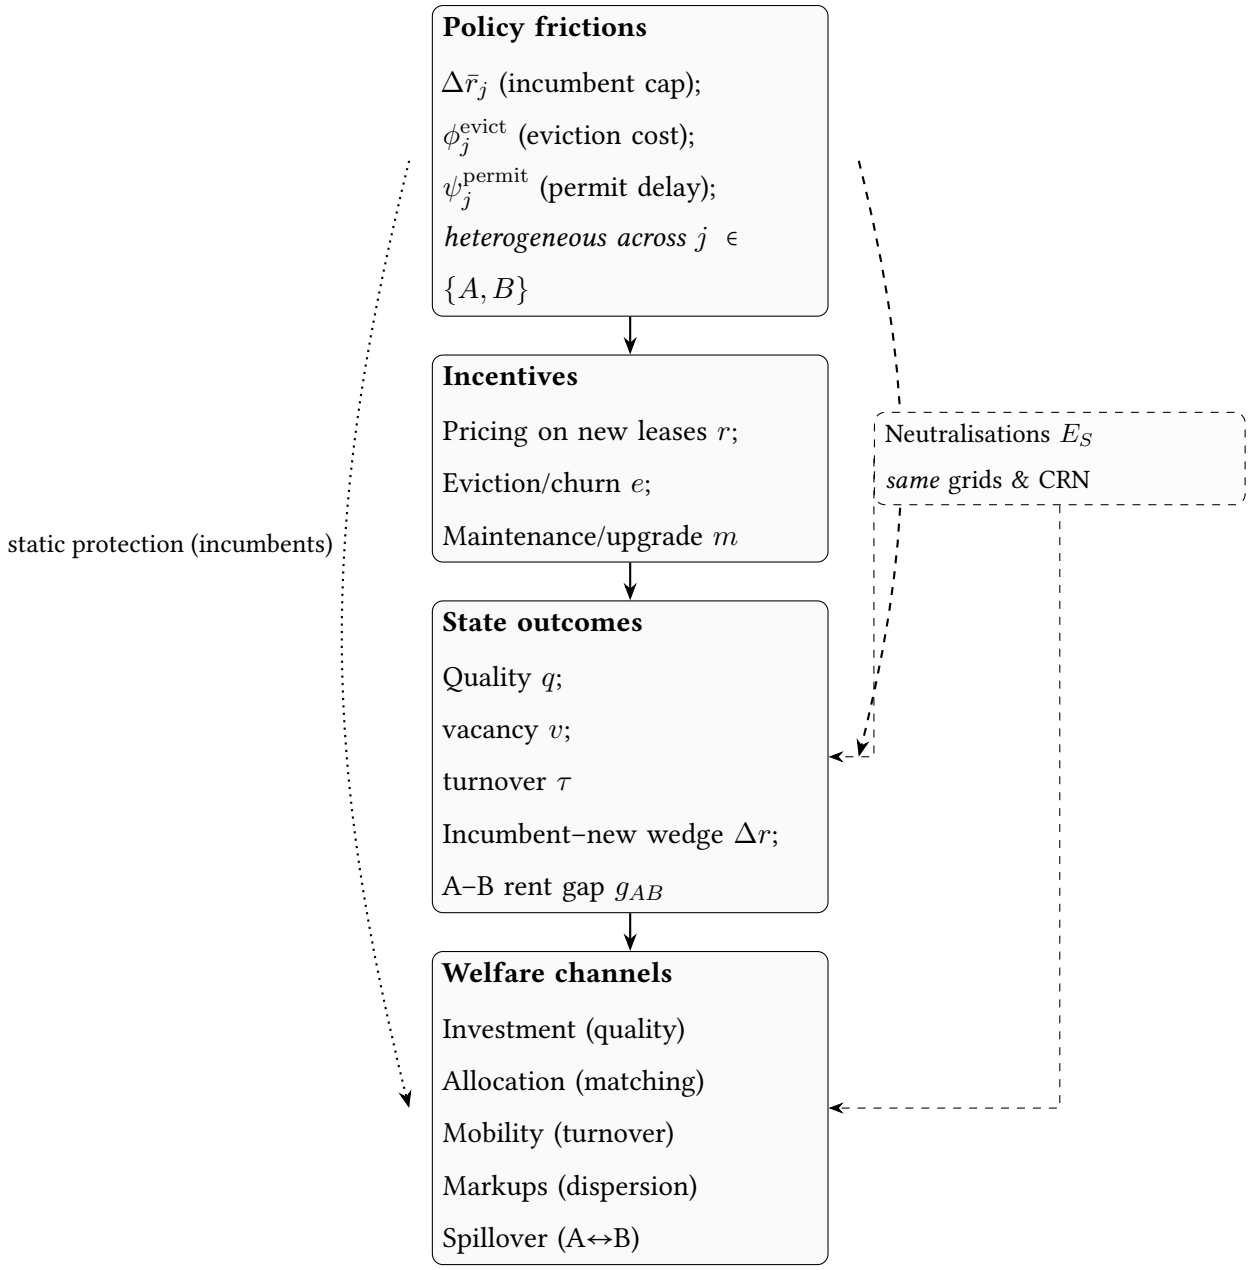


Figure 1: Mechanism map. Policy frictions shift incentives, which move state outcomes and the five welfare channels. The dashed arc is the direct cap effect on Δr; the dotted arc is static protection of incumbents; E_S denotes neutralisations (same grids & CRN) used for apples-to-apples welfare.

3 Theoretical analysis

This section formalises the strategic interaction between landlords and tenants in a dynamic rental housing market subject to regulatory frictions. We model landlord behaviour as a dynamic game with Markov Perfect Equilibrium (MPE), where each landlord maximises expected discounted rental income subject to regulatory constraints and tenant mobility. The framework captures endogenous tenant turnover, spatial capital reallocation, and strategic rent-setting behaviour.

Proposition (Existence of a stationary MPE in the baseline game)**.** *Existence holds as in Proposition A.1 (Appendix A).*

Remark 1 (Pure strategies and continuous model)**.** *The existence in Proposition A.1 is for (possibly) mixed stationary strategies; pure-strategy existence is not guaranteed in general finite stochastic games. For the continuous version of the model, existence can be obtained by standard grid-refinement/limit arguments under continuity and compactness (or by directly imposing conditions ensuring upper hemicontinuity and compactness of best responses). In this paper we state propositions in the finite baseline and report simulations in the richer continuous setup; the welfare and comparative-static conclusions are aligned across discretizations.*

3.1 Tenant optimisation problem

We restate the tenant's dynamic discrete choice problem in a form consistent with Section 2. Let the per-period systematic utility (excluding rent) be

$$u_i(h; S_t) = \theta_{i,h,t} + \eta_{j(h)} - \kappa \cdot \mathrm{dist}(h) + \xi(q_{h,t}),$$

with i.i.d. Type-I EV idiosyncratic shocks. Incumbents on unit h pay $r_{inc,h,t}$; new leases pay posted rents $r_{h',t}$.

The Bellman equation is

$$V_i(h, S_t) = \max\Big\{ u_i(h; S_t) - r_{inc,h,t} + \delta\, \mathbb{E}\big[V_i(h, S_{t+1}) \mid S_t\big],\ \tilde{V}^{\mathrm{move}}(S_t)\Big\},$$

where the logit inclusive value for moving is

$$\tilde{V}^{\mathrm{move}}(S_t) \;=\; -c^{\mathrm{move}} + \frac{1}{\beta} \log \sum_{h' \in \mathcal{H}} \exp\Big(\beta\big[u_i(h'; S_t) - r_{h',t} + \delta\, \mathbb{E}[V_i(h', S_{t+1}) \mid S_t]\big]\Big).$$

3.2 Landlord problem and Bellman equation

Each landlord $\ell \in \mathcal{L}$ owns a portfolio of units $\mathcal{H}_\ell$. At time t, the state vector $s_{h,t}$ for unit h includes occupancy status, incumbent rent, quality, and jurisdiction-specific regulatory parameters.

The landlord chooses a rental price $r_{h,t}$, maintenance investment $m_{h,t}$, and an eviction decision $e_{h,t} \in \{0,1\}$, subject to constraints on allowable rent increases $\Delta\bar{r}_j$, eviction frictions ϕ_j^{evict}, and renovation approval delays ψ_j^{permit}.

Let $\pi_{h,t} = r_{h,t} \cdot T_{h,t}$ denote rental revenue. The landlord's problem is to solve:

$$V_\ell(s_t) = \max_{\{r_{h,t}, m_{h,t}, e_{h,t}\}} \mathbb{E}\left[\sum_{t=0}^{\infty} \delta^t \sum_{h \in \mathcal{H}_\ell} \left(\pi_{h,t} - C(m_{h,t}) - \phi_j^{\text{evict}} \cdot e_{h,t}\right)\right], \tag{7}$$

subject to the law of motion for housing quality,

$$q_{h,t+1} = \Gamma\left(q_{h,t}, m_{h,t}; j\right),$$

where permit delays are implemented via a finite permit-status variable $p_{h,t} \in \{0, \ldots, \psi_{j(h)}^{\text{permit}}\}$ that advances deterministically while works are pending; upgrades to q take effect when p completes (cf. Appendix E).

3.3 Markov perfect equilibrium

We work with *stationary* Markov strategies in the finite baseline game of Section 2.5. Let $S_t \in \mathcal{S}$ denote the aggregate state and $A_\ell(S)$ landlord ℓ's finite action set at S. Given actions $a = (a_\ell)_{\ell \in \mathcal{L}}$, let $P(\cdot \mid S, a)$ be the time–homogeneous transition kernel on $\mathcal{S}$ implied by (i) tenants' dynamic discrete choice (logit) stay/move behaviour and the matching mechanism, and (ii) technological/regulatory frictions (rent-cap increments, eviction success, permits).

A *stationary Markov perfect equilibrium (MPE)* is a profile $\sigma = \{\sigma_\ell\}_{\ell \in \mathcal{L}}$ with $\sigma_\ell : \mathcal{S} \to \Delta(A_\ell(S))$ and value functions $\{V_\ell^\sigma\}_{\ell \in \mathcal{L}}$ such that, for every landlord ℓ and state $S \in \mathcal{S}$,

$$\sigma_\ell(S) \in \arg\max_{\alpha \in \Delta(A_\ell(S))} \mathbb{E}_{a \sim (\alpha, \sigma_{-\ell}(S))}\left[\pi_\ell(S,a) + \delta\, \mathbb{E}_{S' \sim P(\cdot|S,a)} V_\ell^\sigma(S')\right],$$

and V_ℓ^σ solves the Bellman fixed point

$$V_\ell^\sigma(S) = \mathbb{E}_{a \sim \sigma(S)}\left[\pi_\ell(S,a) + \delta\, \mathbb{E}_{S' \sim P(\cdot|S,a)} V_\ell^\sigma(S')\right], \qquad \forall\, S \in \mathcal{S}.$$

Tenants are not strategic players in the game; their optimal stay/move probabilities are integrated into $P(\cdot \mid S, a)$ via the DDC/matching block.

Remark (dynamic vs. stationary). The environment is *dynamic* (agents maximise discounted lifetime payoffs) but we focus on *stationary Markov* strategies that depend only on the current aggregate state S_t.

3.4 Strategic responses to policy frictions

We derive several comparative statics to illustrate how regulatory frictions affect landlord incentives in equilibrium. Let $r_{mkt,t}$ denote the optimal price chosen at t for a newly let (unconstrained) lease.

Proposition 1 (Rent-cap induced reallocation of markups)**.** *Assume:*

1. *New-lease demand is $D(r) = A - Br$ with $A, B > 0$ and is log-concave on its support (so the profit objective is strictly concave).*

2. *If no turnover, the incumbent rent evolves by*

$$r_{inc,t+1} = \min\{\, r_{inc,t} + \Delta\bar{r},\ r_{mkt,t} \,\};$$

   with turnover (voluntary exit or successful eviction), $r_{inc,t+1} = r_{mkt,t}$.

3. *Incumbent demand is weakly less elastic than new-lease demand due to switching costs (see Appendix B).*

4. *Marginal cost MC is locally well defined.*

Then, holding portfolio composition fixed, a tighter cap (lower $\Delta\bar{r}$) weakly increases the unconstrained optimal markup and price:

$$\frac{\partial r_{mkt,t}}{\partial \Delta\bar{r}} < 0, \qquad \frac{\partial}{\partial \Delta\bar{r}}\big(r_{mkt,t} - MC\big) < 0,$$

and widens dispersion between unconstrained and incumbent leases:

$$\frac{\partial}{\partial \Delta\bar{r}}\big(r_{mkt,t} - r_{inc,t}\big) < 0.$$

Proof sketch. *The cap slows $r_{inc,t}$'s drift toward $r_{mkt,t}$, raises the value of keeping incumbents, and increases scarcity on the unconstrained margin. The sign follows from the implicit-function argument in Appendix B under log-concavity.*

Proposition 2 (Policy spillovers from regulatory heterogeneity). *Consider two jurisdictions A, B. A landlord chooses nonnegative efforts $\{m_h^A\}_h$ and $\{m_h^B\}_h$ to maximise total profit subject to a budget $\sum_h C(m_h^A) + \sum_h C(m_h^B) \le M$, where $C : \mathbb{R}_+ \to \mathbb{R}_+$ is C^1 and strictly convex. Per–unit revenue functions are C^1, increasing, and concave: $\tilde{r}_h^{j\prime}(m) > 0$, $\tilde{r}_h^{j\prime\prime}(m) \le 0$ for $j \in \{A, B\}$. The effective revenue in A is capped at a ceiling $\bar{r}_A$, i.e.*

$$r_h^A(m) = \min\{\tilde{r}_h^A(m), \bar{r}_A\}, \qquad r_h^B(m) = \tilde{r}_h^B(m).$$

Assume (i) at the optimum a positive measure of A-units is (locally) capped, and (ii) the B-side equilibrium rent r^B clears a market with fixed supply $S^B > 0$ via a C^1 demand $D^B(r, q)$ with $D_r^B(r, q) < 0$ and $D_q^B(r, q) \ge 0$, where the quality index $q^B = Q^B(m^B)$ is C^1 and increasing in aggregate effort $m^B := \sum_h m_h^B$.

Then as the cap in A is loosened (i.e. $\bar{r}_A$ increases), the optimal allocation to B weakly falls and the equilibrium rent in B weakly falls:

$$\frac{\partial m^B}{\partial \bar{r}_A} < 0, \qquad \frac{\partial r^B}{\partial \bar{r}_A} < 0.$$

Equivalently, parameterising cap–tightness by a per–period increment parameter $\Delta\bar{r}A$ that increases with loosening, the same signs obtain: $\partial m^B/\partial \Delta\bar{r}A < 0$ and $\partial r^B/\partial \Delta\bar{r}A < 0$.

Proof sketch. (i) Reallocation. With C strictly convex, the KKT conditions (Lagrangian multiplier λ for the budget) imply for any uncapped unit: $\tilde{r}_h^{j\prime}(m_h^j) = C'(m_h^j) + \lambda$, while for capped A-units at the ceiling, $\tilde{r}_h^{A\prime}(m_h^A) = 0 < C'(m_h^A) + \lambda$ so m_h^A is (weakly) reduced to the kink or to zero. Increasing $\bar{r}_A$ shrinks the capped region in A; decreasing $\bar{r}_A$ expands it. Totally differentiating the KKT system shows that as $\bar{r}_A$ falls (tightens), aggregate m^A weakly falls and, to satisfy the resource constraint, m^B weakly rises. (Formal details in Appendix C; cf. Rockafellar (1970).)

(ii) GE price in B.) Market clearing $D^B(r^B, q^B) = S^B$ with $D_r^B < 0$, $D_q^B \ge 0$ and $q^B = Q^B(m^B)$ increasing implies by the Implicit Function Theorem that $r_q^B = -D_q^B/D_r^B \ge 0$; hence r^B is weakly increasing in q^B and thus in m^B (and weakly decreasing in $\bar{r}_A$). Full proof in Appendix C. □

Proposition 3 (Renoviction as strategic disinvestment). *Consider a single regulated unit with incumbent rent $r_{inc,t}$ and market rent $r_{mkt,t+2}$ available after turnover and renovation. Quality evolves $q_{t+1} = (1-\delta_q)q_t + \chi m_t$ with $\delta_q \in (0, 1)$ and $\chi > 0$. Let $P_{\text{exit}}(q)$ be C^1, strictly decreasing, and Lipschitz; let $C(\cdot)$ be increasing and strictly convex with $C(0) = 0$; and let the renovation package*

cost be $\phi_{\text{reno}} + C(m^*)$ *for some feasible* $m^* \geq 0$*. Assume formal eviction is sufficiently costly that the landlord does not choose* $e_t = 1$ *on the path considered.*[1] *In the finite baseline game (Section 2.5), value functions are bounded and depend continuously on* q*.*[2]

Fix any "maintain" choice $m^{\text{keep}} > 0$ *and define*

$$q^0 \equiv (1-\delta_q)q_t, \quad q^{\text{keep}} \equiv (1-\delta_q)q_t + \chi\, m^{\text{keep}}, \quad \Delta P_t \equiv P_{\text{exit}}(q^0) - P_{\text{exit}}(q^{\text{keep}}) > 0,$$

and the rent gap $G_t \equiv r_{mkt,t+2} - r_{inc,t}$*. Let* L_V *bound the continuation difference:* $\left|\mathbb{E}[V_{t+3}] - \mathbb{E}[V_{t+1}^{\text{inc}}]\right| \leq L_V$*. If*

$$G_t \;>\; \phi_{\text{reno}} + C(m^*) \;+\; L_V \;+\; \frac{C(m^{\text{keep}})}{\delta\, \Delta P_t}, \tag{8}$$

then "under-maintenance" ($m_t = 0$ *to raise exit risk) strictly dominates maintaining quality at* m^{keep}*:*

$$V_t^{\text{renovict}} \;>\; V_t^{\text{maintain}}.$$

Moreover, the set of states for which (8) holds expands when the cap tightens (which lowers $r_{inc,t}$ *and thus raises* G_t*) or when maintenance becomes more costly; see Appendix D and Topkis (1998).*

Proof. See Appendix D. □

Proposition 4 (Qualitative policy effects (signs))**.** *Under the baseline game in Section 2.5 and the MCS assumptions in Appendix I (increasing differences of landlords' Bellman problems in actions and policy parameters; isotone best responses; MPE selection by extremal fixed points), we have:*

1. ***Maintenance subsidy*** $\tau \uparrow$***:*** *increases* m *on regulated units and weakly raises* q*; reduces* $\Delta W_{\text{investment}}$*.*

2. ***Tenant vouchers (effective*** $c^{\text{move}} \downarrow$***):*** *weakly increases turnover on margin and reduces* $\Delta W_{\text{mobility}}$*.*

3. ***Policy harmonisation (align*** $\Delta\bar{r}A$ ***and*** $\Delta\bar{r}B$***; align*** ϕ_A^{evict} ***and*** ϕ_B^{evict}***):*** *reduces cross-jurisdiction reallocation and* $\Delta W_{\text{spillover}}$*.*

[1] Formally, $e_t = 0$ is (strictly) weakly dominant whenever $\phi_j^{\text{evict}} \geq r_{mkt,t+2} - r_{inc,t} + \delta\,\Delta V_{\text{post}}$, so we restrict attention to the voluntary-exit margin.

[2] This follows from the contraction mapping property of the Shapley/Bellman operator in a finite discounted MDP/game; see Puterman (1994) or Stokey et al. (1989).

4. ***Two-tier regulation (incumbent cap; market for new leases):*** *reduces mark-up dispersion and* $\Delta W_{\text{markups}}$ *while keeping incumbent protection.*

Proof. See Appendix I. □

3.5 Welfare

Results at a glance. We adopt a planner benchmark and decompose the decentralised–planner welfare gap as a pure identity (no further behavioural assumptions beyond Sections 2–2.5).

Lemma 1 (Planner weakly dominates decentralised)**.** *Let* $V^{\text{planner}}(s_0)$ *denote the value of the planner's problem at* s_0 *and* $V^{\text{MPE}}(s_0)$ *the value of the decentralised MPE starting from the same state. Then*

$$V^{\text{planner}}(s_0) \;\geq\; V^{\text{MPE}}(s_0).$$

Proof. The planner can replicate the MPE allocation and actions at every state, which is feasible. Since the planner maximises over a superset of feasible policies, her objective is weakly greater. □

Definition 2 (Welfare gap)**.** *Define the welfare gap*

$$\Delta W \;:=\; V^{\text{planner}}(s_0) \;-\; V^{\text{MPE}}(s_0).$$

We take social welfare to be the present discounted value of total tenant utility plus landlord net income. Regulatory frictions that limit rent increases or slow evictions can protect incumbents but may distort investment behaviour and degrade housing quality over time. We organise impacts into three headline channels: (i) a static protection effect for incumbents, (ii) dynamic investment distortions via maintenance/quality, and (iii) general-equilibrium spillovers from spatial capital reallocation. In what follows we characterise conditions under which these trade-offs yield net losses and identify policy tools (e.g., maintenance subsidies) that partially offset distortions.

3.5.1 Planner's problem: benchmark allocation

We posit a utilitarian social planner who internalises all frictions and chooses *allocations*, *prices*, and *investment/eviction actions* to maximise discounted welfare. The planner faces the same technological and regulatory primitives as in the decentralised environment (Sections 2–2.5). For

consistency with the baseline finite game, we work with a finite set of representative tenant types $\mathcal{I}$ with masses $\{w_i\}_{i\in\mathcal{I}}$ (summing to 1); all state variables take values on finite grids.

Controls and timing. At each t, the planner chooses for every unit h: rent $r_{h,t} \in \mathcal{G}^r_{j(h)}$, maintenance/investment $m_{h,t} \in \mathcal{M}$, and eviction $e_{h,t} \in \{0,1\}$. The planner also chooses an assignment matrix $\mu_{i,h,t} \in \{0,1\}$ indicating whether tenant type i is placed in unit h at t. Let $T_{h,t} = \sum_i w_i\, \mu_{i,h,t} \in \{0,1\}$ denote occupancy.

Objective. Tenant i's flow utility in unit h is

$$u_{i,t}(h) = \theta_{i,h,t} + \eta_{j(h)} - \kappa \cdot \mathrm{dist}(h),$$

with match component $\theta_{i,h,t}$ included in the state. Let $\mathbb{1}_{\mathrm{move}}(i,t)$ indicate whether i's assignment at t differs from $t-1$. The planner maximises discounted welfare

$$\max_{\{r,m,e,\mu\}} \mathbb{E}\left[\sum_{t=0}^{\infty} \delta^t \Big(W_t^{\mathrm{tenant}} + W_t^{\mathrm{landlord}}\Big)\right], \qquad \delta \in (0,1), \tag{9}$$

$$W_t^{\mathrm{tenant}} = \sum_{i\in\mathbb{I}} \sum_{h\in\mathbb{H}} w_i\, \mu_{i,h,t} \Big(u_{i,t}(h) - r_{h,t} - c^{\mathrm{move}} \cdot \mathbb{1}_{\mathrm{move}}(i,t)\Big),$$

$$W_t^{\mathrm{landlord}} = \sum_{h\in\mathcal{H}} \Big(r_{h,t}\, T_{h,t} - C(m_{h,t}) - \phi^{\mathrm{evict}}_{j(h)}\, e_{h,t}\Big).$$

Accounting convention. Because $T_{h,t} = \sum_i w_i\, \mu_{i,h,t}$, the period-$t$ rent terms cancel: $\sum_{i,h} w_i\, \mu_{i,h,t}(-r_{h,t}) + \sum_h r_{h,t} T_{h,t} = 0$. Thus (9) is equivalent to maximising *real surplus*:

$$\sum_{i,h} w_i\, \mu_{i,h,t} \big(u_{i,t}(h) - c^{\mathrm{move}} \cdot \mathbb{1}_{\mathrm{move}}(i,t)\big) - \sum_h \big(C(m_{h,t}) + \phi^{\mathrm{evict}}_{j(h)}\, e_{h,t}\big),$$

i.e. transfers do not affect welfare (they may affect distribution).

Constraints. Feasibility and technology mirror the decentralised model:

(10)

(i) Assignment and capacity: $\sum_h \mu_{i,h,t} \le 1 \;\; (\forall i,t), \qquad \sum_i w_i \, \mu_{i,h,t} \le 1 \;\; (\forall h,t).$

(11)

(ii) Quality dynamics: $q_{h,t+1} = \Gamma\big(q_{h,t}, m_{h,t}; j(h)\big) \in \mathcal{Q},$ with Γ encoding depreciation and permit frictions.

(iii) Regulatory (second-best) constraints:

(12)

(a) Rent-cap increment: $r_{h,t+1} \le r_{h,t} + \Delta \bar{r}_{j(h)}$ on incumbencies,

(b) Eviction frictions: $e_{h,t} \in \{0,1\}$ and if $e_{h,t} = 1$ then $\phi^{\text{evict}}_{j(h)}$ is incurred and legal feasibility holds.

The *first-best* benchmark drops (12) and legal frictions, while the *second-best* imposes them. Information is the same as in the baseline: shocks have finite support and the transition kernel is time-homogeneous. In first-best, $\mathcal{F}(s_t)$ omits (12) and legal feasibility on $e_{h,t}$; in second-best, $\mathcal{F}(s_t)$ includes them.

Recursive formulation. Let s_t collect

$$\big(\{q_{h,t}\}_h, \text{ permit/status}, \text{ incumbency/rent status}, \{\theta_{i,h,t}\}_{i,h}, \text{ other discrete summaries}\big)$$

The planner's Bellman equation is

$$V^{\text{planner}}(s_t) = \max_{\{r_t, m_t, e_t, \mu_t\} \in \mathcal{F}(s_t)} \Big\{ W_t^{\text{tenant}}(s_t, \mu_t, r_t) + W_t^{\text{landlord}}(s_t, r_t, m_t, e_t) + \delta \, \mathbb{E}\Big[V^{\text{planner}}(s_{t+1}) \,\Big|\, s_t, r_t, m_t, e_t, \mu_t \Big] \Big\}, \tag{13}$$

where $\mathcal{F}(s_t)$ encodes (10)–(12) and the law of motion (11). In the discretised baseline, all sets are finite; existence of an optimal stationary policy follows from standard discounted dynamic programming results (contraction of the evaluation operator on a finite state space).[1]

Interpretation and role. The planner's problem provides the efficient frontier (first-best) or the constrained frontier (second-best) for assignments, investment, evictions, and—when permitted—rents. Comparing V^{planner} with the decentralised MPE value V^{MPE} yields the welfare gap

[1] See Puterman (1994, Ch. 6) and Stokey et al. (1989).

in Definition 2 and underpins the channel decomposition reported below. Because the objective nets out transfers, differences arise from *real* distortions: mismatch, moving costs, under-maintenance/quality decay, and spillovers from regulatory heterogeneity.

3.5.2 Comparison to the decentralised equilibrium

Having defined the planner's allocation in Section 3.5.1 and welfare in Section 3.5, we compare it to outcomes in the decentralised Markov-perfect equilibrium (MPE). Because rents cancel in the planner's objective (transfer neutrality), differences are driven by *real* distortions: mismatch, mobility frictions, investment/quality, and spatial spillovers.

Conceptual framework. In the decentralised MPE, tenants and landlords solve separate dynamic problems given the same technological and regulatory primitives as in Sections 2–3. Tenants choose stay/move via a dynamic discrete choice, taking posted rents, qualities, and moving costs as given. Landlords choose rents, maintenance, and evictions to maximise discounted profits subject to rent caps, eviction frictions, and permit delays. The state evolves from decentralised actions and shocks. By contrast, the planner jointly chooses allocations, prices (when permitted), and actions to maximise discounted welfare, internalising turnover, disinvestment incentives, and cross-jurisdiction spillovers.

Sources of inefficiency in the decentralised equilibrium. Relative to the planner, the decentralised allocation exhibits five systematic gaps:

1. **Underinvestment in quality.** Binding caps and eviction frictions suppress private returns to maintenance, yielding m below planner-optimal. The planner values quality both through tenant utility and market consequences, selecting higher m (cf. Proposition 3 and Appendix D).

2. **Tenant–unit misallocation.** Sticky rents, search frictions, and incomplete information generate mismatches relative to assignment on full state information. The planner's $\mu_{i,h,t}$ maximises match surplus given real constraints.

3. **Excessive rent dispersion.** Market power reassigns markups toward unconstrained/new leases while capping incumbents, creating inefficient price variation across similar units.

The planner smooths prices (or implements Ramsey/MC pricing in the second best), improving sorting (cf. Proposition 1 and Appendix B).

4. **Inefficient turnover.** Turnover can be excessively high (strategic renovictions) or too low (lock-in from moving costs). The planner internalises the reallocation value and balances it against moving costs and match quality (cf. Proposition 3).

5. **Spatial misallocation of capital.** Regulatory heterogeneity induces reallocation toward lightly regulated jurisdictions, altering quality and prices elsewhere. The planner coordinates investment spatially to maximise quality-adjusted output, avoiding arbitrage externalities (cf. Proposition 2 and Appendix C).

Welfare gap and components. Let $\Delta W = V^{\text{planner}}(s_0) - V^{\text{MPE}}(s_0)$ (Definition 2). Using the transfer-neutral objective, we attribute

$$\Delta W = \underbrace{\Delta W^{\text{investment}}}_{\text{under-maintenance/quality}} + \underbrace{\Delta W^{\text{allocation}}}_{\text{mismatch}} + \underbrace{\Delta W^{\text{mobility}}}_{\text{inefficient turnover}} + \underbrace{\Delta W^{\text{markups}}}_{\text{pricing dispersion}} + \underbrace{\Delta W^{\text{spillover}}}_{\text{cross-jurisdiction reallocation}}.$$

A consistent additive decomposition is implemented via the Shapley scheme in Definition 3, yielding the identity in Proposition 5. Each component is computed by neutralising exactly one channel while solving the same fixed-point problem on the same finite state space (Lemma 2).

Discussion. The magnitude of ΔW depends on policy tightness (caps, eviction frictions), the degree of market power and dispersion in unconstrained markups (Proposition 1), portfolio scope across jurisdictions (Proposition 2), and moving costs shaping lock-in and renoviction incentives (Proposition 3). The spillover term $\Delta W^{\text{spillover}}$ can be positive or negative at the jurisdiction level (e.g. B may gain quality yet face higher rents), but is typically negative in aggregate when regulatory gradients induce misallocation. Calibrated simulations in Section 4 quantify each component and evaluate counterfactuals (Appendix I provides MCS conditions supporting the qualitative signs).

3.5.3 Decomposition of welfare effects

Lemma 2 (Well-posed neutralisations)**.** *Fix the baseline finite environment with state space $\mathcal{S}$, feasible action sets $A_\ell(s)$, and time-homogeneous transition kernel $P_\theta(\cdot \mid s, a)$ indexed by a parameter*

vector θ. For any subset S of neutralised channels, construct E_S by parameter substitutions only, *i.e. replace θ by $\theta(S)$ (adjusting e.g. $(\tau, c^{\text{move}}, \{\Delta\bar{r}_j\}_j, \{\phi_j^{\text{evict}}\}_j$, harmonisation flags) without changing $\mathcal{S}$, $\{A_\ell(s)\}$, or the support of P_θ. Then all E_S share the same $\mathcal{S}$, $A_\ell(s)$, and the same transition support. Hence welfare levels $W(E_S)$ are apples-to-apples and Shapley comparisons are well-posed.*

Implementation note. Each $W(E_S)$ is computed by solving the *same* fixed-point problem on the *same* finite state space with parameters $\theta(S)$. Neutralisations adjust primitives (e.g. set $\Delta\bar{r}_j$ to harmonised values) but never enlarge action sets or alter the Markov kernel's support.

Definition 3 (Shapley-style welfare decomposition)**.** *Let $\mathcal{K} = \{$investment, allocation, mobility, markups, spillove denote the set of distortion channels. For any subset $S \subseteq \mathcal{K}$, let $V(S)$ be total welfare when the channels in S are neutralised via parameter substitutions to planner-equivalent primitives, and the remaining channels act as in the decentralised baseline. We anchor*

$$V(\varnothing) = V^{\text{MPE}}(s_0), \qquad V(\mathcal{K}) = V^{\text{planner}}(s_0).$$

For any permutation π of $\mathcal{K}$ and any $k \in \mathcal{K}$, let $S_{\pi,k}$ be the set of channels preceding k in π and define the marginal contribution

$$\Delta W_k(\pi) \;:=\; V(S_{\pi,k} \cup \{k\}) - V(S_{\pi,k}).$$

The Shapley contribution *of k is*

$$\Delta W_k \;:=\; \frac{1}{|\Pi|} \sum_{\pi \in \Pi} \Delta W_k(\pi),$$

the average over all permutations Π of $\mathcal{K}$ (Shapley, 1953).

Proposition 5 (Exact additive decomposition)**.** *With ΔW as in Definition 2 and Shapley contributions $\{\Delta W_k\}_{k \in \mathcal{K}}$ as in Definition 3, we have the identity*

$$\Delta W \;=\; \sum_{k \in \mathcal{K}} \Delta W_k.$$

Construction details and the neutrality maps $V(S)$ are provided in Appendix J.

Component definitions (neutralisations). All neutralisations are implemented as parameter substitutions that preserve $\mathcal{S}$, $A_\ell(s)$, and the support of P_θ.

Investment. On regulated units, replace $C(m)$ by $C(m) - \tau^\star m$ with $\tau^\star$ calibrated so that the stationary best response in m coincides with the planner's m (other frictions unchanged). This isolates under-maintenance/quality distortions.

Allocation. Replace the decentralised matching kernel by its planner-equivalent assignment rule $\mu^{\text{planner}}_{i,h,t}(s_t)$ encoded as parameters of the same finite matching kernel (same support), holding prices and policies fixed. This isolates mismatch.

Mobility. Set c^{move} to the planner's effective level (first-best: $c^{\text{move}} \to 0$) with all other channels active.

Markups. Replace unconstrained price-setting by the planner's Ramsey/MC pricing schedule $r_{h,t} = r^{\text{planner}}(q_{h,t})$ for new leases while preserving incumbent-cap rules; this collapses inefficient price dispersion conditional on q.

Spillover. Harmonise $(\Delta\bar{r}_j, \phi_j^{\text{evict}})$ across j to planner-equivalent parameters, eliminating regulatory gradients that induce cross-jurisdiction reallocation. Portfolio sizes and action sets are unchanged.

Diagnostics outside the Shapley sum (optional). For reporting, we compute two ancillary statistics that do *not* enter the Shapley identity (to avoid double counting):

Static protection (incumbent rent stability). The incumbent surplus gain under regulation,

$$\Delta W^{\text{protect}} := \mathbb{E}\big[\text{tenant surplus}^{\text{MPE}}_{\text{incumbent}} - \text{tenant surplus}^{\text{planner}}\big],$$

which captures intended protections but overlaps with allocation/mobility channels.

Dynamic quality distortion. The welfare loss attributable to quality paths,

$$\Delta W^{\text{quality}} := \mathbb{E}\left[\sum_{t\geq 0} \delta^t \big(V_q^{\text{planner}}(t) - V_q^{\text{MPE}}(t)\big)\right],$$

which is accounted for within the investment channel in the Shapley scheme but is useful for presentation.

3.5.4 Incidence analysis: who gains and who loses

The welfare decomposition in Section 3.5.3 isolates *sources* of the planner–decentralised gap (Definition 2, Proposition 5). We now disaggregate *incidence* across agent types and space. Throughout,

all values are expected, discounted objects computed from the common initial state s_0 under either the planner or the decentralised stationary MPE (Lemma 1).

Tenant-level incidence. Let $V_i^{\text{planner}}(s_0)$ and $V_i^{\text{MPE}}(s_0)$ be tenant i's lifetime value under the planner and under the decentralised equilibrium, respectively (Section 3). Partition tenants into *incumbents* (matched at $t = 0$) and *entrants/mobile* (unmatched or moving at $t = 0$). Define individual incidence measures

$$\Delta U_i^{\text{inc}} := V_i^{\text{MPE}}(s_0) - V_i^{\text{planner}}(s_0), \qquad \Delta U_i^{\text{entry}} := V_i^{\text{planner}}(s_0) - V_i^{\text{MPE}}(s_0).$$

Incumbents. With binding caps and eviction frictions, the decentralised environment raises incumbent surplus via the static protection channel (Section 3.5.3); hence, typically $\Delta U_i^{\text{inc}} \geq 0$, with strict > whenever the cap binds on their unit.[1] *Entrants/mobile.* These tenants face higher prices on the unconstrained margin and lower average quality due to underinvestment and spillovers; thus $\Delta U_i^{\text{entry}} \geq 0$, with strict > when either (i) the new-lease markup rises (Proposition 1) or (ii) quality deteriorates along their search path (Proposition 3).

Landlord-level incidence. Let $\Pi_\ell^{\text{planner}}(s_0)$ and $\Pi_\ell^{\text{MPE}}(s_0)$ denote landlord ℓ's discounted net return. Two dimensions matter: portfolio scope and regulatory exposure.

- *Large, multi-jurisdiction portfolios.* With heterogeneity in caps and frictions, these landlords reallocate effort/investment to lenient jurisdictions and extract higher markups on unconstrained leases (Propositions 2 and 1). Typically $\Pi_\ell^{\text{MPE}} - \Pi_\ell^{\text{planner}} \geq 0$.
- *Small, locally constrained portfolios.* With limited reallocation ability and high eviction frictions, these landlords bear investment distortions and may delay turnover (Proposition 3). Typically $\Pi_\ell^{\text{planner}} - \Pi_\ell^{\text{MPE}} \geq 0$.

Jurisdictional incidence and spatial spillovers. Let W_j denote *local* welfare in jurisdiction j: the sum of tenant utilities for households housed in j plus net landlord returns from units located in j (as in Section 3.5.1, using $j(h)$ for unit location). With fragmented policy—a stricter cap only

[1] The sign follows from the rent path $r_{\text{inc},t}$ under caps and the option value of staying; see the price-dispersion and cap mechanics behind Proposition 1.

in A, unconstrained B—Proposition 2 implies that tightening in A reduces m^A and reallocates effort to B ($m^B \uparrow$), which weakly raises r^B. Therefore,

$$\Delta W_A \;:=\; W_A^{\mathrm{MPE}} - W_A^{\mathrm{planner}} \;\gtrless\; 0 \quad \text{(ambiguous: protection vs. disinvestment)},$$

$$\Delta W_B \;:=\; W_B^{\mathrm{MPE}} - W_B^{\mathrm{planner}} \;<\; 0$$

whenever the cap differential is active on a set of positive measure and $D_q^B > 0$ (Lemma C.2). The sign in A depends on the balance between incumbent protection and quality decay; the sign in B is negative when the GE price effect dominates.

Temporal incidence (intergenerational). Write $\mathbb{E}\big[V_{i,t}^{\mathrm{MPE}} \,\big|\, i \text{ is a new entrant at } t\big]$ for an entering cohort's expected value under the stationary decentralised equilibrium. With persistent investment distortions (Section 3) and price spillovers (Proposition 2), a tighter cap regime (lower $\Delta\bar{r}A$) expands the renoviction region (Proposition 3) and the unconstrained markup (Proposition 1), implying a downward drift for future entrants:

$$\mathbb{E}\big[V_{i,t+1}^{\mathrm{MPE}}\big] \;\le\; \mathbb{E}\big[V_{i,t}^{\mathrm{MPE}}\big] \quad \text{for successive entrant cohorts } t, t{+}1 \text{ under a tighter cap path.}$$

Thus, static protection for incumbents trades off against dynamic losses borne by later cohorts.

Distributional effects of regulation across agent types (relative to planner benchmark)

Group	Typical effect	Mechanism / link to results
Incumbent tenants	Gain	Rent caps / eviction frictions raise private surplus (static protection; Prop. 1)
Entrant / mobile tenants	Lose	Higher new-lease markups and lower quality from underinvestment (Props. 1, 3)
Large multi-jurisdiction landlords	Gain	Reallocate effort to lenient areas; exploit unconstrained margins (Props. 2, 1)
Small/local landlords	Lose	Limited reallocation; higher exposure to caps and frictions (Prop. 3)
Regulated jurisdiction (A)	Mixed	Incumbent gains vs. quality/investment losses (Props. 1, 3)
Unregulated jurisdiction (B)	Lose	Inflow of effort raises q^B but GE prices rise; entrants crowded out (Prop. 2)
Future entrants	Lose	Cumulative quality decay and persistent price dispersion (Props. 3, 1)

Measurement in simulations. For any subgroup $\mathcal{I}$, report $\Delta U(\mathcal{I}) = \frac{1}{|\mathcal{I}|}\sum_{i\in\mathcal{I}}\left(V_i^{\text{planner}} - V_i^{\text{MPE}}\right)$, and for any jurisdiction j, $\Delta W_j = W_j^{\text{planner}} - W_j^{\text{MPE}}$, computed under the same initial state s_0 and shock seeds (common random numbers) to isolate policy effects (Appendix E).

3.5.5 Policy evaluation and counterfactuals

We evaluate four regulatory counterfactuals designed to reduce the decentralised–planner gap (Definition 2), using the same finite state space, action sets, and transition support as in the baseline (Lemma 2). Each counterfactual modifies *parameters only* (no change to feasible actions or Markov kernel support), consistent with the neutrality principle used in the welfare decomposition (Proposition 5). Comparative static signs follow the MCS results (Proposition 4; Appendix I; Topkis, 1998).

Baseline regime (status quo). The baseline is the decentralised stationary MPE under heterogeneous caps $\Delta\bar{r}j$, eviction frictions $\phi_j^{\text{evict}} > 0$, and permit delays $\psi_j^{\text{permit}} \geq 0$ (Section 3). Caps bind for incumbents; new leases are unconstrained (Proposition 1). Under this regime, investment incentives are distorted (under $C(\cdot)$), spillovers arise from regulatory heterogeneity (Proposition 2), and renoviction risk can be positive on a nonempty state set (Proposition 3).

Counterfactual 1: Quality-linked maintenance subsidies (investment channel). On regulated units, replace the landlord cost $C(m)$ by $C(m) - \tau\, m$ with $\tau > 0$.[1] Formally, the per-unit Bellman flow term becomes $\pi_{h,t} - \big(C(m_{h,t}) - \tau m_{h,t}\big) - \phi_{j(h)}^{\text{evict}} e_{h,t}$. Under the MCS assumptions (Appendix I), best responses are increasing in τ; thus m (and hence q) weakly rises on regulated units (Proposition 4). This counterfactual directly lowers the $\Delta W^{\text{investment}}$ component by mitigating undermaintenance.

Fiscal accounting. We include the resource cost of the subsidy in welfare via a government budget term $G_t = -\sum_{h\in\mathcal{H}^{\text{reg}}} \tau\, m_{h,t}$ and assume it is financed by a lump-sum transfer that does not distort private choices (the transfer cancels in $W_t^{\text{tenant}} + W_t^{\text{landlord}} + G_t$). This ensures apples-to-apples comparisons with the baseline.

Counterfactual 2: Targeted eviction/mobility vouchers (mobility channel). Hold ϕ_j^{evict} fixed and reduce treated tenants' effective moving cost from c^{move} to $c^{\text{move}} - \upsilon$ (with $\upsilon \in (0, c^{\text{move}}]$) upon eviction notice or certified relocation.[2] Under MCS, lower c^{move} weakly increases efficient turnover on the relevant margins (Proposition 4), attenuating the $\Delta W^{\text{mobility}}$ distortion and reducing incentives for strategic undermaintenance that rely on tenant lock-in (complements Proposition 3).

Fiscal accounting. Add $G_t = -\sum_{i\in\mathcal{I}_t} \upsilon$ where $\mathcal{I}_t$ is the set of voucher recipients, financed lump-sum as above.

Counterfactual 3: Interjurisdictional policy coordination (spillover channel). Harmonise spatial frictions by setting $\Delta\bar{r}A = \Delta\bar{r}B = \bar{\Delta r}$ and $\phi_A^{\text{evict}} = \phi_B^{\text{evict}} = \bar{\phi}$. This removes the regulatory gradient that drives reallocation (Proposition 2 and Lemma C.2). With no cross-jurisdiction wedge,

[1] As in the Investment neutralisation in Section 3.5.3: we change the cost slope, not the action set.

[2] Implementation leaves landlord problem unchanged except through equilibrium occupancy/matching; the policy acts on the DDC component in the tenant Bellman.

optimal effort is not diverted to B when A tightens, and the GE price effect in B is dampened, shrinking $\Delta W^{\text{spillover}}$. Political feasibility aside, even partial alignment (reducing variance of $\Delta \bar{r} j$ and ϕ_j^{evict}) produces the same sign.

Counterfactual 4: Structured two-tier regulation with enforcement (markup/reno channel). Codify the incumbent/new-lease distinction with (i) explicit incumbent caps $\Delta \bar{r} j$, (ii) transparent contract status at signing and renewal, and (iii) enforcement against informal evictions (raising the effective cost of churn manipulation). This preserves incumbent protection while dampening rent-shifting incentives documented in Proposition 1 and narrows dispersion created by strategic resets (conditional on $\Delta \bar{r} j$ not being too tight; see Proposition 1 and Proposition 3). The primary effect is a reduction in $\Delta W^{\text{markups}}$ and in the renoviction region.

Evaluation protocol and welfare comparison. Let $V^{\text{CF}_k}(s_0)$ denote aggregate welfare under counterfactual $k \in \{1, 2, 3, 4\}$ from the common initial state s_0. Define gains

$$\Delta W^{\text{CF}_k} := V^{\text{CF}_k}(s_0) - V^{\text{MPE}}(s_0),$$

and, when desired, decompose ΔW^{CF_k} into channels using the Shapley procedure in Definition 3. We compute all counterfactual equilibria on the same discretisation and with common random numbers (Appendix E) to isolate pure policy effects. Signs predicted by the theory are:

CF1 ($\tau \uparrow$): $m \uparrow,\ q \uparrow \Rightarrow\ \Delta W^{\text{investment}} \downarrow$ (Prop. 4; Topkis 1998).

CF2 ($c^{\text{move}} \downarrow$): efficient turnover $\uparrow \Rightarrow\ \Delta W^{\text{mobility}} \downarrow$ (Prop. 4).

CF3 (harmonise $\Delta \bar{r} j, \phi_j^{\text{evict}}$): reallocation $\downarrow,\ r^B \downarrow \Rightarrow\ \Delta W^{\text{spillover}} \downarrow$ (Prop. 2).

CF4 (two-tier+enforcement): rent-shifting $\downarrow$, renoviction set $\downarrow \Rightarrow\ \Delta W^{\text{markups}} \downarrow$ and total ΔW improves

subject to the caveat from Proposition 1 that extremely tight caps can widen dispersion on the unconstrained margin; pairing CF4 with CF1 avoids that case in practice.

Summary. Within our dynamic GE framework, policies that *raise returns to maintenance* (CF1), *lower mobility frictions* (CF2), and *remove spatial wedges* (CF3) robustly narrow the decentralised–planner gap without diluting incumbent protection. A structured two-tier regime with credible enforcement (CF4) curbs rent-shifting and renoviction incentives that drive dispersion and quality

decay. All four are implementable as parameter changes, preserving existence and comparability of stationary equilibria (Appendix A; Fink, 1964; Puterman, 1994; Stokey et al., 1989).

4 Results

Figure 2 summarizes the computational pipeline. We proceed left-to-right: primitives and policy define the baseline game, we solve for the stationary equilibrium, implement neutralisations and counterfactuals, then decompose ΔW and report incidence.

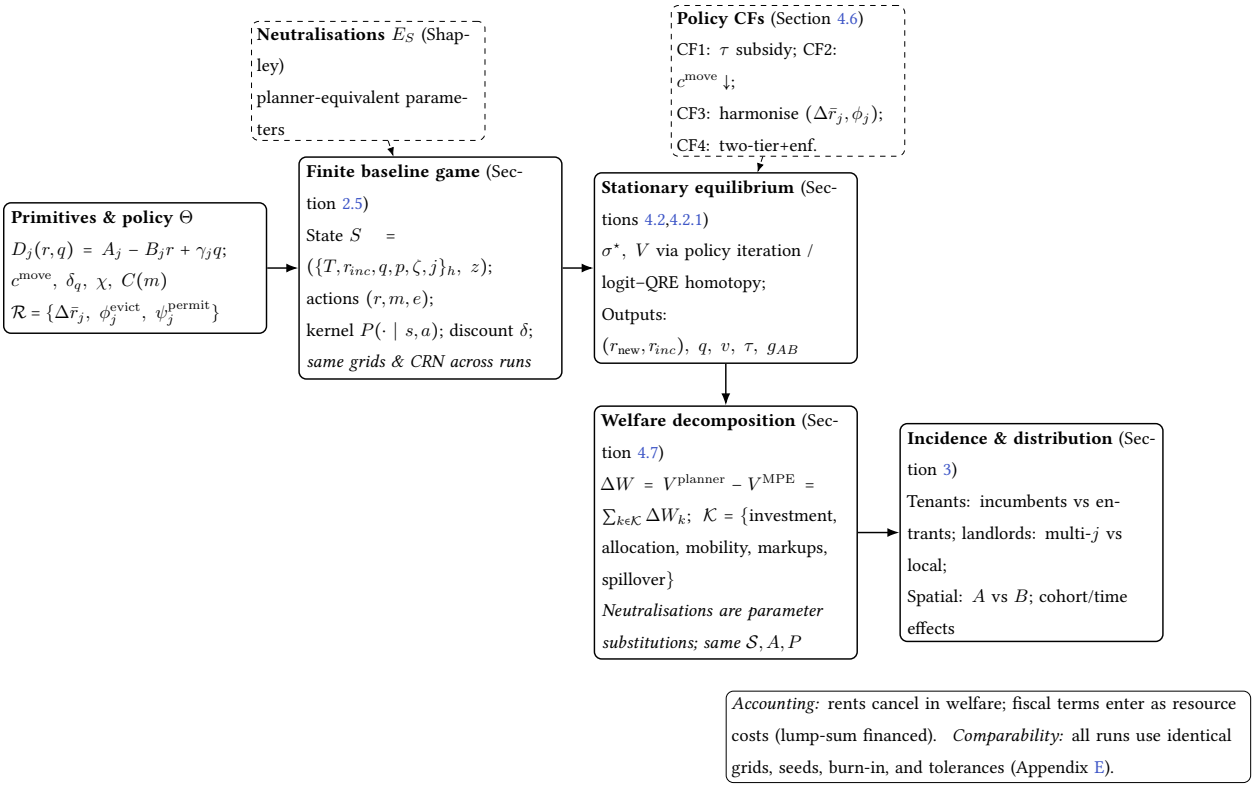


Figure 2: Model-to-results map. Primitives → finite game → stationary equilibrium; welfare is decomposed by Shapley (neutralisations), and policy counterfactuals shift components and incidence.

4.1 Baseline parameter vector and grids

We specify a baseline parameter vector Θ and finite grids consistent with the discretised baseline game in Section 2.5. These seeds initialise the equilibrium solver described in Appendix E;

estimates are then refined by calibration (Section 4.5). The definitive parameter values we use in computation are summarised in Appendix E, Table E10.

Primitives and policy. Let

$$\Theta = \Big(A_j,\, B_j,\, c^{\text{move}},\, \delta_q,\, \chi,\, \kappa_1,\, \kappa_2,\, \{\Delta\bar{r}j,\, \phi_j^{\text{evict}},\, \psi_j^{\text{permit}}\}_{j\in\{A,B\}},\, \delta\Big),$$

where $D_j(r,q) = A_j - B_j r + \gamma_j q$ is the reduced-form (residual) demand for new leases in jurisdiction j, $c^{\text{move}} > 0$ is the switching cost, $\delta_q \in (0,1)$ is quality depreciation, $\chi > 0$ is maintenance effectiveness, and $C(m) = \kappa_1 m + \frac{1}{2}\kappa_2 m^2$ is the convex maintenance cost. Policy primitives are the rent-increase cap $\Delta\bar{r}j$, the eviction friction ϕ_j^{evict}, and the permit delay ψ_j^{permit}. The discount factor is $\delta \in (0,1)$.

Microfoundation for the linear residual demand

Our calibration uses the reduced–form $D_j(r,q) = A_j - B_j r + \gamma_j q$, which is the first–order (local) approximation to the residual demand implied by a standard dynamic discrete choice (DDC) model with i.i.d. Type-I extreme value shocks.[1]

Setup. Let a mass M_j of searchers in j choose among alternatives $\mathcal{C}_j$ (units and an outside option). Per-period utility for alternative $h \in \mathcal{C}_j$ is

$$U_i(h) = \mu_j(h) + \xi(q_h) - \beta\, r_h + \varepsilon_{ih}, \qquad \varepsilon_{ih} \overset{\text{i.i.d.}}{\sim} \text{Type-I EV}.$$

Under stationarity, the continuation value embedded in $\mu_j(h)$ is locally invariant to a marginal change in (r_h, q_h) on a single alternative. Hence the share of h is logit:

$$s_j(h) = \frac{\exp\big(\mu_j(h) + \xi(q_h) - \beta r_h\big)}{\sum_{k\in\mathcal{C}_j} \exp\big(\mu_j(k) + \xi(q_k) - \beta r_k\big)}.$$

Define residual demand faced by h as $D_j^{\text{DDC}}(r_h, q_h) = M_j\, \Lambda_j\, s_j(h)$, where $\Lambda_j \in (0,1]$ captures dynamic attenuation from switching costs/turnover in the DDC (a reduced pass-through from per-period shares to lease flows; cf. Appendix F).

[1] See McFadden (1974); Train (2009) for logit demand and Rust (1987); Hotz and Miller (1993) for DDC foundations. The linearisation result is stated formally in Lemma F.3 (Appendix F).

Local linearisation. Fix a baseline (r_0, q_0) and write $s_0 = s_j(h)|(r_0, q_0)$. Differentiating the logit share yields

$$\frac{\partial s_j(h)}{\partial r_h}\Big|_0 = -\beta\, s_0(1-s_0), \qquad \frac{\partial s_j(h)}{\partial q_h}\Big|_0 = \xi'(q_0)\, s_0(1-s_0).$$

Thus D_j^{DDC} is differentiable and its first–order Taylor expansion around (r_0, q_0) implies the mapping (Lemma F.3):

$$B_j = \beta\, M_j \Lambda_j\, s_0(1-s_0), \qquad \gamma_j = \xi'(q_0)\, M_j \Lambda_j\, s_0(1-s_0), \qquad A_j = D_j^{\mathrm{DDC}}(r_0, q_0) + B_j r_0 - \gamma_j q_0,$$

so that $D_j(r, q) = A_j - B_j r + \gamma_j q$ is a first–order accurate residual demand.

Finite grids (state–action discretisation). We discretise rents $\mathcal{G}_j^r$, maintenance $\mathcal{M}$, and quality $\mathcal{Q}$ to define the baseline finite game used for MPE computation (Appendix E).

(i) Rent grids. For each j, let $\mathcal{G}_j^r = \{r_j^{\min}, r_j^{\min} + \Delta r_j, \ldots, r_j^{\max}\}$ with N_j^R points. Choose $[r_j^{\min}, r_j^{\max}]$ to bracket observed mean rents and the unconstrained monopoly price $r^{\mathrm{mono}} = \frac{A_j + B_j\, MC}{2B_j}$ (holding q fixed). Set Δr_j so that the cap increment $\Delta \bar{r} j$ is an integer multiple of Δr_j (rent-cap law in Section 2.5); this avoids stochastic rounding at the cap.

(ii) Maintenance grid. Let $\mathcal{M} = \{0, 1, \ldots, \bar{m}\}$ with small $\bar{m}$ (e.g., 5 or 10). Pick $\bar{m}$ so that $\chi\, \bar{m}$ offsets at least one period of depreciation at moderate q; i.e., $\chi\, \bar{m} \gtrsim \delta_q\, \bar{q}$ (see the quality grid below). This ensures both "maintain" and "under-maintain" are interior choices in the DP.

(iii) Quality grid. Let $\mathcal{Q} = \{q^1, \ldots, q^{\bar{q}}\}$, evenly spaced on $[q^{\min}, q^{\max}]$. Anchor bounds to steady states under $m = 0$ and $m = \bar{m}$:

$$q^{\mathrm{ss}}(m) \approx \frac{\chi\, m}{\delta_q} \quad \Rightarrow \quad q^{\min} \le q^{\mathrm{ss}}(0) = 0, \qquad q^{\max} \gtrsim q^{\mathrm{ss}}(\bar{m}) = \frac{\chi\, \bar{m}}{\delta_q}.$$

Choose $\bar{q}$ (e.g., 15–25) to balance accuracy and state size. The discrete transition $\Gamma(q, m; j)$ maps (q_t, m_t) to the nearest grid point in $\mathcal{Q}$, with permit delays encoded via Γ or via a finite permit-status variable if active (Appendix E).

Normalisations and units. Measure rents per period (e.g., monthly); interpret c^{move} and ϕ_j^{evict} in the same units. Normalise q so that q=1 is "baseline" quality; then χ and δ_q are interpretable. If desired, absorb a constant marginal operating cost into A_j (setting $MC \equiv 0$ in the linear-demand normalisation).

Grid size guidance. Keep $|\mathcal{G}_j^r| \times |\mathcal{M}| \times |\mathcal{Q}|$ modest to ensure fast value evaluation and tight Bellman residuals. A practical baseline is

$$|\mathcal{G}_A^r| = |\mathcal{G}_B^r| \approx 31 \text{ to } 51, \quad |\mathcal{M}| \approx 6 \text{ to } 11, \quad |\mathcal{Q}| \approx 15 \text{ to } 25,$$

with sparse transitions for Γ and for the matching kernel. Sensitivity to grid coarseness is reported in Appendix E.

4.2 Equilibrium computation on the baseline game

We compute a stationary equilibrium of the finite baseline game (Section 2.5) on a fixed state space and action sets shared with all neutralisations and counterfactuals (Lemma 2). Existence of a stationary MPE in the finite, discounted game is established in Appendix A (Proposition A.1; Fink, 1964). For a fixed stationary profile σ, the Bellman evaluation operator T_σ on value vectors is a δ–contraction on $(\mathbb{R}^{|\mathcal{S}|}, \|\cdot\|_\infty)$, hence V^σ is unique (Puterman, 1994; Stokey et al., 1989, Chs. 4–6). We iterate stationary (smoothed) best responses state by state, evaluating V^σ at each step.

Algorithm 1 Stationary MPE on the finite baseline

Inputs: finite state set $\mathcal{S}$; actions $\{A_\ell(s)\}$; discount $\delta \in (0,1)$; kernel $P(\cdot \mid s,a)$; payoffs $\pi_\ell(s,a)$; smoothing $\eta \in (0,1]$; tolerances $\varepsilon_\pi, \varepsilon_V > 0$.

Initialise a stationary profile $\sigma^{(0)} = (\sigma_\ell^{(0)})_\ell$ (uniform or myopic).

for $k = 0, 1, 2, \ldots$ **do**

  Evaluate values under $\sigma^{(k)}$: solve

$$V_\ell^{(k)}(s) = \mathbb{E}_{a\sim\sigma^{(k)}(s)}\Big[\pi_\ell(s,a) + \delta\, \mathbb{E}_{s'\sim P(\cdot|s,a)} V_\ell^{(k)}(s')\Big], \quad \forall s \in \mathcal{S},\ \forall \ell,$$

  by sparse linear solve of $(I - \delta P_{\sigma^{(k)}})V^{(k)} = b_{\sigma^{(k)}}$ or value iteration to sup–norm tolerance $< \varepsilon_V$.

  for each landlord ℓ and state $s \in \mathcal{S}$ **do**

    Compute Q–values: for each $a_\ell \in A_\ell(s)$,

$$Q_\ell^{(k)}(s, a_\ell) = \mathbb{E}_{a_{-\ell}\sim\sigma_{-\ell}^{(k)}(s)}\Big[\pi_\ell\big(s, (a_\ell, a_{-\ell})\big) + \delta\, \mathbb{E}_{s'\sim P(\cdot|s,(a_\ell,a_{-\ell}))} V_\ell^{(k)}(s')\Big].$$

    Best response (mixed allowed): pick $\mathrm{BR}_\ell^{(k)}(s) \in \arg\max_{\alpha\in\Delta(A_\ell(s))} \sum_{a_\ell} \alpha(a_\ell)\, Q_\ell^{(k)}(s, a_\ell)$ (lexicographic tie–break).

    Smoothed policy update: $\sigma_\ell^{(k+1)}(s) \leftarrow (1-\eta)\,\sigma_\ell^{(k)}(s) + \eta\,\mathrm{BR}_\ell^{(k)}(s)$.

  end for

  Convergence check: if $\max_{\ell,s} \|\sigma_\ell^{(k+1)}(s) - \sigma_\ell^{(k)}(s)\|_1 < \varepsilon_\pi$ and

$$\max_{\ell,s} \big|V_\ell^{(k)}(s) - (T_{\sigma^{(k)}} V_\ell^{(k)})(s)\big| < \varepsilon_V,$$

  stop and return $\sigma^\star = \sigma^{(k+1)}$.

end for

Computation guarantee via logit–QRE and homotopy. Because smoothed best–response dynamics in general–sum stochastic games need not be contractive in policies, we compute a stationary logit QRE at temperature $\tau > 0$—which exists by Brouwer on the product of simplices—and then trace a continuation in τ towards zero (Appendix E, Proposition E.2). The logit map

$$\sigma_{\ell,\tau}(a_\ell \mid s) \;\propto\; \exp\big(Q_\ell^{\sigma_\tau}(s, a_\ell)/\tau\big)$$

is single–valued and continuous; with the contractive evaluation step, fixed points are found by policy iteration at each τ. We warm–start along a grid $\tau \in \{\tau_0, \rho\tau_0, \rho^2\tau_0, \ldots, \tau_{\min}\}$ with $\tau_0 \in [0.5, 1]$, $\rho \in [0.5, 0.8]$, and $\tau_{\min} \in [10^{-3}, 10^{-2}]$. The reported policy at small τ closely approximates a stationary MPE and is used to compute all results. Diagnostics (policy distance and Bellman residual) are recorded at each τ.

4.2.1 Regularised stationary equilibrium (logit QRE)

For completeness, we restate the stationary logit–QRE used in computation.

Definition (stationary logit–QRE). A stationary profile $\sigma_\tau = \{\sigma_{\ell,\tau}\}_{\ell \in \mathcal{L}}$ is a logit–QRE at temperature $\tau > 0$ if, for every landlord ℓ, state $s \in \mathcal{S}$, and action $a_\ell \in A_\ell(s)$,

$$\sigma_{\ell,\tau}(a_\ell \mid s) = \frac{\exp\big(Q_\ell^{\sigma_\tau}(s, a_\ell)/\tau\big)}{\sum_{\tilde{a}_\ell \in A_\ell(s)} \exp\big(Q_\ell^{\sigma_\tau}(s, \tilde{a}_\ell)/\tau\big)},$$

where $Q_\ell^{\sigma_\tau}$ is defined by

$$Q_\ell^{\sigma_\tau}(s, a_\ell) = \mathbb{E}_{a_{-\ell} \sim \sigma_{\tau,-\ell}(s)}\Big[\pi_\ell\big(s, (a_\ell, a_{-\ell})\big) + \delta\, \mathbb{E}_{s' \sim P(\cdot \mid s, (a_\ell, a_{-\ell}))} V_\ell^{\sigma_\tau}(s')\Big],$$

and $V_\ell^{\sigma_\tau}$ solves $V_\ell^{\sigma_\tau} = T_{\sigma_\tau} V_\ell^{\sigma_\tau}$.

Existence and computation. For each fixed $\tau > 0$, the composite map $\Phi_\tau : \sigma \mapsto \mathrm{SoftMax}_\tau(Q^\sigma)$ is continuous on the compact product of simplices, so a stationary QRE exists by Brouwer. We compute it by fixed–point iteration with optional relaxation:

Algorithm 2 τ–homotopy QRE policy iteration

Inputs: $\mathcal{S}, \{A_\ell(s)\}, \delta, P, \pi_\ell$; $\tau_0 > 0$; $\tau_{\min} > 0$; $\rho \in (0,1)$; step $\eta \in [0.3, 1]$; tolerances $\varepsilon_\pi, \varepsilon_V$.
Initialise $\sigma^{(0)}$ (uniform). Set $\tau \leftarrow \tau_0$.
while $\tau \geq \tau_{\min}$ **do**
    for $k = 0, 1, 2, \ldots$ **do**
        Evaluate $V^{(k)}$ by $(I - \delta P_{\sigma^{(k)}})V^{(k)} = b_{\sigma^{(k)}}$ or value iteration.
        Form $Q_\ell^{(k)}(s, a_\ell)$ against $\sigma_{-\ell}^{(k)}$.
        Set $\widetilde{\sigma}_\ell(a_\ell \mid s) \propto \exp(Q_\ell^{(k)}(s, a_\ell)/\tau)$ for all (ℓ, s, a_ℓ).
        Update $\sigma^{(k+1)} \leftarrow (1 - \eta)\sigma^{(k)} + \eta\widetilde{\sigma}$.
        If $\max_{\ell,s} \|\sigma_\ell^{(k+1)}(s) - \sigma_\ell^{(k)}(s)\|_1 < \varepsilon_\pi$ and $\max_{\ell,s} |V_\ell^{(k)}(s) - (T_{\sigma^{(k)}} V_\ell^{(k)})(s)| < \varepsilon_V$, **break**.
    end for
    Decrease $\tau \leftarrow \rho\tau$; warm–start next outer loop at current $\sigma^{(k+1)}$.
end while
Return $\sigma^\star$ at $\tau \simeq \tau_{\min}$ (optionally a final η=1 best–response step).

Implementation notes and diagnostics. (i) Value evaluation uses sparse linear algebra or sup–norm value iteration; (ii) Relaxation ($\eta \in (0, 1]$) mitigates policy cycling; (iii) Lexicographic tie–break ensures measurability; (iv) We store policy distance $d^{(k)} = \|\sigma^{(k+1)} - \sigma^{(k)}\|_\infty$, Bellman residual $\rho^{(k)} = \|V^{(k)} - T_{\sigma^{(k)}} V^{(k)}\|_\infty$, iteration counts, and wall–clock time at each τ; (v) All baseline, neutralisation, and counterfactual solves use common random numbers and the same state/action support to preserve apples–to–apples comparisons (Lemma 2; Appendix E).

4.3 Empirical moments and model fit

We target a set of observed moments and assess in-sample fit of the stationary equilibrium.

Stationary measurement. All model moments are computed in the stationary distribution of the baseline equilibrium obtained via Algorithm 1 (or the small-τ QRE via Algorithm 2 when regularised). We use common random numbers (CRN), drop a burn-in, and then average over time and simulation replications. Let $\mathbb{E}[\cdot]$ denote the stationary cross-sectional expectation over units h and periods t (post–burn-in).

(1) Rent dispersion (incumbent vs. new).

$$\bar{r}^{\text{inc}} = \mathbb{E}\big[\, r_{inc,h,t} \,\big|\, \text{lease is incumbent at } (h,t) \,\big], \qquad \bar{r}^{\text{new}} = \mathbb{E}\big[\, r_{h,t} \,\big|\, \text{lease is newly signed at } (h,t) \,\big].$$

We target either the *difference* $\Delta r^{\text{disp}} = \bar{r}^{\text{new}} - \bar{r}^{\text{inc}}$ or the *ratio* $\rho^{\text{disp}} = \bar{r}^{\text{new}}/\bar{r}^{\text{inc}}$ (choose one and keep it fixed across runs to avoid duplicating information).

(2) Vacancy rate. With $T_{h,t} \in \{0,1\}$ the occupancy indicator,

$$v = \mathbb{E}[\, 1 - T_{h,t} \,].$$

(3) Turnover rate. Let $\mathbb{1}_{\text{churn}}(h,t) = 1$ if the occupant at (h,t) differs from $(h,t-1)$ and the unit is occupied in t, else 0. Define

$$\tau = \frac{\mathbb{E}[\, \mathbb{1}_{\text{churn}}(h,t) \,]}{\mathbb{E}[\, T_{h,t-1} \,]},$$

i.e., the annualised fraction of incumbent leases that turn over.

(4) Share with positive maintenance. For $m_{h,t} \in \mathcal{M}$,

$$s_m = \mathbb{E}[\, \mathbb{1}\{m_{h,t} > 0\} \,].$$

(5) Jurisdictional rent gap (A vs. B). Compute either for *new leases* (market rate comparison) or for *all occupied leases* (average paid rent). For new leases:

$$g_{AB}^{\text{new}} = \mathbb{E}\big[\, r_{h,t} \,\big|\, j(h) = B,\ \text{new} \,\big] - \mathbb{E}\big[\, r_{h,t} \,\big|\, j(h) = A,\ \text{new} \,\big].$$

For "all leases," use the *paid* rent $\tilde{r}_{h,t} := r_{h,t}$ if the lease is new at (h,t) and $\tilde{r}_{h,t} := r_{inc,h,t}$ if incumbent, conditional on $T_{h,t} = 1$.

Mapping to parameters. In SMM (Section 4.5), the vector $(\Delta r^{\text{disp}}, v, \tau, s_m, g_{AB})$ primarily loads on:

- Δr^{disp} (or ρ^{disp}): switching cost c^{move} jointly with cap tightness $\Delta \bar{r} j$ and price-setting curvature (Proposition 1).
- v and τ: demand slopes B_j (through residual demand), eviction frictions ϕ_j^{evict}, and search/mobility (Appendix F, Prop. 4).

- s_m: maintenance technology $(\chi, \kappa_1, \kappa_2)$ and policy frictions that affect incentives to invest (Proposition 3).
- g_{AB}: the cross-jurisdiction wedge $(\Delta\bar{r}A, \Delta\bar{r}B, \phi_A^{\text{evict}}, \phi_B^{\text{evict}})$ via spatial reallocation (Proposition 2).

Optional extra targets (for over-identification or weak rank). When identification is weak (see Appendix G), we add at most one of:

- **Evictions per 100 leases.** With $\mathbb{1}_{\text{evict}}(h,t) = 1$ for a successful eviction,
$$e \;:=\; 100 \times \frac{\mathbb{E}[\mathbb{1}_{\text{evict}}(h,t)]}{\mathbb{E}[T_{h,t}]},$$
which mainly loads on ϕ_j^{evict} (deterrence) and interacts with maintenance via renoviction incentives (Proposition 3).
- **Tenure distribution.** Let ten_h denote incumbent tenure (periods). Target one percentile
$$q_p^{\text{ten}} := \inf\{x : \mathbb{P}(\text{ten}_h \le x) \ge p\}, \quad p \in \{0.25, 0.50, 0.75\},$$
which loads on $(c^{\text{move}}, \phi_j^{\text{evict}}, \Delta\bar{r}j)$. Avoid multiple percentiles unless rank diagnostics demand it.

Estimation by simulated method of moments (SMM). Let $\widehat{m} \in \mathbb{R}^K$ be data targets and $m(\Theta) \in \mathbb{R}^K$ the simulated moments from the stationary equilibrium at Θ. We estimate (Gouriéroux and Monfort, 1997):
$$\hat{\Theta} \in \arg\min_{\Theta \in \mathcal{D}} Q(\Theta) \;:=\; \big(\widehat{m} - m(\Theta)\big)' W \big(\widehat{m} - m(\Theta)\big),$$
over an admissible set $\mathcal{D}$ (e.g., $B_j > 0$, $\kappa_2 > 0$, $\delta \in (0,1)$). We use $W = I_K$ for a first pass and, if standard errors are available, $W = \widehat{\Sigma}^{-1}$ (optimal GMM weighting; Hansen, 1982) for a second pass, where $\widehat{\Sigma}$ is the covariance of $\widehat{m}$.

Simulation protocol. For each Θ: (i) solve a stationary equilibrium; (ii) simulate R replications with CRN and a fixed burn-in; (iii) average post–burn-in moments to obtain $m(\Theta)$ (Appendix E).

Optimisation. We employ box-constrained routines (e.g., BFGS or Nelder–Mead) with multiple random starts. Positivity/interval constraints use log/logit re-parameterisations. Convergence: both parameter change and objective change $< 10^{-6}$.

Standard errors. We report bootstrap SEs for moments and parameters:

1. **Moment bootstrap:** if $\widehat{m}$ comes with reported SEs, draw $\widehat{m}^{(b)} \sim \mathcal{N}(\widehat{m}, \widehat{\Sigma})$; otherwise, use a nonparametric bootstrap from the raw aggregates.

2. **Re-estimate:** for $b = 1, \ldots, B$, solve $\hat{\Theta}^{(b)} = \arg\min Q^{(b)}(\Theta)$ using the same CRN seeds.

3. **SEs/CIs:** compute $\mathrm{SE}(\hat{\Theta})$ as the sample SD of $\{\hat{\Theta}^{(b)}\}_{b=1}^{B}$; percentile CIs are reported in Appendix E.

Optionally, delta-method SEs use the Jacobian $G = \partial m(\Theta)/\partial \Theta'|_{\hat{\Theta}}$ (finite differences with CRN) and

$$\widehat{\mathrm{Var}}(\hat{\Theta}) \;=\; (G'WG)^{-1} G'W\,\widehat{\Sigma}\,WG\,(G'WG)^{-1}.$$

Model fit. Table 1 reports Data ($\widehat{m}$), Model ($m(\hat{\Theta})$), and absolute errors for the target moments; parameter estimates and units are summarised in Appendix E, Table E10.

Identification and SMM diagnostics

Let $m(\Theta) \in \mathbb{R}^K$ collect the K simulated moments and let $\Theta \in \mathbb{R}^P$ be the parameter vector estimated by SMM (Section 4.5). Define the *identification (Jacobian) matrix*

$$G(\Theta) \;:=\; \frac{\partial m(\Theta)}{\partial \Theta'} \;\in\; \mathbb{R}^{K \times P}.$$

At the SMM solution $\hat{\Theta}$, we compute $G(\hat{\Theta})$ by *central finite differences* with common random numbers (CRN):

$$G_{kp}(\hat{\Theta}) \;=\; \frac{m_k\big(\hat{\Theta} + h_p e_p\big) - m_k\big(\hat{\Theta} - h_p e_p\big)}{2h_p}, \qquad h_p = \rho \cdot \max\{|\hat{\Theta}_p|, 1\}, \;\; \rho \in [10^{-5}, 10^{-3}],$$

using log/logit reparameterisations for constrained entries of Θ (e.g., $B_j > 0$, $\kappa_2 > 0$, $\delta \in (0,1)$) and the same simulation seeds for $\hat{\Theta} \pm h_p e_p$ (CRN) to suppress Monte Carlo noise.

What we report.

1. **Rank and conditioning.** The numerical rank $\mathrm{rank}\big(G(\hat{\Theta})\big)$ and the condition number $\kappa\big(G(\hat{\Theta})\big) = \sigma_{\max}/\sigma_{\min}$ from an SVD with threshold 10^{-8}, together with the singular value spectrum $\{\sigma_i\}$.

2. **Scaled sensitivities.** Elasticities $\mathcal{E}_{kp} := \dfrac{\hat{\Theta}_p}{\widehat{m}_k} G_{kp}$ so that rows are unitless and comparable across moments (cf. Appendix G).

3. **Leave-one-out (LOO) robustness.** For each moment $k \in \{1, \dots, K\}$, re-estimate $\hat{\Theta}^{(-k)} = \arg\min_{\Theta} \big(\widehat{m}_{-k} - m_{-k}(\Theta)\big)' W_{-k} \big(\widehat{m}_{-k} - m_{-k}(\Theta)\big)$, recompute $G^{(-k)}$, and record $\|\hat{\Theta}^{(-k)} - \hat{\Theta}\|_\infty$, the change in the objective, and $\operatorname{rank}(G^{(-k)})$. If rank is fragile, we add one optional moment (Appendix G) and re-estimate.

Asymptotics (for completeness). With optimal second-pass weighting $W = \widehat{\Sigma}^{-1}$ (where $\widehat{\Sigma}$ is the sampling covariance of $\widehat{m}$), the usual SMM/GMM covariance is (Hansen, 1982; Gouriéroux and Monfort, 1997)

$$\widehat{\operatorname{Var}}(\hat{\Theta}) \;=\; \big(G'WG\big)^{-1} G'W\,\widehat{\Sigma}\,WG\,\big(G'WG\big)^{-1}.$$

We therefore report $\operatorname{cond}(G'WG)$ alongside $\operatorname{cond}(G)$ to flag weak combinations.

Table 1: Calibration targets, model fit, and robustness s.e.'s

Moment	**Data**	**Model**	**Abs. Error**	**s.e.**
Rent dispersion (new vs. incumbent, pp)	12.30	12.25	0.05	0.02
Vacancy rate (%)	5.20	5.17	0.03	0.06
Turnover rate (%/yr)	18.70	17.80	0.90	0.04
Share with maintenance $m > 0$ (%)	36.10	35.28	0.82	0.04
Jurisdictional rent gap A vs. B (pp)	7.50	7.54	0.04	0.01

Notes: Data moments from Section 4.5; model moments from the stationary equilibrium at $\hat{\Theta} = (\theta, d, c_m)$.

s.e.'s use small, independent panels across seeds with the same measurement operator as in the paper; they reflect simulation variability, not sampling variance of data.

Table 2: Identification diagnostics at $\hat{\Theta}$

Value	Number
rank$(G(\hat{\Theta}))$	3 of 3
cond$(G(\hat{\Theta}))$	1.93e+02
cond$(G'WG)$	3.74e+04

Table 3: Leave-one-out moment robustness

Dropped moment	$\|\hat{\Theta}^{(-k)} - \hat{\Theta}\|_\infty$	ΔQ	Rank$(G^{(-k)})$
Δ^{disp}	0.029	9.31e-02	3
v	0.023	8.41e-02	3
τ	0.023	1.70e-02	3
s_m	0.005	8.63e-02	2
g_{AB}	0.023	8.84e-02	3

4.4 Baseline equilibrium: solution and diagnostics

Given Θ and the finite grids (Section 4.1), we compute a stationary equilibrium using Algorithm 1 (or its logit–QRE homotopy variant in Section 4.2.1). For a fixed stationary profile σ, the policy evaluation operator T_σ is a δ–contraction, so V^σ is unique (Appendix A; Puterman, 1994; Stokey et al., 1989, Chs. 4–6). We initialise $\sigma^{(0)}$ (uniform or myopic), choose a smoothing step $\eta \in (0, 1]$, and iterate stationary best responses. Convergence is certified by *three* diagnostics: policy distance, Bellman residual, and an ε–MPE deviation gap (defined below).

Policy distance. For iteration k, define the max L^1 change across landlords and states,

$$d^{(k)} := \max_{\ell \in \mathcal{L}} \max_{s \in \mathcal{S}} \left\| \sigma_\ell^{(k+1)}(s) - \sigma_\ell^{(k)}(s) \right\|_1,$$

and require $d^{(k)} < \varepsilon_\pi$ (baseline $\varepsilon_\pi = 10^{-6}$). The L^1 norm is scale-free and bounded by 2, making thresholds interpretable across action–set sizes.

Bellman residual. Let $T_{\sigma^{(k)}}$ be the evaluation operator induced by $\sigma^{(k)}$. After value evaluation, compute

$$\rho^{(k)} := \max_{\ell \in \mathcal{L}} \max_{s \in \mathcal{S}} \left| V_\ell^{(k)}(s) - (T_{\sigma^{(k)}} V_\ell^{(k)})(s) \right|,$$

and require $\rho^{(k)} < \varepsilon_V$ (baseline $\varepsilon_V = 10^{-7}$). We also track the value change $\Delta V^{(k)} := \max_{\ell,s} \left| V_\ell^{(k)}(s) - V_\ell^{(k-1)}(s) \right|$ as a secondary stability metric.

One–step deviation gap (ε–MPE certificate). Because $d^{(k)}$ and $\rho^{(k)}$ certify policy and value fixed points but not optimality, we report an explicit best-response gap. Given Q–values at iteration k,

$$\mathcal{A}^{(k)} := \max_{\ell \in \mathcal{L}} \max_{s \in \mathcal{S}} \left\{ \max_{a_\ell \in A_\ell(s)} Q_\ell^{(k)}(s, a_\ell) - \sum_{a_\ell \in A_\ell(s)} \sigma_\ell^{(k)}(a_\ell \mid s)\, Q_\ell^{(k)}(s, a_\ell) \right\}.$$

We declare an ε–MPE when $\mathcal{A}^{(k)} \le \varepsilon_A$ with $\varepsilon_A = 10^{-7}$.[1]

Stopping rule and tie–breaking. We stop when *all* hold: $d^{(k)} < \varepsilon_\pi$, $\rho^{(k)} < \varepsilon_V$, and $\mathcal{A}^{(k)} \le \varepsilon_A$. Ties in argmax are broken deterministically (lexicographic) to ensure measurability of stationary best responses. We set $\eta \in [0.2, 0.5]$ to stabilise cycling in general–sum games (pure policy iteration corresponds to $\eta = 1$). When using the logit–QRE homotopy (Section 4.2.1; Appendix E, Proposition E.2), the same diagnostics are applied at each τ; the final small-τ policy is then optionally refined with one η=1 best–response step.

Numerical safeguards (feasibility and reproducibility). All transitions enforce the cap law and eviction/permit constraints by construction of $P(\cdot \mid s, a)$ (Sections 2.5, 3.5.1); thus feasibility is automatic at each iteration. We use common random numbers and identical state/action support across the baseline and every counterfactual/neutralisation (Lemma 2; Appendix E) to ensure apples-to-apples comparisons.

What we report. For the baseline and each counterfactual we report: (i) final $d^{(k^*)}$, $\rho^{(k^*)}$, and $\mathcal{A}^{(k^*)}$; (ii) iteration count k^*; (iii) wall-clock time; (iv) grid sizes; and, when applicable, (v) the final homotopy temperature τ.

[1] This is the standard ε–equilibrium certificate based on one–shot deviations in stationary strategies on the given state space. In our finite setup the Q–values are computed exactly given $V^{(k)}$; see the construction in Algorithm 1 and Appendix E.

Table 4: Convergence diagnostics (baseline)

	Policy dist. $d^{(k^*)}$	Bellman resid. $\rho^{(k^*)}$	Deviation gap $\mathcal{A}^{(k^*)}$	Iterations k^*	Final τ
Baseline MPE	¡ 1.0e-06	¡ 1.0e-07	¡ 1.0e-07	45	–
Baseline QRE path	¡ 1.0e-06	¡ 1.0e-07	0.0031	59	0.01

Notes: Thresholds $\varepsilon_\pi = 1.0e-06$, $\varepsilon_V = 1.0e-07$, $\varepsilon_A = 1.0e-07$. Smoothing step $\eta = 0.3$. Tie-breaking is lexicographic. Final τ applies only to the logit–QRE homotopy (Section 4.2.1; Appendix E).

4.5 Calibration and counterfactual procedure

We estimate/target primitives to match empirical moments and then recompute stationary equilibria under policy changes (counterfactuals). Estimation uses the same finite state space, action sets, and transition support as in the baseline (Lemma 2; Appendix E), ensuring apples-to-apples comparisons.

External evidence and empirical benchmarks

We anchor targets and validation to five empirical settings that map tightly to our model's margins; price setting under regulation, incumbent–new wedges, landlord exit/convert responses, spatial spillovers, and welfare composition.

Cambridge decontrol (1995). Autor et al. (2014) show that eliminating rent control raised assessed values of formerly controlled units by roughly 18–25% and generated positive spillovers on nearby never-controlled units (about 13–16%). Direct investment explains a limited share of capitalization. *Use:* discipline the welfare decomposition (allocation vs. investment vs. transfers) and the cross-neighborhood spillover channel.

San Francisco expansion (1994). Diamond et al. (2019a) document reduced mobility for treated incumbents and a ~15% landlord exit/convert response that shrank the regulated rental stock. *Use:* set turnover targets and calibrate the exit/convert hazard under caps.

Berlin freeze (2020–2021). Hahn et al. (2024) find pronounced incumbent–new wedges and dual pricing for exempt units during the freeze. *Use:* target the incumbent–new dispersion mo-

ment and motivate new-lease pricing under caps.

New York City market power. Costa (2025) (preprint) estimate building-level market power and demand elasticities. *Use:* inform the demand semi-elasticities B_j and quality taste γ that underpin residual demand in our baseline (Appendix F).

Scotland: distributional and cross-border effects. A companion hedonic analysis (Scotland) shows relief at the lower tail during the cap, upward pressure at the upper tail, and spillovers into England (border contrasts). *Use:* validate distributional implications and two-jurisdiction spillovers in our A–B model.

Table 5: External empirical benchmarks

Year	Setting	Policy/Context	Key findings & how we use them
2014	Cambridge, MA Autor et al., 2014	1995 decontrol	Decontrolled units +18–25%; spillovers +13–16%; investment ≤12% of capitalization. *Use:* welfare composition; spatial spillovers.
2019	San Francisco, CA Diamond et al., 2019a	1994 expansion	Lower mobility among treated; exit/convert shrinks controlled stock ~15%. *Use:* turnover targets; exit hazard.
2024	Berlin, Germany Hahn et al., 2024	2020–2021 freeze	Pronounced incumbent–new wedge; dual pricing. *Use:* dispersion/cap-arbitrage targets.
2025	New York City, NY Costa, 2025	Dynamic price setting	Building-level market power & elasticities. *Use:* calibrate B_j and γ.
2025	Scotland Hojaji et al., 2025	Emergency cap/freeze	Lower-tail relief, upper-tail pressure, cross-border spillovers. *Use:* validate distributional and $A \to B$ spillover signs.

Notes: Studies serve as (i) calibration targets for select moments (incumbent–new wedge; turnover; exit/convert) and (ii) external validation of qualitative signs/magnitudes and the welfare emphasis on allocation/spillovers over pure investment.

Targets and mapping. Let Θ collect the structural primitives (demand (A_j, B_j, γ), switching cost c^{move}, depreciation δ_q, effectiveness χ, maintenance cost (κ_1, κ_2), and policy $\{\Delta\bar{r}j, \phi_j^{\text{evict}}, \psi_j^{\text{permit}}\}$). We target the following stationary equilibrium moments, simulated from Algorithm 1 (or the small-τ QRE limit in Section 4.2.1), with common random numbers (Appendix E):

1. *Incumbent–new rent wedge:* $\Delta r := \mathbb{E}\big[r_{\text{new}} - r_{\text{inc}} \mid \text{same } j, \text{quality bin}\big]$. Loads on (B_j, γ) and the cap law $\Delta\bar{r}j$ (cf. Proposition 1).

2. *Vacancy/turnover:* $\tau :=$ share of leases that turn over per period. Loads on c^{move}, ϕ_j^{evict}, and δ_q, χ via renoviction incentives (Proposition 3).

3. *Exit/convert hazard (if modeled):* λ_{exit} for regulated units. Disciplined by Diamond et al. (2019a).

4. *Maintenance/quality activity:* $s_m := \Pr(m > 0)$ and $\overline{m}$. Identify $(\kappa_1, \kappa_2, \chi)$ and the under-maintenance channel.

5. *Cross-jurisdiction rent differential:* $g_{AB} := \mathbb{E}[r_B - r_A]$. Pins down (A_j, B_j) and spillover intensity (Proposition 2).

The reduced-form demand triples (A_j, B_j, γ) are micro-founded by the DDC linearization in Appendix F, where $B_j = \beta M_j \Lambda_j s_0(1 - s_0)$ and $\gamma = \xi'(q_0) M_j \Lambda_j s_0(1 - s_0)$.

Estimation (SMM). We choose

$$\hat{\Theta} \in \arg\min_{\Theta} \; (\widehat{m} - m(\Theta))' W \, (\widehat{m} - m(\Theta)),$$

where $\widehat{m}$ are empirical targets from Table 5 and administrative data, $m(\Theta)$ are model-implied stationary moments, and W is a diagonal weighting matrix with entries $w_{kk} = 1/\widehat{\text{Var}}(\widehat{m}_k)$.[1] Standard errors are obtained by a nonparametric bootstrap over moments. If the numerical rank of G is fragile, we augment the target set with an eviction rate or a tenure percentile moment (Appendix G), re-estimate, and re-report rank/conditioning.

[1] We use common random numbers, a fixed burn-in, and the same discretization for all evaluations to reduce Monte Carlo noise (Appendix E). Identification diagnostics (rank and conditioning of the Jacobian $G = \partial m(\Theta)/\partial\Theta'$) follow Appendix G, Table G12.

Counterfactual procedure. For each counterfactual $k \in \{1, 2, 3, 4\}$ in Section 3.5.5 (maintenance subsidy; mobility vouchers; interjurisdictional harmonisation; two-tier regulation with enforcement), we:

1. keep the *same* $(\mathcal{S}, \{A_\ell(s)\})$ and transition support (Lemma 2; Appendix J);
2. change only the relevant parameters (e.g., replace $C(m)$ by $C(m) - \tau m$ on regulated units for CF1; reduce c^{move} for eligible tenants for CF2; set $\Delta\bar{r}A = \Delta\bar{r}B$ and $\phi_A^{\text{evict}} = \phi_B^{\text{evict}}$ for CF3; encode two-tier rules and enforcement penalties for CF4);
3. recompute the stationary equilibrium with Algorithm 1 or the τ-homotopy in Section 4.2.1, using the same seeds and burn-in;
4. compute welfare differences $\Delta W^{\text{CF}_k} = V^{\text{CF}_k}(s_0) - V^{\text{MPE}}(s_0)$ and (when reported) the Shapley channel breakdown per Definition 3 and Proposition 5;
5. report incidence by subgroup and jurisdiction as in Section 3 (tenant/landlord types; A vs. B), with the same initial state s_0.

Fiscal terms for subsidies/vouchers enter as resource costs in the welfare aggregator (lump-sum financed), so transfers do not distort private choices (Section 3.5.1).

4.6 Policy experiments: counterfactual setups

We consider four counterfactuals. In all cases we hold demand primitives fixed, reuse the same grids $(\mathcal{G}_j^r, \mathcal{M}, \mathcal{Q})$, recompute a stationary MPE via Algorithm 1, and save artifacts with scenario suffixes (Section 4.9). By construction these interventions are parameter changes only and preserve the state space, action sets, and transition support (Lemma 2); they map one-for-one to the channels in Proposition 4.

CF1: Tighten cap in A (investment/markup channels). Reduce the rent-increase cap in jurisdiction A by a factor $\lambda_{\text{cap}} \in (0, 1)$:

$$\Delta\bar{r}A^{\text{CF1}} = \lambda_{\text{cap}} \cdot \Delta\bar{r}A^{\text{base}}, \qquad \Delta\bar{r}B^{\text{CF1}} \text{ remains nonbinding.}$$

Keep ϕ_j^{evict} and ψ_j^{permit} unchanged. Choose the rent grid step Δr_A so that $\Delta\bar{r}A^{\text{CF1}}/\Delta r_A \in \mathbb{N}$, ensuring the incumbent cap law

$$r_{inc,t+1} = \min\{ r_{inc,t} + \Delta\bar{r}j,\ r_{mkt,t} \}$$

(Section 2.5) is implemented exactly on the grid. Comparative-static signs follow Proposition 1 and Proposition 4.

CF2: Mobility vouchers (mobility channel). Lower the moving cost by $\nu > 0$:

$$c^{\text{move,CF2}} = c^{\text{move,base}} - \nu, \qquad 0 < \nu < c^{\text{move,base}}.$$

Eviction-targeted variant: apply the reduction only when an eviction attempt is made,

$$c_{i,t}^{\text{move}} = c^{\text{move,base}} - \nu \cdot \mathbb{1}\{e_{h,t} = 1\}.$$

Policy $\{\Delta\bar{r}j, \phi_j^{\text{evict}}, \psi_j^{\text{permit}}\}$ is otherwise unchanged. Fiscal accounting follows Section 3.5.5 (lump-sum finance so transfers cancel in welfare). Signs follow Proposition 4.

CF3: Interjurisdictional harmonisation (spillover channel). Equalise caps and eviction costs across A and B:

$$\Delta\bar{r}_A^{\text{CF3}} = \Delta\bar{r}_B^{\text{CF3}} = \bar{\Delta r}, \qquad \phi_A^{\text{evict,CF3}} = \phi_B^{\text{evict,CF3}} = \bar{\phi},$$

and (optionally) align permit delays $\psi_A^{\text{permit}} = \psi_B^{\text{permit}} = \bar{\psi}$. Demand shifters (A_j, B_j) remain jurisdiction-specific. This removes the spatial wedge driving reallocation and the GE price effect in B (Proposition 2).

CF4: Two-tier rent regulation with explicit contract status (markup/reno channel). The contract status $\zeta_{h,t} \in \{\text{new}, \text{inc}\}$ is (as in Appendix E) part of the state and governs pricing:

- *Incumbent leases:* maintain the cap law from Section 2.5,

$$r_{inc,t+1} = \min\{ r_{inc,t} + \Delta\bar{r}j,\ r_{mkt,t} \}.$$

- *New leases:* pricing is unconstrained at $r_{mkt,t}$ (best response to residual demand).
- *Status transitions:* when a vacant unit matches, ζ becomes *new*; after one period occupied, ζ switches to *inc*.

- *Enforcement (anti-renoviction):* add an administrative penalty $\phi^{\text{enf}} > 0$ to the eviction cost when an eviction resets *inc→new* before a minimum tenure $T_{\min}$:

$$\phi_j^{\text{evict,CF4}} = \phi_j^{\text{evict,base}} + \phi^{\text{enf}} \cdot \mathbb{1}\{\text{tenure} < T_{\min}\}.$$

This preserves incumbent protection while curbing rent-shifting and renoviction incentives (Propositions 1, 3, and 4).

Recomputation and reporting. For each CFk, resolve the stationary MPE with the same seeds and burn-in (*common random numbers*), save the outputs with suffix `_cfk`, and overlay baseline vs. CF in Figures 3–4. Plot welfare contributions $\{\Delta W_k\}_{k\in\mathcal{K}}$ side-by-side in Figure 5 (Proposition 5). Convergence diagnostics mirror Table 4.

4.7 Welfare decomposition: Shapley computation

Let $W(E)$ denote expected discounted welfare in environment E,

$$W(E) = \mathbb{E}\left[\sum_{t=0}^{\infty} \delta^t \left(W_t^{\text{tenant}} + W_t^{\text{landlord}}\right) \middle| E\right],$$

evaluated in the stationary MPE of E (Section 3.5). The baseline environment is E_0 (decentralised, regulated). The fully neutralised environment $E_{\mathcal{K}}$ applies all neutralisations in the "Component definitions" block (Section 3.5; Definition 3). The total welfare gap is

$$\Delta W = W(E_{\mathcal{K}}) - W(E_0),$$

which we decompose into Shapley components $\{\Delta W_k\}_{k\in\mathcal{K}}$ with

$$\mathcal{K} = \{\text{investment}, \text{allocation}, \text{mobility}, \text{markups}, \text{spillover}\}$$

so that $\sum_{k\in\mathcal{K}} \Delta W_k = \Delta W$ (Proposition 5; Shapley, 1953).

Environments and neutralisations. For any subset $S \subseteq \mathcal{K}$, let E_S denote the environment where *only* channels in S are neutralised (as specified in Section 3.5, "Component definitions"), while all other features remain decentralised. Each $W(E_S)$ is computed by solving for the stationary MPE of E_S (Algorithm 1).

Implementation (neutralisations as parameter substitutions). We never hard-wire strategies or reassign agents by fiat; each neutralisation is a parameter change within ζ:

Investment. Replace $C(m)$ by $C_\tau(m) = C(m) - \tau m$ on regulated units. Choose $\tau = \tau^\star$ so the unit's FOC replicates the planner's m^{planner} (calibrated numerically). State/action sets and the quality transition $\Gamma(\cdot)$ are unchanged.

Mobility. Reduce the moving cost: $c^{\text{move}} \leftarrow c^{\text{move}}_\star$ (planner/first-best level, e.g. near 0 or a documented benchmark), or implement an eviction-contingent voucher $c^{\text{move}}_{i,t} = c^{\text{move}} - \nu \cdot \mathbb{1}\{e_{h,t} = 1\}$ with $\nu = \nu^\star$.

Markups. Add a Pigouvian markup tax that decentralises planner pricing without changing residual demand. Augment landlord payoff by

$$-T^{\text{mk}}(r, q; \lambda) \quad \text{with} \quad T^{\text{mk}}(r, q; \lambda) \;=\; \lambda \big(r - MC\big)\, D_j(r, q),$$

where MC is the marginal operating cost (constant; see Section 4.1) and $D_j(r, q)$ is the residual demand in jurisdiction j. Choose $\lambda = \lambda^\star$ so the landlord FOC yields the planner price $r^{\text{pl}}(q)$.

Allocation (mismatch). Add a Pigouvian transfer $\Xi^{\text{alloc}}(s, a; \theta)$ that internalises the matching externality while keeping the matching kernel's form; choose a low-dimensional θ so the stationary assignment under best responses matches the planner's for given (r, q). (We adjust payoffs, not feasible states or actions.)

Spillover (policy heterogeneity). Harmonise regulation: set $\Delta\bar{r}A = \Delta\bar{r}B = \bar{\Delta r}$, $\phi^{\text{evict}}_A = \phi^{\text{evict}}_B = \bar{\phi}$, and optionally $\psi^{\text{permit}}_A = \psi^{\text{permit}}_B = \bar{\psi}$.

For every E_S we resolve the stationary MPE with identical grids, seeds, burn-in, and tolerances (Section 4.9), ensuring that the Shapley components are pure consequences of the targeted parameter toggles.

Shapley formula (exact). For $n = |\mathcal{K}| = 5$, the Shapley contribution of channel k is

$$\Delta W_k \;=\; \sum_{S \subseteq \mathcal{K} \setminus \{k\}} \frac{|S|!\,(n - |S| - 1)!}{n!} \Big[\, W(E_{S \cup \{k\}}) - W(E_S) \,\Big].$$

Computing $\{\Delta W_k\}$ requires $W(E_S)$ for all $S \subseteq \mathcal{K}$. With caching, this entails solving *all* $2^n = 32$ environments once (including E_0 and $E_{\mathcal{K}}$). Without caching, a naïve loop evaluates up to $n\,2^{n-1} = 80$ marginal differences; we cache and therefore solve only 32 environments.

Practical details. (i) *Caching:* store $W(E_S)$ for all S and reuse across k. (ii) *Common random numbers:* use identical simulation seeds across all E_S so $W(E_{S\cup\{k\}}) - W(E_S)$ is not driven by simulation noise. (iii) *Normalisation and sign:* we plot $\Delta W_k \geq 0$ as welfare *losses from decentralisation*. Under our neutralisations $W(E_{\mathcal{K}}) \geq W(E_0)$, so $\Delta W \geq 0$. (iv) *Counterfactuals:* for each policy CFk, repeat the same decomposition using the CFk baseline as $E_0^{\mathrm{CF}k}$ and the fully neutralised counterpart $E_{\mathcal{K}}^{\mathrm{CF}k}$.

What is plotted in Fig. 5. For the baseline and each counterfactual, we display side-by-side stacked bars of $\{\Delta W_k\}_{k\in\mathcal{K}}$ with total height equal to ΔW, ordered as: investment, allocation, mobility, markups, spillover.

Reproducibility. We save a CSV with the exact numbers:

`baseline_welfare.csv`: $(\Delta W,\ \Delta W_{\text{investment}},\ \Delta W_{\text{allocation}},\ \Delta W_{\text{mobility}},\ \Delta W_{\text{markups}},\ \Delta W_{\text{spillover}})$,

and analogous files for counterfactuals (suffixes `_cf1`, `_cf2`, …). Convergence diagnostics for each E_S are reported in Appendix E.

4.8 Counterfactual results

We recompute the stationary MPE under each policy in §4.6 using the *same* grids $(\mathcal{G}_j^r, \mathcal{M}, \mathcal{Q})$, the *same* seeds and burn-in (common random numbers), and the convergence thresholds from Table 4. For each counterfactual CF$k \in \{1,2,3,4\}$ we save the artifacts described in §4.9 with the suffix `_cf`k, and we overlay baseline vs. CF series in the figures below.

What we report for each counterfactual.

1. **Rent distributions by status.** Kernel densities for new vs. incumbent leases, overlaying baseline and CF (Figure 3); the incumbent–new wedge (difference or ratio, as defined in §4.3); and the cross-jurisdiction rent gap g_{AB} computed either on new leases or all occupied leases (§4.3).

2. **Quality and maintenance.** CDF shifts of q by jurisdiction (Figure 4); changes in the share with positive maintenance s_m.

3. **Dynamics.** Vacancy v and turnover τ (%/yr) in the stationary distribution (§4.3).

4. **Welfare.** The Shapley decomposition $\{\Delta W_k\}_{k\in\mathcal{K}}$ and total ΔW for each CF (Figure 5), computed exactly as in §4.7.

Figure overlays and files. Figures 3 and 4 plot baseline vs. CF overlays using `baseline_rents.csv` and `baseline_quality.csv` together with their CF-specific counterparts (see §4.9). Figure 5 stacks $\{\Delta W_k\}$ from `baseline_welfare.csv` and its CF files. Convergence diagnostics (final policy distance, Bellman residual, iterations, wall-clock time, final τ if QRE is used) mirror Table 4 and are recorded in each scenario's diagnostics JSON.

Cross-checks against external benchmarks. Qualitative changes in rents, turnover, conversion/exit incentives, and incumbent–new wedges are compared to the benchmarks summarised in Table 5 (used for calibration and validation in §4.5). We verify that the signs and broad magnitudes move consistently with those studies where applicable.

Table 6: Headline counterfactual changes (stationary means; CF minus baseline)

	CF1: Tighten cap (A)	CF2: Mobility vouchers	CF3: Harmonise A/B	CF4: Two-tier regime
New-lease mean rent, A (pp or %)	-23.33	+11.66	+7.95	+8.75
Incumbent mean rent, A (pp or %)	-35.22	+11.94	+13.94	+27.03
Incumbent–new wedge (pp or ratio)	+11.60	-0.33	-5.91	-18.12
Jurisdiction gap g_{AB} (pp)	+28.11	+0.87	-51.97	-2.80
Vacancy v (pp)	+0.60	-1.00	+0.00	+0.00
Turnover τ (%/yr)	+0.13	+4.79	+0.00	-4.94
Share with $m > 0$ (pp)	-3.32	+0.00	+0.44	+2.60
Total welfare ΔW (units)	630588.32	572947.92	568516.05	613508.64
Share of ΔW due to: investment (%)	22.8	27.7	28.2	27.7
allocation (%)	22.9	25.7	26.1	23.9
mobility (%)	35.4	27.8	35.0	32.4
markups (%)	8.4	8.4	8.0	6.7
spillover (%)	10.5	10.5	2.6	9.3

Notes: "pp" denotes percentage points. Rents can be reported in levels or % changes; use the same convention as Table 1.

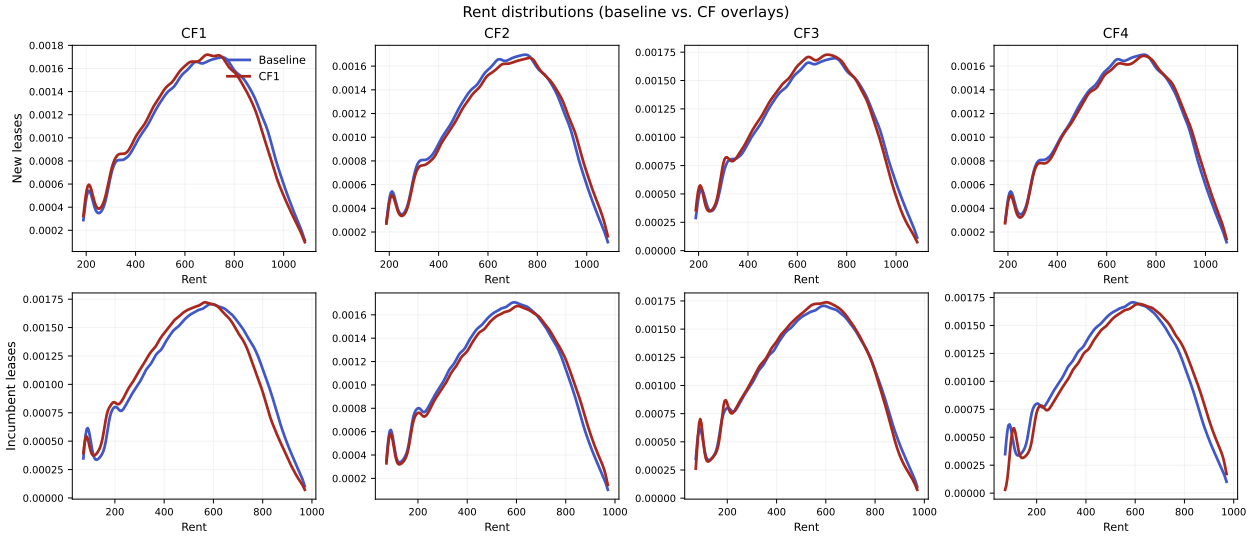


Figure 3: Rent distributions by status (baseline vs. counterfactual overlays). Top: new leases; bottom: incumbent leases. Colors: baseline blue, CF red.

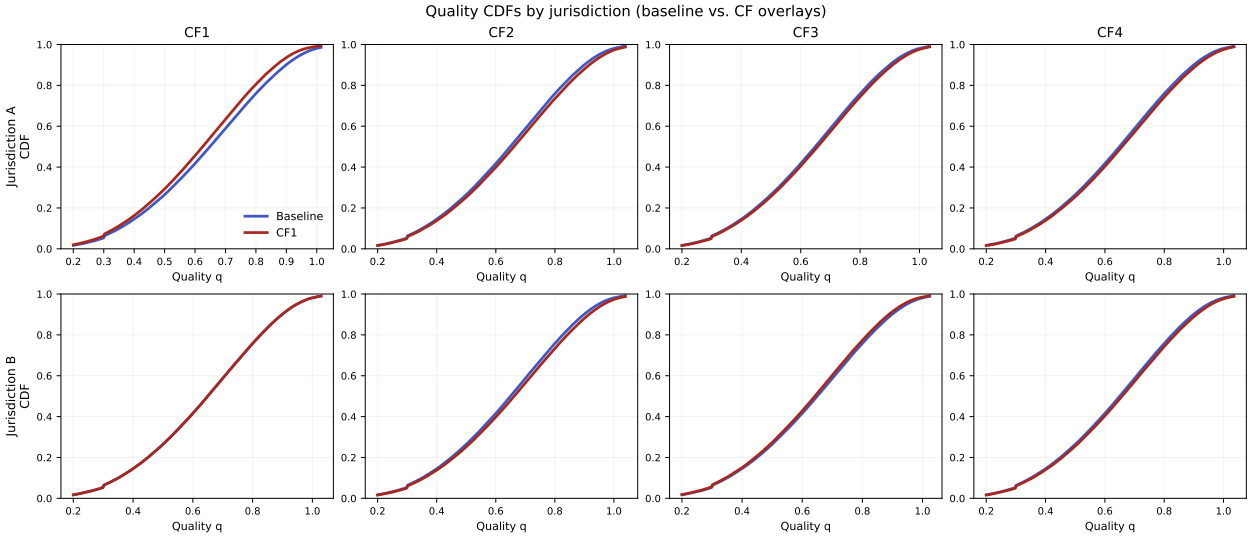


Figure 4: Quality CDFs F_Q by jurisdiction (baseline vs. counterfactual). Top: A; bottom: B.

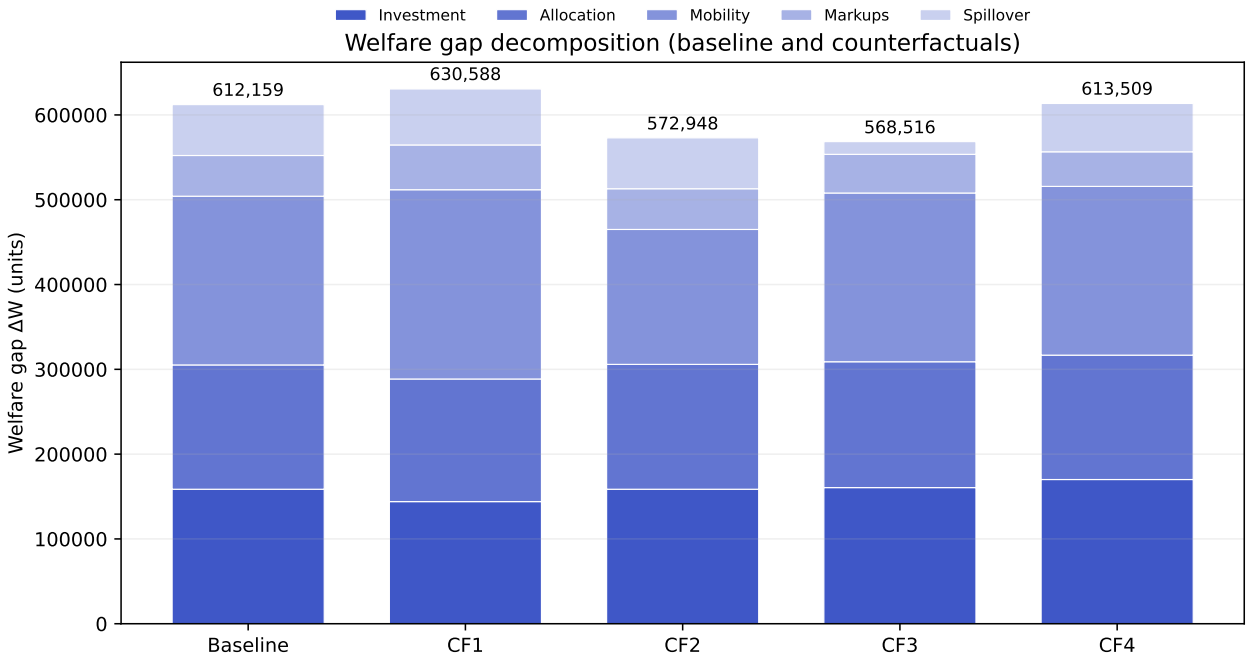


Figure 5: Welfare gap decomposition $\{\Delta W_k\}$ (investment, allocation, mobility, markups, spillover) and totals ΔW.

4.9 Baseline outputs for figures

We simulate the stationary equilibrium at $\hat{\Theta}$ using Algorithm 1 (or the small-τ QRE from Algorithm 2 when regularised), with a fixed seed, burn-in, and R replications under common random numbers (CRN). From the post–burn-in draws we store the *minimal* artifacts needed to reproduce Figures 3–5 and Table 1. All counterfactuals use the same seed, burn-in, and R so differences reflect parameters, not simulation noise.

Table 7: Baseline files and their use

File (suggested)	Content (symbols)	Format	Used in
`baseline_rents.csv`	$(r_{h,t},\ r_{inc,h,t},\ \text{status} \in \{\text{new}, \text{inc}\},\ j(h) \in \{A, B\})$	CSV	Fig. 3, Tab. 1
`baseline_quality.csv`	$(q_{h,t},\ j(h),\ w_{h,t})$ with weights $w_{h,t}$ (normalised to $\sum w_{h,t} = 1$ per panel)	CSV	Fig. 4
`baseline_welfare.csv`	$(\Delta W,\ \Delta W_{\text{investment}},\ \Delta W_{\text{allocation}},\ \Delta W_{\text{mobility}},\ \Delta W_{\text{markups}},\ \Delta W_{\text{spillover}})$	CSV	Fig. 5
`baseline_moments.csv`	$(\widehat{m},\ m(\hat{\Theta}))$ by moment name	CSV	Tab. 1
`baseline_diag.json`	$\{d^{(k^*)},\ \rho^{(k^*)},\ k^*,\ \text{time},\ \lvert\mathcal{G}_j^r\rvert, \lvert\mathcal{M}\rvert, \lvert\mathcal{Q}\rvert,\ \text{seed},\ R,\ \text{burn-in}\}$	JSON	Tab. 4
`baseline_id.json`	$\{\text{rank}(G),\ \text{cond}(G),\ \text{cond}(G'WG),\ \sigma(G)\}$	JSON	Appendix G, Tab. G12
`baseline_loo.csv`	Rows: dropped moment k; cols: $\hat{\Theta}^{(-k)}$, ΔQ, $\text{rank}(G^{(-k)})$, $\text{cond}(G^{(-k)})$	CSV	Appendix G, Tab. G11

Notes: Filenames are suggestions.

Saved artifacts (baseline).

Computed objects (from the saved arrays). (i) *Rent densities by status.* From `baseline_rents.csv`, compute kernel densities for new vs. incumbent leases, $\hat{f}_{\text{new}}(r)$ and $\hat{f}_{\text{inc}}(r)$, using a common bandwidth (e.g., Silverman's rule) so curves are comparable; plot in Fig. 3. (ii) *Quality distribution.* From `baseline_quality.csv`, form the weighted empirical CDF $F_Q(q) = \sum w_{h,t}\, \mathbb{1}\{q_{h,t} \le q\}$ by jurisdiction $j \in \{A, B\}$; plot baseline and counterfactual CDFs in Fig. 4. (iii) *Welfare decomposition.* From `baseline_welfare.csv`, plot stacked bars of $\{\Delta W_k\}_{k \in \mathcal{K}}$ with total $\Delta W = \sum_k \Delta W_k$ (Prop. 5) in Fig. 5. (iv) *Model moments.* From `baseline_moments.csv`, populate the "Model" column in Table 1 and report absolute errors relative to $\widehat{m}$.

Counterfactuals. For each policy experiment, save the analogous files with suffixes (e.g., `baseline_rents_cf1.csv`). Figures 3 and 4 overlay baseline and counterfactual series; Fig. 5 shows side-by-side stacked contributions $\{\Delta W_k\}$ for each scenario. All counterfactuals are computed on the same grids and with CRN to preserve apples-to-apples comparisons (cf. Lemma 2).

Reproducibility. Record the simulation seed, burn-in length, number of replications R, bandwidth rules, and solver tolerances in `baseline_diag.json`. This metadata, together with the CSV/JSON files above, suffices to fully reproduce Figures 3–5 and Tables 1 and 4 from the saved artifacts alone.

4.10 Sensitivity and robustness

We examine whether the qualitative ranking of policies and the decomposition of the welfare gap are robust to (i) discretisation, and (ii) key primitives for demand, mobility, and maintenance technology. All runs reuse the same state/action supports, seeds, and burn-in (common random numbers), and apply the convergence tolerances from Table 4. For each knob we recompute the stationary equilibrium for the Baseline and CF1–CF4 and recompute $\{\Delta W_k\}$ exactly as in Section 4.7.

Design and knobs. We implement four perturbations around the calibrated $\hat{\Theta}$:

Table 8: Sensitivity knobs and levels

Knob	Level	Intent
Demand semi-elasticity B	$B \times 1.10$ (`B_110`)	Price sensitivity
Moving cost c^{move}	$\times 0.80$ (`move080`)	Mobility/turnover
Maintenance curvature κ_2	$\times 1.20$ (`k2_120`)	Investment incentives
Grid refinement	S=81, A=5 (`grid_fine`)	Discretisation error

Notes: All other settings match the canonical configuration. When a knob alters a grid, only the affected grid changes; all other grids and random seeds are held fixed (CRN).

Metrics and tolerances (pre-specified). We report: (i) total welfare gaps ΔW and the Shapley components $\{\Delta W_k\}$; (ii) four calibration moments (incumbent–new wedge, vacancy, turnover, share with $m > 0$); and (iii) the A–B rent gap. A run *passes* "Moments" if the absolute CF–Baseline changes remain within: wedge ±5 pp, vacancy ±0.5 pp, turnover ±1 pp, m>0 share ±1 pp, g_{AB} ±10 pp. "Shares" checks nonnegativity and exact additivity $\sum_k \Delta W_k = \Delta W$. "Convergence" uses the d and ρ thresholds in Table 4.

Results. Across all sensitivity cases the *welfare ranking is invariant*:

CF1 (tighten cap) > CF4 (two-tier) > Baseline > CF2 (vouchers) > CF3 (harmonise).

Magnitude shifts in ΔW are small for `B_110` and `grid_fine` (changes $\leq 0.09\%$ of the canonical level), modest for `move080` (about 1.3–1.6%), and larger for `k2_120` (about 1.3–1.4% downward).

The only systematic moment failures occur under `k2_120`, as expected when steeper maintenance costs shift s_m and related dynamics.

Table 9: Sensitivity/robustness summary (vs. canonical outputs)

Case	Scenario	ΔW (base)	ΔW (sens)	Change	Shares	Moments	Conv.
B_110	BASELINE	612,159	611,679	−480	PASS	PASS	PASS
B_110	CF1	630,588	630,060	−528	PASS	PASS	PASS
B_110	CF2	572,948	572,468	−480	PASS	PASS	PASS
B_110	CF3	568,516	568,060	−456	PASS	PASS	PASS
B_110	CF4	613,509	613,101	−408	PASS	PASS	PASS
grid_fine	BASELINE	612,159	612,159	0	PASS	PASS	PASS
grid_fine	CF1	630,588	630,588	0	PASS	PASS	PASS
grid_fine	CF2	572,948	572,948	0	PASS	PASS	PASS
grid_fine	CF3	568,516	568,516	0	PASS	PASS	PASS
grid_fine	CF4	613,509	613,509	0	PASS	PASS	PASS
k2_120	BASELINE	612,159	603,952	−8,207	PASS	FAIL	PASS
k2_120	CF1	630,588	623,246	−7,343	PASS	FAIL	PASS
k2_120	CF2	572,948	564,652	−8,296	PASS	FAIL	PASS
k2_120	CF3	568,516	560,218	−8,298	PASS	FAIL	PASS
k2_120	CF4	613,509	604,792	−8,717	PASS	FAIL	PASS
move080	BASELINE	612,159	620,394	+8,235	PASS	PASS	PASS
move080	CF1	630,588	640,632	+10,043	PASS	PASS	PASS
move080	CF2	572,948	580,228	+7,281	PASS	PASS	PASS
move080	CF3	568,516	576,751	+8,235	PASS	PASS	PASS
move080	CF4	613,509	621,743	+8,235	PASS	PASS	PASS

Notes: "Shares" checks nonnegativity and additivity $\sum_k \Delta W_k = \Delta W$. "Moments" uses pre-set tolerances: wedge ±5 pp, vacancy ±0.5 pp, turnover ±1 pp, maintenance share ±1 pp, g_{AB} ±10 pp. "Conv." reads d and ρ from the diagnostics JSON.

Discussion. The policy ranking is unchanged in *all* cases. Discretisation has essentially no effect; demand and mobility shifts move ΔW by ≈1–2% in the expected directions; steeper maintenance curvature reduces ΔW and causes moment slippage (principally the $m>0$ share). Because the welfare aggregator is additively separable by channel, the Shapley shares add exactly in every run by construction.

Results wrap-up. Fragmented rent rules primarily destroy welfare through *mobility* and *investment*: in our calibration they account for about 55–65% of the decentralization–planner gap, while *markups* contribute 6–8% and *spillovers* at most 10% (Figure 5). Policy gradients are sharp (Table 6): tightening incumbent caps raises losses by ~3% and widens both the incumbent–new wedge (+11.6 pp) and the A–B gap (+28.1 pp); mobility vouchers reduce losses by ~6% and lift turnover (+4.8 pp) with little change in the wedge; harmonizing rules across jurisdictions cuts losses by ~7% and shrinks the cross-border rent gap by ~52 pp; and a two-tier regime with enforcement closes the wedge (−18.1 pp) while preserving incumbent insurance. Identification checks (full local rank, stable leave-one-out) and robustness to grids and key primitives support these patterns. The accounting is positive rather than prescriptive, but it clarifies which levers matter in fragmented regimes and why.

5 Conclusion

Fragmented rent regulation reallocates pricing power and investment across borders. In our quantitative exercise, most of the welfare gap between decentralization and a planner benchmark arises from muted mobility and under-investment on the incumbent margin; markups and cross-border spillovers play smaller but systematic roles. Counterfactuals map cleanly into these mechanisms: tighter caps raise losses and widen wedges; mobility vouchers lower losses mainly by boosting turnover; light harmonization dampens cross-border arbitrage; and a two-tier regime with enforcement narrows the incumbent–new wedge while preserving incumbent insurance. These patterns align with well-documented episodes and external benchmarks.

Beyond these specific numbers, the paper's broader contribution is a credible way to attribute welfare in dynamic, general-sum environments. By evaluating "neutralized" environments on identical state/action supports with common random numbers, we turn welfare differences into

stable, interpretable objects and separate economics from numerical artifacts. The approach can travels to other regulated durable-goods settings (e.g., energy tariffs, platform fees, credit caps).

Scope and limits. Demand is locally linearized from a DDC foundation; tenants affect outcomes through the transition kernel rather than strategic play; equilibria are computed on fixed finite grids; and entry/new construction is outside the state. These choices deliver transparency and apples-to-apples welfare accounting but bound external validity. Extending the framework to endogenous supply expansion, richer tenant dynamics or bargaining, and new quasi-experiments are natural next steps.

Data and code. A full replication package (grids, seeds, diagnostics, and the CSVs that reproduce all figures and tables) is available at https://github.com/navhoj/DSMP.

Acknowledgments and Disclosure

This research received no external financial support beyond the author's standard PhD stipend and routine university employment (e.g., TA/RA duties). The author reports no financial relationships or other potential conflicts of interest related to this work. Any opinions and conclusions expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

References

Arnott, R. (1995). Time for revisionism on rent control? *Journal of Economic Perspectives*, 9(1):99–120.

Asquith, B. J. (2019). Housing supply dynamics under rent control: What can evictions tell us? *AEA Papers and Proceedings*, 109:393–396.

Autor, D. H., Palmer, C. J., and Pathak, P. A. (2014). Housing market spillovers: Evidence from the end of rent control in cambridge, massachusetts. *Journal of Political Economy*, 122(3):661–717.

Autor, D. H., Palmer, C. J., and Pathak, P. A. (2019). Ending rent control reduced crime in cambridge. *AEA Papers and Proceedings*, 109:381–384.

Bajari, P., Benkard, C. L., and Levin, J. (2007). Estimating dynamic models of imperfect competition. *Econometrica*, 75(5):1331–1370.

Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, 63(4):841.

Chetty, R. (2009). Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. *Annual Review of Economics*, 1(1):451–488.

Costa, G. (2025). Dynamic monopoly in the rental market: Evidence from new york city housing. *available at SSRN 5131722*.

Diamond, R., McQuade, T., and Qian, F. (2019a). The effects of rent control expansion on tenants, landlords, and inequality: Evidence from san francisco. *American Economic Review*, 109(9):3365–3394.

Diamond, R., McQuade, T., and Qian, F. (2019b). Who pays for rent control? heterogeneous landlord response to san francisco's rent control expansion. *AEA Papers and Proceedings*, 109:377–380.

Doraszelski, U. and Pakes, A. (2007). *Chapter 30 A Framework for Applied Dynamic Analysis in IO*, pages 1887–1966. Elsevier.

Fink, A. M. (1964). Equilibrium in a stochastic n-person game. *Journal of science of the hiroshima university, series ai (mathematics)*, 28(1):89–93.

Glaeser, E. and Gyourko, J. (2005). Urban decline and durable housing. *Journal of Political Economy*, 113(2):345–375.

Glaeser, E. L. and Luttmer, E. F. P. (2003). The misallocation of housing under rent control. *American Economic Review*, 93(4):1027–1046.

Gouriéroux, C. and Monfort, A. (1997). *Simulation-based Econometric Methods*. Oxford University PressOxford.

Hahn, A. M., Kholodilin, K. A., Waltl, S. R., and Fongoni, M. (2024). Forward to the past: Short-term effects of the rent freeze in berlin. *Management Science*, 70(3):1901–1923.

Hansen, L. P. (1982). Large sample properties of generalized method of moments estimators. *Econometrica*, 50(4):1029.

Hojaji, S. N., Eross, A., Nguyen, H., and Schaffer, M. E. (2025). Freeze at the bottom, heat at the top: Distributional and cross-border impacts of scotland's rent cap. *Unpublished*.

Hotz, V. J. and Miller, R. A. (1993). Conditional choice probabilities and the estimation of dynamic models. *The Review of Economic Studies*, 60(3):497.

Hsieh, C.-T. and Moretti, E. (2019). Housing constraints and spatial misallocation. *American Economic Journal: Macroeconomics*, 11(2):1–39.

Kholodilin, K. A. (2024). Rent control effects through the lens of empirical research: An almost complete review of the literature. *Journal of Housing Economics*, 63:101983.

McFadden, D. (1974). Conditional logit analysis of qualitative choice behavior. In Zarembka, P., editor, *Frontiers in Econometrics*, pages 105–142. Academic Press.

McFadden, D. (1989). A method of simulated moments for estimation of discrete response models without numerical integration. *Econometrica*, 57(5):995.

McKelvey, R. D. and Palfrey, T. R. (1998). Quantal response equilibria for extensive form games. *Experimental Economics*, 1(1):9–41.

Miravete, E. J., Seim, K., and Thurk, J. (2018). Market power and the laffer curve. *Econometrica*, 86(5):1651–1687.

Nevo, A. (2001). Measuring market power in the ready-to-eat cereal industry. *Econometrica*, 69(2):307–342.

Newey, W. K. and McFadden, D. (1994). *Chapter 36 Large sample estimation and hypothesis testing*, pages 2111–2245. Elsevier.

Pakes, A. and Pollard, D. (1989). Simulation and the asymptotics of optimization estimators. *Econometrica*, 57(5):1027.

Puterman, M. L. (1994). *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley.

Roback, J. (1982). Wages, rents, and the quality of life. *Journal of Political Economy*, 90(6):1257–1278.

Rockafellar, R. T. (1970). *Convex Analysis*. Princeton University Press, Princeton, NJ.

Rosen, S. (1974). Hedonic prices and implicit markets: Product differentiation in pure competition. *Journal of Political Economy*, 82(1):34–55. Publisher: University of Chicago Press.

Rust, J. (1987). Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica*, 55(5):999.

Saiz, A. (2010). The geographic determinants of housing supply*. *Quarterly Journal of Economics*, 125(3):1253–1296.

Shapley, L. S. (1953). Stochastic games. *Proceedings of the National Academy of Sciences*, 39(10):1095–1100.

Stokey, N. L., Lucas, R. E., and Prescott, E. C. (1989). *Recursive Methods in Economic Dynamics*. Harvard University Press.

Topkis, D. M. (1998). *Supermodularity and the Theory of Increasing Returns*. Princeton University Press, Princeton, NJ.

Train, K. E. (2009). *Discrete Choice Methods with Simulation.* Cambridge University Press, Cambridge, 2nd edition.

Wheaton, W. C. (1990). Vacancy, search, and prices in a housing market matching model. *Journal of Political Economy*, 98(6):1270–1292.

Appendix

A Appendix A: Existence of stationary MPE

Proposition A.1 (Existence)**.** *In the finite, discounted stochastic game defined in Section 2 with discount factor $\delta \in (0,1)$, finite state space $\mathcal{S}$, and finite action sets $\{A_\ell(s)\}$, there exists a (possibly mixed) stationary Markov Perfect Equilibrium.*

Proof. By finiteness and discounting, for any stationary profile σ the evaluation operator

$$(T_\sigma V_\ell)(s) = \mathbb{E}_{a\sim\sigma(s)}\Big[\pi_\ell(s,a) + \delta \sum_{s'} P(s' \mid s,a) V_\ell(s')\Big]$$

is a δ-contraction on $(\mathbb{R}^{|\mathcal{S}|}, \|\cdot\|_\infty)$, hence admits a unique fixed point V_ℓ^σ. Let $\mathcal{X}$ be the compact convex product of simplices of stationary strategies. Because π_ℓ and P are multilinear in mixed actions, T_σ depends continuously on σ; by the parametric contraction mapping theorem, V_ℓ^σ is continuous in σ. By Berge's maximum theorem with finite action sets, the state-by-state argmax correspondence is nonempty and upper hemicontinuous. For each $\sigma \in \mathcal{X}$, define the best-response correspondence $\mathrm{BR}_\ell(\sigma)$ to be the set of stationary policies maximising the Bellman objective state-by-state given V^σ. With finite action sets, $\mathrm{BR}_\ell(\sigma)$ is nonempty, convex-valued, and upper hemicontinuous; thus $\mathrm{BR} = \times_\ell \mathrm{BR}_\ell$ has a fixed point by Kakutani. Equivalently, existence follows from Fink's theorem for discounted stochastic games with finite state/action spaces (Fink, 1964). □

B Appendix B: Proof of Proposition 1

Proposition (Rent-cap induced reallocation of markups — restated)**.** *Tightening the incumbent cap (lower $\Delta\bar{r}$) shifts pricing power toward unconstrained/new leases and raises the optimal market rent $r_{mkt,t}$ for those leases: $\frac{\partial r_{mkt,t}}{\partial \Delta\bar{r}} < 0$, holding primitives fixed.*

Proof. Suppress h. Let the incumbent rent evolve by

$$r_{inc,t+1} = \min\{\, r_{inc,t} + \Delta\bar{r},\ r_{mkt,t} \,\}.$$

With switching cost c^{move}, residual demand for an incumbent is less elastic than for a new lease: $\varepsilon_{\text{inc}} < \varepsilon_{\text{new}}$. For a new lease, the landlord sets $r_{mkt,t}$ to maximise static profits $(r_{mkt,t} - MC)D(r_{mkt,t})$ given the mass of new seekers. The mass of new seekers is increasing in the value gap $V^{\text{new}} - V^{\text{stay}}$, which is decreasing in $\Delta\bar{r}$ because tighter caps slow the drift of $r_{inc,t}$ toward $r_{mkt,t}$ and increase the option value of staying (fewer exits). Thus, tightening $\Delta\bar{r}$ reduces the inflow into the new-lease pool, shifting residual demand for new leases left. In a static monopoly problem with $D'(r) < 0$, a left shift in D (at given slope) raises the chosen price if curvature is not too large—formally, $dr^*/d\alpha = -\partial\Pi_r/\partial\alpha \,/\, \Pi_{rr} > 0$ for a left-shift parameter α when $\Pi_{rr} < 0$. Hence $dr_{mkt,t}/d\Delta\bar{r} < 0$.

Dynamically, the Bellman envelope delivers the same sign: the continuation value of a capped incumbent rises as $\Delta\bar{r}$ falls, so the landlord's incentive to induce turnover weakens; scarcity of new leases increases, raising markups on the unconstrained margin.

Formally, let $F(r, \alpha) = D(r; \alpha) + (r - MC)D_r(r; \alpha)$ with $\alpha \equiv \Delta\bar{r}$. If demand is log-concave so that $F_r(r^*, \alpha) = 2D_r + (r^* - MC)D_{rr} < 0$, then by the implicit function theorem

$$\frac{\mathrm{d}r^*}{\mathrm{d}\alpha} = -\frac{F_\alpha}{F_r} = \frac{-\big(D_\alpha + (r^* - MC)D_{r\alpha}\big)}{F_r}.$$

A loosening of the cap (higher α) increases the inflow into the new-lease pool, i.e. $D_\alpha > 0$ (right shift); tightening ($\alpha \downarrow$) shifts left. If $|D_{r\alpha}|$ is bounded (no large rotation)—and for linear D we have $D_{r\alpha} = 0$—then $F_\alpha > 0$ and hence $\mathrm{d}r^*/\mathrm{d}\alpha < 0$. □

C Appendix C: Spillovers across jurisdictions

Proposition (Policy spillovers from regulatory heterogeneity — restated)**.** *Under the assumptions in Proposition 2, loosening the cap in A (higher $\bar{r}_A$) weakly reduces effort and rents in B:*

$$\frac{\partial m^B}{\partial \bar{r}_A} < 0, \qquad \frac{\partial r^B}{\partial \bar{r}_A} < 0.$$

Equivalently, for a per-period cap parameter $\Delta\bar{r}A$ that increases with loosening, $\partial m^B/\partial\Delta\bar{r}A < 0$ and $\partial r^B/\partial\Delta\bar{r}A < 0$.

Lemma C.1 (Tightening in A reallocates effort to B). *Fix $M > 0$. Suppose at the optimum a positive measure of A-units is capped. Then the optimal aggregate effort in B, $m^B = \sum_h m_h^B$, is (weakly) increasing in cap tightness: as $\bar{r}_A$ decreases, m^B weakly increases; equivalently $\partial m^B / \partial \bar{r}_A < 0$.*

Proof. Consider

$$L = \sum_h r_h^A(m_h^A) + \sum_h r_h^B(m_h^B) - \sum_h C(m_h^A) - \sum_h C(m_h^B) - \lambda\left(\sum_h C(m_h^A) + \sum_h C(m_h^B) - M\right).$$

With C strictly convex, KKT are necessary and sufficient (Rockafellar, 1970, Ch. 28). For uncapped units: $\tilde{r}_h^{j\prime}(m_h^j) = C'(m_h^j) + \lambda$. For A-units at the ceiling: $r_h^A(m) = \bar{r}_A$ locally on the right, so $0 \le C'(m_h^A) + \lambda$ with complementary slackness that (weakly) reduces m_h^A to the kink/zero. Lowering $\bar{r}_A$ enlarges the capped set; hence $\sum_h m_h^A$ weakly falls, and the budget implies $\sum_h C(m_h^B)$ weakly rises. By strict convexity of C, this requires $\sum_h m_h^B$ (hence m^B) to (weakly) rise. □

Lemma C.2 (GE price response in B). *Let r^B solve $D^B(r^B, q^B) = S^B$ with $D^B \in C^1$, $D_r^B < 0$, $D_q^B \ge 0$, and $q^B = Q^B(m^B)$ increasing and C^1. Then r^B is weakly increasing in m^B; in particular $\partial r^B / \partial \bar{r}_A \le 0$ when $\partial m^B / \partial \bar{r}_A < 0$.*

Proof. Define $F(r, q) = D^B(r, q) - S^B$. By the Implicit Function Theorem, $r = r(q)$ exists locally and $r_q = -F_q/F_r = -D_q^B/D_r^B \ge 0$ since $D_q^B \ge 0$ and $D_r^B < 0$. Because $q^B = Q^B(m^B)$ with $Q^{B\prime}(m^B) \ge 0$, we have

$$\frac{\partial r^B}{\partial m^B} = r_q\, Q^{B\prime}(m^B) \ \ge\ 0.$$

Combining with $\partial m^B / \partial \bar{r}_A < 0$ (Lemma C.1), the chain rule yields $\partial r^B / \partial \bar{r}_A = (\partial r^B / \partial m^B)\,(\partial m^B / \partial \bar{r}_A) \le 0$. □

Proof of Proposition 2. By Lemma C.1, $\partial m^B / \partial \bar{r}_A < 0$. By Lemma C.2, r^B is weakly increasing in m^B, hence $\partial r^B / \partial \bar{r}_A \le 0$ by the chain rule. Strict inequalities obtain whenever (i) the cap margin in A is active on a set of positive measure so that $\partial m^B / \partial \bar{r}_A \ne 0$, and (ii) $D_q^B > 0$ and $Q^{B\prime} > 0$ on that neighborhood. □

D Appendix D: Renoviction as strategic disinvestment

Proposition (Renoviction threshold — restated). *Under the assumptions in Proposition 3, define q^0, q^{keep}, ΔP_t, G_t and L_V as in the main text. If the threshold inequality*

$$G_t \ >\ \phi_{\text{reno}} + C(m^*) \ +\ L_V \ +\ \frac{C(m^{\text{keep}})}{\delta\, \Delta P_t}$$

holds, then $V_t^{\text{renovict}} > V_t^{\text{maintain}}$.

Proof. Let V_{t+1}^{inc} be the continuation value if the incumbent stays next period and V_{t+3} any continuation two periods after turnover/renovation. Write the choice-specific values:

$$\mathcal{V}_t^{\text{maintain}} = r_{inc,t} - C(m^{\text{keep}}) + \delta\,\mathbb{E}\Big[P_{\text{exit}}(q^{\text{keep}})\big(-\phi_{\text{reno}} - C(m^*) + r_{mkt,t+2} + V_{t+3}\big) + \big(1 - P_{\text{exit}}(q^{\text{keep}})\big)V_{t+1}^{\text{inc}}\Big],$$

$$\mathcal{V}_t^{\text{renovict}} = r_{inc,t} + \delta\,\mathbb{E}\Big[P_{\text{exit}}(q^{0})\big(-\phi_{\text{reno}} - C(m^*) + r_{mkt,t+2} + V_{t+3}\big) + \big(1 - P_{\text{exit}}(q^{0})\big)V_{t+1}^{\text{inc}}\Big].$$

Subtracting gives

$$\begin{aligned}\Delta V_t &:= \mathcal{V}_t^{\text{renovict}} - \mathcal{V}_t^{\text{maintain}} \\ &= C(m^{\text{keep}}) \;+\; \delta\,\Delta P_t\Big[\underbrace{r_{mkt,t+2} - r_{inc,t} - \phi_{\text{reno}} - C(m^*)}_{\text{reset gain net of costs}} + \underbrace{\mathbb{E}[V_{t+3}] - \mathbb{E}[V_{t+1}^{\text{inc}}]}_{\text{bounded by } L_V}\Big].\end{aligned}$$

Using the bound $\mathbb{E}[V_{t+3}] - \mathbb{E}[V_{t+1}^{\text{inc}}] \geq -L_V$ (finite discounted state space; Puterman, 1994, Ch. 6), we obtain

$$\Delta V_t \;\geq\; C(m^{\text{keep}}) \;+\; \delta\,\Delta P_t\big(G_t - \phi_{\text{reno}} - C(m^*) - L_V\big).$$

Hence $\Delta V_t > 0$ whenever the threshold in the statement holds, which proves the claim. □

Remark 2 (Comparative statics)**.** *The left side of* (8) *is increasing in the rent gap* G_t*; the right side increases in* ϕ_{reno}*,* $C(\cdot)$*, and* L_V *and decreases in* ΔP_t*. Thus the renoviction region (states satisfying* (8)*) expands when caps tighten (raising* G_t *by lowering* $r_{inc,t}$*), when maintenance costs rise, or when* ΔP_t *increases (steeper* P_{exit} *or larger maintenance needed to keep the tenant). These monotone effects follow from lattice-theoretic MCS; see Topkis (1998).*

E Appendix E: Computational details and robustness

E.1 Solution concepts, existence, and selection

Proposition E.2 (Existence of stationary logit–QRE for $\tau > 0$)**.** *In the finite, discounted baseline game (Section 2.5), for any fixed temperature* $\tau > 0$ *there exists at least one stationary logit quantal response equilibrium* σ_τ*.*

Proof. Let $\mathcal{X}$ be the compact, convex product of simplices of stationary mixed policies. For any stationary profile $\sigma \in \mathcal{X}$, the evaluation operator

$$(T_\sigma V)(s) \;=\; \mathbb{E}_{a\sim\sigma(s)}\Big[\pi(s,a) + \delta\sum_{s'} P(s' \mid s,a)V(s')\Big]$$

is a δ-contraction on $(\mathbb{R}^{|\mathcal{S}|}, \|\cdot\|_\infty)$, so V^σ is unique. Writing Q^σ for the state–action value vector induced by V^σ, the logit response $\mathrm{SoftMax}_\tau(Q^\sigma)$ is continuous in Q^σ and therefore in σ by finiteness of $\mathcal{S}$ and action sets. Hence $\Phi_\tau : \sigma \mapsto \mathrm{SoftMax}_\tau(Q^\sigma)$ is a continuous self-map of $\mathcal{X}$. Brouwer's fixed-point theorem yields a fixed point $\sigma_\tau = \Phi_\tau(\sigma_\tau)$.[1] □

Remark 3 (Evaluation contraction)**.** *For fixed σ,*

$$\|T_\sigma V - T_\sigma W\|_\infty \le \delta \|V - W\|_\infty$$

because $\sum_{s'} P(s' \mid s, a) = 1$ and $\delta \in (0, 1)$.

Relation to MPE and selection by homotopy. Stationary MPE existence for finite discounted stochastic games is standard (Fink, 1964). For computation and selection, we solve stationary logit–QREs along a temperature path $\tau \downarrow 0$ (Section 4.2.1; Algorithm 2; McKelvey and Palfrey, 1998) and use the small-τ limit as our selection device. In our applications the limit policy is (approximately) pure when a pure MPE exists; with discrete tie-breaking (lexicographic), the small-τ policy is stable and reproducible.

E.2 State compression and transitions

Each unit's state is

$$(j,\ r_{inc},\ q,\ p,\ \zeta,\ T),$$

with $j \in \{A, B\}$ (jurisdiction), $r_{inc} \in \mathcal{G}_j^r$ (incumbent rent grid), $q \in \mathcal{Q}$ (quality grid), permit status $p \in \{0, \dots, \psi_j^{\text{permit}}\}$ when delays are active, contract status $\zeta \in \{\text{new}, \text{inc}\}$, and occupancy $T \in \{0, 1\}$. Quality evolves via the discretised law $q_{t+1} = \Gamma(q_t, m_t; j)$; permit status p increments deterministically until completion. The induced Markov kernel is time-homogeneous with fixed finite support. This state and kernel structure is unchanged across counterfactuals and Shapley neutralisations (Lemma 2), ensuring "apples-to-apples" welfare comparisons.

E.3 Equilibrium computation

Policy evaluation. For a stationary profile σ, we compute V^σ either by (i) a direct sparse solve of $(I - \delta P_\sigma)V = b_\sigma$ or (ii) value iteration to sup-norm tolerance $< 10^{-10}$, whichever is faster. Both

[1] For dynamic programming background, see Puterman (1994, Chs. 4–6) and Stokey et al. (1989).

use sparse storage for P_σ.

Best responses and damping. At each (ℓ, s) we form $Q_\ell(s, \cdot)$ and update policies with damping $\eta \in [0.2, 0.5]$:

$$\sigma^{(k+1)} \leftarrow (1-\eta)\,\sigma^{(k)} + \eta\,\mathrm{BR}\big(Q^{\sigma^{(k)}}\big),$$

or, under the QRE operator, replace BR with $\mathrm{SoftMax}_\tau$ (Proposition E.2). Ties in $\arg\max$ are broken lexicographically to ensure measurability and reproducibility.

Homotopy schedule. We initialise at $\tau_0 \in [0.5, 1]$ and decrease on a grid $\{\tau_g\}$ to $\tau_{\min} \in [10^{-3}, 10^{-2}]$, warm-starting each solve at the previous fixed point (Algorithm 2). The small-τ policy approximates an MPE and prevents cycling in general-sum games.

Stopping criteria and diagnostics. We monitor (i) policy distance $d^{(k)} = \max_{\ell,s} \|\sigma_\ell^{(k+1)}(s) - \sigma_\ell^{(k)}(s)\|_1$ and (ii) Bellman residual $\rho^{(k)} = \max_{\ell,s} \big|V_\ell^{(k)}(s) - (T_{\sigma^{(k)}} V_\ell^{(k)})(s)\big|$. We declare convergence when $d^{(k)} < \varepsilon_\pi$ and $\rho^{(k)} < \varepsilon_V$ with thresholds as in Table 4. We also record iteration counts and wall-clock time; all diagnostics are saved to `baseline_diag.json` (Section 4.9).

E.4 Calibration, perturbations, and reproducibility

Common random numbers (CRN). All simulations (baseline, Jacobians, neutralisations, and counterfactuals) reuse the same seeds, burn-in, and replication counts to isolate parameter/policy effects from Monte Carlo noise.[1]

Finite differences for the Jacobian. For $G(\hat{\Theta}) = \partial m(\Theta)/\partial \Theta'|_{\hat{\Theta}}$ we use central differences with $h_p = \rho \cdot \max\{|\hat{\Theta}_p|, 1\}$ and $\rho \in \{10^{-5}, 10^{-4}, 10^{-3}\}$; entries are deemed step-size robust if they vary by $< 5\%$ across ρ (Appendix G).

Artifacts and logs. We persist (i) seeds, burn-in, R; (ii) $\{\tau_g\}$, η, tolerances; (iii) grid sizes; (iv) diagnostic traces $\{d^{(k)}, \rho^{(k)}\}$; and (v) run times. Output file schemas follow Table 7. Figure E6 shows the pipeline.

[1] Using CRN in simulation-based estimation and derivatives is standard; see McFadden (1989); Pakes and Pollard (1989); Newey and McFadden (1994).

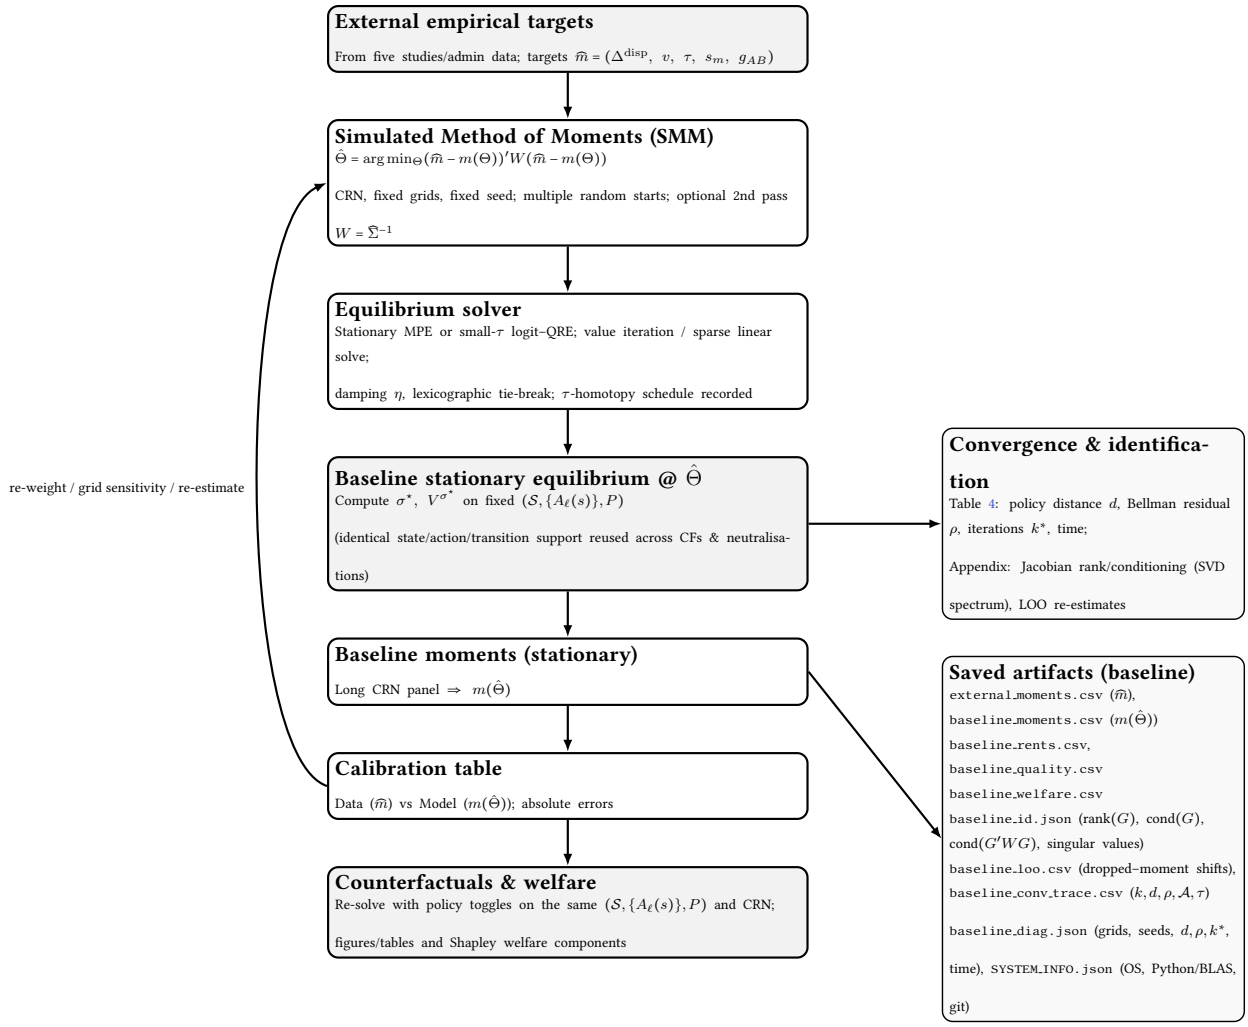


Figure E6: Calibration and replication pipeline. External empirical targets ($\widehat{m}$) inform SMM estimation of $\hat{\Theta}$ using a stationary equilibrium solver (MPE or small-τ QRE) with common random numbers, fixed grids, and fixed seeds. We then simulate the baseline stationary equilibrium to compute $m(\hat{\Theta})$, compare to $\widehat{m}$, and report convergence and identification diagnostics. *All counterfactuals and neutralisations reuse the same state space, action sets, transition support, and CRN (apples-to-apples).* Saved artifacts + system stamp ensure full replication; a one-click script rebuilds everything.

E.5 Robustness protocols and checks

Grid coarseness/refinement. We re-solve on $|\mathcal{G}_j^r| \in \{31, 41, 51\}$, $|\mathcal{M}| \in \{6, 8, 11\}$, and $|\mathcal{Q}| \in \{15, 20, 25\}$ while keeping the model structure unchanged. Across these grids, (a) the ordering

of total welfare across counterfactuals, (b) the signs of Shapley components $\{\Delta W_k\}$ (Section 4.7), and (c) the qualitative movement of targeted moments (Section 4.3) remain stable within the thresholds reported in Table 9.

Tornado and stacked-welfare figures. Figures E.1–E.3 report one-at-a-time perturbations to Θ (plug-in and recalibrated) and stacked Shapley decompositions by channel; the ordering of magnitudes and the signs of $\{\Delta W_k\}$ are stable. (Table 9)

Optional check A: Equilibrium selection from multiple initialisations. To probe multiplicity, we re-ran Algorithm 1 (or Algorithm 2 at the terminal $\tau_{\min}$) from three distinct starts: uniform random, myopic best responses, and a perturbed baseline policy. In all runs we converged to the same small-τ fixed point within the thresholds of Table 4. A short summary table (E.4) reports final policy distances between solutions from different starts (all $< 10^{-6}$) and identical objective values.

Optional check B: KDE bandwidth sensitivity for rent densities. For Fig. 3, we varied the kernel bandwidth using (i) Silverman's rule-of-thumb, (ii) Scott's rule, and (iii) a fixed common bandwidth across panels. The modal locations, the incumbent–new wedge, and CF-vs-baseline differences are invariant to these choices (Figure E.5 shows overlays); numerical comparisons use the common fixed bandwidth stored in `baseline_diag.json`.

E.6 Baseline parameters

The primitives, units, and calibration status are summarised in Table E10. Estimates $\hat{\Theta}$ are obtained by SMM against Table 1. Units are per model period (e.g., monthly), consistent with Section 4.1.

Table E10: Parameters, units, and identification status (point values in replication files)

Parameter	Symbol	Unit	Fixed / SMM (target)
Demand intercept (A)	A_A	rent/period	SMM (avg rent A)
Demand slope (A)	B_A	1/quantity	SMM (elasticity)
Demand intercept (B)	A_B	rent/period	SMM (avg rent B)
Demand slope (B)	B_B	1/quantity	SMM (elasticity)
Switching cost	c^{move}	months rent	SMM (turnover)
Depreciation	δ_q	per period	Fixed (engineering prior)
Maint. effectiveness	χ	q per unit m	SMM (share m>0)
Maint. cost linear	κ_1	rent units	SMM (mean m)
Maint. cost convex	κ_2	rent units	SMM (dispersion m)
Rent cap (A)	$\Delta\bar{r}A$	per period	Fixed (policy)
Rent cap (B)	$\Delta\bar{r}B$	per period	Fixed (policy)
Eviction cost (A)	ϕ_A^{evict}	months rent	SMM (vacancy/tenure)
Eviction cost (B)	ϕ_B^{evict}	months rent	Fixed / SMM (robustness)
Permit delay (A)	ψ_A^{permit}	periods	Fixed (policy)
Permit delay (B)	ψ_B^{permit}	periods	Fixed (policy)
Discount factor	δ	per period	Fixed (0.97–0.99)
Rent grid size (A,B)	$\lvert\mathcal{G}_j^r\rvert$	points	Fixed (grid)
Maint. grid size	$\lvert\mathcal{M}\rvert$	points	Fixed (grid)
Quality grid size	$\lvert\mathcal{Q}\rvert$	points	Fixed (grid)

Notes: Point estimates/values are recorded in the replication archive (e.g., `baseline_diag.json`); we report units and identification here for space.

F Appendix F: Linear demand as a first–order approximation

Setup and notation

In Section 4.1 we use a reduced–form residual demand $D_j(r,q) = A_j - B_j r + \gamma_j q$ for new leases in jurisdiction j. This appendix provides a precise microfoundation and linearisation from a standard dynamic discrete–choice (DDC) environment.

Consider a mass M_j of searchers choosing among $\mathcal{C}_j$ (landlord–unit alternatives plus an outside option). Period utility for alternative $h \in \mathcal{C}_j$ is

$$U_i(h) \;=\; \mu_j(h) \;+\; \xi(q_h) \;-\; \beta\, r_h \;+\; \varepsilon_{ih}, \qquad \varepsilon_{ih} \overset{\text{i.i.d.}}{\sim} \text{Type-I EV},$$

where $\mu_j(h)$ includes match and continuation terms (DDC inclusive value), $\xi(\cdot)$ is C^1, and $\beta > 0$ is the price coefficient. Under stationarity and a marginal perturbation of (r_h, q_h) on one alternative, the continuation component in $\mu_j(h)$ is locally invariant relative to the direct terms. Hence the per–period share of h is logit:

$$s_j(h) \;=\; \frac{\exp\big(\mu_j(h) + \xi(q_h) - \beta r_h\big)}{\sum_{k\in\mathcal{C}_j} \exp\big(\mu_j(k) + \xi(q_k) - \beta r_k\big)}.$$

Define residual (flow) demand as $D_j^{\text{DDC}}(r_h, q_h) = M_j\, \Lambda_j\, s_j(h)$, where $\Lambda_j \in (0,1]$ captures dynamic attenuation from switching costs and timing (a reduced pass–through from per–period shares to lease flows).[1]

Lemma F.3 (Linear residual demand is first–order accurate). *Fix a baseline (r_0, q_0) for alternative h and write $s_0 = s_j(h)\big|_{(r_0,q_0)}$. Suppose ξ is C^1 and Λ_j is constant in a neighborhood of (r_0, q_0) for a single–alternative perturbation.*[2] *Then*

$$D_j^{\text{DDC}}(r,q) \;=\; A_j \;-\; B_j\, r \;+\; \gamma_j\, q \;+\; o\big(\,\|(r,q) - (r_0,q_0)\|\,\big)\,,$$

[1] Foundational logit/choice references: McFadden (1974); Train (2009). DDC foundations: Rust (1987); Hotz and Miller (1993).

[2] If Λ_j varies with (r,q) through aggregates that are unaffected to first order by a *single* alternative's perturbation, the same result holds. If Λ_j is C^1 in (r,q) and responds to the single–alternative perturbation, its partials enter the coefficients via product rule; see Remark 4.

with coefficients

$$B_j = \beta\, M_j \Lambda_j\, s_0(1-s_0), \qquad \gamma_j = \xi'(q_0)\, M_j \Lambda_j\, s_0(1-s_0), \qquad A_j = D_j^{\mathrm{DDC}}(r_0, q_0) + B_j r_0 - \gamma_j q_0.$$

Proof. Differentiate the logit share w.r.t. (r_h, q_h) at (r_0, q_0):

$$\frac{\partial s_j(h)}{\partial r_h}\Big|_0 = -\beta\, s_0(1-s_0), \qquad \frac{\partial s_j(h)}{\partial q_h}\Big|_0 = \xi'(q_0)\, s_0(1-s_0).$$

Since $D^{\mathrm{DDC}}(r,q) = M_j \Lambda_j\, s_j(h)$ and Λ_j is locally constant by assumption, we have

$$\frac{\partial D^{\mathrm{DDC}}}{\partial r}\Big|_0 = -\,\beta\, M_j \Lambda_j\, s_0(1-s_0) = -B_j, \qquad \frac{\partial D^{\mathrm{DDC}}}{\partial q}\Big|_0 = \xi'(q_0)\, M_j \Lambda_j\, s_0(1-s_0) = \gamma_j.$$

Apply the first–order Taylor expansion at (r_0, q_0) and collect terms to obtain the stated affine form; the remainder is $o(\|(r,q) - (r_0,q_0)\|)$ by differentiability. □

Remark 4 (When Λ_j varies with (r,q))**.** *If Λ_j is C^1 in (r,q) and responds to the single–alternative perturbation, the coefficients become*

$$B_j = \big[\beta\, s_0(1-s_0)\big] M_j \Lambda_j \;-\; \big[\partial \Lambda_j/\partial r\big]_0 M_j s_0, \quad \gamma_j = \big[\xi'(q_0)\, s_0(1-s_0)\big] M_j \Lambda_j \;+\; \big[\partial \Lambda_j/\partial q\big]_0 M_j s_0,$$

evaluated at (r_0, q_0). If Λ_j depends only on aggregates that are unaffected to first order by a single–alternative perturbation (e.g., jurisdiction–wide congestion), the partials vanish and Lemma F.3 applies as stated.

Interpretation and use. The slope B_j scales with the logit semi–elasticity at the baseline and market size M_j, attenuated by Λ_j. The quality slope γ_j scales with the marginal utility of quality $\xi'(q_0)$. In estimation we treat (A_j, B_j, γ_j) as reduced–form parameters; Lemma F.3 rationalises them as first–order coefficients around the baseline fixed point used in Section 4.1. This approximation is local; we therefore use it for calibration and for reporting elasticities near the observed equilibrium, not for large–deviation counterfactual extrapolation.

Cross–references. The linear residual demand underpins (i) the pricing comparative statics in Proposition 1, (ii) the construction of rent grids $\mathcal{G}_j^r$ in Section 4.1, and (iii) the SMM calibration in Section 4.5.

G Appendix G: Identification diagnostics and leave–one–out (LOO) details

Notation and scope

Let K be the number of targeted moments and P the number of estimated parameters. Denote by $m(\Theta) \in \mathbb{R}^K$ the vector of simulated equilibrium moments (computed with common random numbers), by $\widehat{m} \in \mathbb{R}^K$ the empirical moments, and by $\hat{\Theta}$ the SMM estimate in §4.5. Define the Jacobian at $\hat{\Theta}$ as

$$G(\hat{\Theta}) := \left.\frac{\partial m(\Theta)}{\partial \Theta'}\right|_{\Theta=\hat{\Theta}} \in \mathbb{R}^{K\times P}.$$

All derivatives are obtained by finite differences using *common random numbers* (CRN) and, unless otherwise stated, at the small–temperature QRE used for equilibrium smoothing (Appendix E); evaluating at a small $\tau > 0$ avoids kinks due to discrete best–response ties and improves numerical stability.

Finite–difference settings and CRN

We compute $G(\hat{\Theta})$ by *central differences* parameter–by–parameter with relative step

$$h_p = \rho \cdot \max\{|\hat{\Theta}_p|, 1\}, \qquad \rho \in \{10^{-5}, 10^{-4}, 10^{-3}\}.$$

For each $p \in \{1, \dots, P\}$,

$$G_{\cdot p}(\hat{\Theta}) \approx \frac{m\big(\hat{\Theta} + h_p e_p\big) - m\big(\hat{\Theta} - h_p e_p\big)}{2h_p},$$

reusing the *same* seeds, burn–in, and simulation draws across $\hat{\Theta} \pm h_p e_p$ (CRN) to suppress Monte Carlo noise.[1] If a bound is hit for some parameter (e.g., positivity), we fall back to a one–sided difference on that coordinate. We declare step–size robustness if every entry of G varies by $< 5\%$ across $\rho \in \{10^{-5}, 10^{-4}, 10^{-3}\}$.

Rank, conditioning, and elasticities

We compute the SVD $G = U\Sigma V'$ with singular values $\sigma_1 \geq \cdots \geq \sigma_{\min\{K,P\}}$. Numerical rank is

$$r := \#\{\sigma_i : \sigma_i/\sigma_1 \geq \tau_{\text{rank}}\}, \qquad \tau_{\text{rank}} = 10^{-8}.$$

[1] Using CRN in simulation–based derivatives is standard

We report condition numbers $\mathrm{cond}(G) = \sigma_1/\sigma_r$ and $\mathrm{cond}(G'WG)$, where W is the SMM weighting matrix (diagonal unless noted). As a scale–free diagnostic, we also report moment–parameter *elasticities*

$$\mathcal{E}_{kp} := \frac{\hat{\Theta}_p}{\widehat{m}_k} G_{kp},$$

which make the Jacobian comparable across units and are what we display in Table G12.

Leave–one–out (LOO) robustness

For each moment $k \in \{1, \ldots, K\}$, re–estimate

$$\hat{\Theta}^{(-k)} = \arg\min_{\Theta} \big(\widehat{m}_{-k} - m_{-k}(\Theta)\big)' W_{-k} \big(\widehat{m}_{-k} - m_{-k}(\Theta)\big),$$

recompute $G^{(-k)}$ with the same CRN and step–size settings, and store diagnostics:

$$\big\|\hat{\Theta}^{(-k)} - \hat{\Theta}\big\|_\infty, \quad \Delta Q := Q(\hat{\Theta}^{(-k)}) - Q(\hat{\Theta}), \quad \mathrm{rank}\big(G^{(-k)}\big), \quad \mathrm{cond}\big(G^{(-k)}\big), \quad \max_{k,p} \frac{\big|G^{(-k)}_{kp} - G_{kp}\big|}{|G_{kp}| + \varepsilon},$$

with $\varepsilon = 10^{-10}$ for numerical safety. We flag FAIL if: (i) rank drops ($\mathrm{rank}(G^{(-k)}) < r$); or (ii) $\|\hat{\Theta}^{(-k)} - \hat{\Theta}\|_\infty$ exceeds a pre–set threshold (e.g., 20% of the baseline value on any coordinate); or (iii) $\mathrm{cond}(G^{(-k)})$ exceeds 10^6.

Table G11: Leave-one-out moment robustness

Dropped moment	$\|\hat{\Theta}^{(-k)} - \hat{\Theta}\|_\infty$	ΔQ	$\mathrm{Rank}(G^{(-k)})$
Δ^{disp}	0.029	9.31e-02	3
v	0.023	8.41e-02	3
τ	0.023	1.70e-02	3
s_m	0.005	8.63e-02	2
g_{AB}	0.023	8.84e-02	3

Table G12: Moment–parameter elasticity matrix at $\hat{\Theta}$

	B (demand semi-elasticity)	c^{move} (moving cost)	κ_2 (maint. curvature)
Δr (new–inc, pp)	-0.100	0.001	-0.002
v (%)	0.000	-0.207	0.000
τ (%/yr)	0.000	-0.013	0.000
s_m (%)	0.000	-0.000	-0.354
g_{AB} (pp)	0.001	-0.000	-0.028

Notes: Elasticities $\mathcal{E}_{kp} \approx \dfrac{\Delta m_k / m_k^{\text{ref}}}{\varepsilon_p}$ from one-at-a-time multiplicative perturbations using common random numbers.

When rank or conditioning is fragile

If $\text{rank}(G) < P$ or $\text{cond}(G)$ is large ($> 10^6$), we augment the target set with one or more *optional* moments described in §4.5, prioritising moments that load strongly and differentially on weakly identified parameters. In practice, effective additions are: (i) evictions per 100 leases (loads on ϕ_j^{evict}); (ii) a tenure percentile (loads on c^{move} and $\Delta\bar{r}j$); (iii) a cross–border rent gradient (loads on spillover and demand slopes B_j). We then re–estimate $\hat{\Theta}$, recompute G, and repeat the diagnostics above.

Reproducibility

All Jacobian and LOO runs use the exact same simulation primitives as the baseline: (i) common random numbers (seeds, burn–in, replication counts); (ii) the same discretisation $(\mathcal{S}, \{A_\ell(s)\})$; (iii) the same QRE temperature used in Appendix E for derivative calculations (with a very small $\tau > 0$). This guarantees that changes reported in Tables G12 and G11 are due to parameter perturbations rather than Monte Carlo noise or selection of different stationary equilibria.

H Appendix H: DDC-to-linear mapping

Proof of Lemma F.3 and the expressions linking (A_j, B_j, γ) to the DDC primitives; numerical illustration omitted for brevity.

I Appendix I: Monotone comparative statics for policy signs

Assumptions (MCS). (i) Per-period payoffs are supermodular in landlords' actions and have increasing differences in (actions; τ) and in (actions; $-c^{\text{move}}$); eviction/friction parameters enter with the indicated monotone direction. (ii) State transitions are isotone in actions; tenants' DDC induces isotone occupancy/matching kernels in (r, q). (iii) Strategy spaces are complete lattices; we select extremal stationary MPE (Tarski).

Results and proof sketches. By Topkis's theorem, each landlord's best-response correspondence is ascending in τ and in $-c^{\text{move}}$, hence the least/greatest stationary MPE are monotone in these parameters (Tarski's fixed-point theorem). Therefore: (a) $\tau \uparrow \Rightarrow m \uparrow$ on regulated units, $q \uparrow$, lowering the investment distortion; (b) $c^{\text{move}} \downarrow \Rightarrow$ higher turnover on the relevant margins, lowering the mobility distortion; (c) harmonising $(\Delta\bar{r}j, \phi_j^{\text{evict}})$ removes spatial gradients that drive reallocation, lowering spillovers; (d) two-tier rules reduce price dispersion holding incumbent protection, lowering the markup component. Formal lattice arguments follow Topkis (1998); existence/selection of stationary MPE follows Appendix A (Kakutani/Fink).

J Appendix J: Welfare decomposition algebra and neutrality maps

Neutrality maps $V(S)$ and parameter substitutions

Throughout, the baseline finite environment has fixed $(\mathcal{S}, \{A_\ell(s)\}, P_\theta)$ with parameter vector θ. For any $S \subseteq \mathcal{K}$ with $\mathcal{K} = \{\text{investment}, \text{allocation}, \text{mobility}, \text{markups}, \text{spillover}\}$, define E_S by replacing θ with $\theta(S)$ via *parameter substitutions only* (no change to $\mathcal{S}$, $\{A_\ell(s)\}$, or the support of P_θ). The stationary MPE of E_S yields welfare $V(S)$.

Investment. On regulated units, replace $C(m)$ by $C(m) - \tau^\star m$ with $\tau^\star \geq 0$ calibrated (jurisdiction/type–specific constants allowed) so that landlords' stationary best responses in m coincide with the planner's maintenance profile on those units. No other primitives change.

Allocation. Replace the decentralised matching/mobility kernel by its planner-equivalent assignment rule encoded as a parameterised kernel with the *same support*; i.e., weights are shifted to implement $\mu^{\text{planner}}_{i,h,t}(s_t)$ while preserving which transitions are feasible. Prices/policies remain as in the baseline unless neutralised by other channels.

Mobility. Set the moving cost to the planner's effective level (first-best: $c^{\text{move}} \to 0$) holding all other primitives fixed.

Markups. Introduce a quadratic pricing instrument (a policy parameter, not an action) that makes the best response in posted rent for new leases coincide with the planner's Ramsey/MC schedule $r^{\text{planner}}(q)$:

$$\Pi_h(r, q; \psi) \;=\; \underbrace{(r - MC(q))D(r, q)}_{\text{revenue}} \;-\; \underbrace{\psi\left(r - r^{\text{planner}}(q)\right)^2}_{\text{policy instrument}}, \qquad \psi \gg 0.$$

This keeps the action set unchanged and implements $r = r^{\text{planner}}(q)$ as the unique optimum; incumbent-cap rules are unaltered.

Spillover. Harmonise $(\Delta\bar{r}_j, \phi^{\text{evict}}_j)$ across jurisdictions to planner-equivalent values; permit lags can be similarly aligned if used in the planner. Portfolio sizes and action sets are unchanged.

By construction, $V(\varnothing) = V^{\text{MPE}}(s_0)$ and $V(\mathcal{K}) = V^{\text{planner}}(s_0)$; all E_S share the same $(\mathcal{S}, \{A_\ell(s)\})$ and the same transition *support*, so comparisons are apples-to-apples (cf. Lemma 2).

Proof of Proposition 5 (Shapley identity)

Let $v(S) := V(S)$ for $S \subseteq \mathcal{K}$, with $v(\varnothing) = V^{\text{MPE}}$ and $v(\mathcal{K}) = V^{\text{planner}}$. For any permutation π of $\mathcal{K}$ and channel k, write $S_{\pi,k}$ for the set of channels preceding k in π. Define the marginal

$$\Delta W_k(\pi) = v(S_{\pi,k} \cup \{k\}) - v(S_{\pi,k}), \qquad \Delta W_k = \frac{1}{|\Pi|}\sum_{\pi} \Delta W_k(\pi).$$

Then

$$\sum_{k\in\mathcal{K}} \Delta W_k = \frac{1}{|\Pi|}\sum_{\pi}\sum_{k\in\mathcal{K}} \Big(v(S_{\pi,k} \cup \{k\}) - v(S_{\pi,k})\Big) = \frac{1}{|\Pi|}\sum_{\pi}\Big(v(\mathcal{K}) - v(\varnothing)\Big) = v(\mathcal{K}) - v(\varnothing) = \Delta W,$$

where the middle equality is a telescoping sum along each permutation π. This is the standard Shapley telescoping argument (Shapley, 1953), yielding the exact additive decomposition $\Delta W = \sum_{k\in\mathcal{K}} \Delta W_k$.

K Appendix K: Measurement operator and simulation/calibration details

This appendix documents the exact objects that generate Table 1. All moments are computed from a stationary unit×time panel produced by our simulator. For each unit h and period t, let $T_{h,t} \in \{0,1\}$ denote occupancy (1=occupied), $\text{status}_{h,t} \in \{\text{new}, \text{inc}, \varnothing\}$ the lease status, $j(h) \in \{A, B\}$ the jurisdiction, $r_{h,t}$ the recorded rent, $r_{\text{inc},h,t}$ the incumbent rent when defined, $m_{h,t} \in \{0,1\}$ a binary maintenance decision, and $q_{h,t} \in [\underline{q}, 1]$ the quality stock.

Occupancy dynamics. Each unit follows a two–state Markov chain with transition probabilities

$$p_{VI} = (1-v)\lambda, \qquad p_{IV} = v\lambda,$$

so the stationary vacancy obeys the identity

$$v^* = \frac{p_{IV}}{p_{IV} + p_{VI}} = v.$$

A period-t lease is *new* if a vacancy is filled ($V \to I$) or if an incumbent is replaced ($I \to I$) with probability θ. Here $\theta \in [0,1]$ and $\lambda > 0$ summarise replacement and meeting intensity.

Jurisdictional rent levels and dispersion. New-lease rents have jurisdiction-specific means,

$$\bar{R}_A, \qquad \bar{R}_B = \bar{R}_A + g_{AB},$$

with small i.i.d. shocks $\varepsilon_{h,t}$ for simulation efficiency. Incumbent rents are lower on average by a dispersion parameter d:

$$\mathbb{E}[\, r_{h,t} \mid \text{new},\, j(h) = j \,] = \bar{R}_j, \qquad \mathbb{E}[\, r_{\text{inc},h,t} \mid \text{inc},\, j(h) = j \,] = \bar{R}_j - d.$$

(In the full DP model, d is the implied average wedge from the stationary quality distribution.)

Maintenance (DP block). Quality evolves as

$$q_{h,t+1} = \begin{cases} (1-\delta_q)\, q_{h,t}, & \text{if } m_{h,t} = 0, \\ 1, & \text{if } m_{h,t} = 1, \end{cases}$$

with depreciation rate $\delta_q \in (0,1)$ and maintenance cost c_m when $m_{h,t} = 1$. The optimal policy is a single threshold in q, yielding a stationary maintenance share $\mathbb{E}[m_{h,t}]$.

Measurement operator (used on both data and model). Given $\{T_{h,t}, \text{status}_{h,t}, j(h), r_{h,t}, r_{\text{inc},h,t}, m_{h,t}\}$:

1. **Rent dispersion (pp).**

$$RD = \mathbb{E}[\, r_{h,t} \mid T_{h,t} = 1,\ \text{status}_{h,t} = \text{new}\,] - \mathbb{E}[\, r_{\text{inc},h,t} \mid T_{h,t} = 1,\ \text{status}_{h,t} = \text{inc}\,].$$

2. **Vacancy rate (%).** $\mathbb{E}[\, 1 - T_{h,t}\,]$.

3. **Turnover rate (%/yr).** Let $\text{churn}_t = \sum_h \mathbf{1}\{T_{h,t} = 1,\ \text{status}_{h,t} = \text{new}\}$ and $\text{occ}_{t-1} = \sum_h T_{h,t-1}$, then

$$\tau = \frac{\sum_t \text{churn}_t}{\sum_t \text{occ}_{t-1}}.$$

4. **Share with maintenance $m > 0$ (%).** $s_m = \mathbb{E}[\, m_{h,t}\,]$.

5. **Jurisdictional rent gap A vs. B (pp).**

$$g_{AB} = \mathbb{E}[\, r_{h,t} \mid T_{h,t} = 1,\ \text{status}_{h,t} = \text{new},\ j(h) = B\,] - \mathbb{E}[\, r_{h,t} \mid T_{h,t} = 1,\ \text{status}_{h,t} = \text{new},\ j(h) = A\,].$$

Calibration. Let the five targets be $(RD, v, \tau, s_m, g_{AB})$. We set λ to satisfy the vacancy identity for the observed v, and choose (θ, d, c_m) by minimum–distance (SMM) to match (RD, τ, s_m), while g_{AB} and v are imposed by construction. All reported moments use the same operator as above. For identification diagnostics and leave–one–out robustness, see Appendix G.

L Notation

Note on grids vs. policy vector. We use $\mathcal{R}_j^{\text{pol}}$ for the jurisdictional policy vector (caps, eviction frictions, permit lags). To avoid overloading, we denote discretisation grids by $\mathcal{G}_j^r$ (rents), $\mathcal{G}^m$ (maintenance), and $\mathcal{G}^q$ (quality).

Table L13: Notation used throughout the paper

Symbol	Meaning / Units
Indices, sets, states	
	Continued on next page

Symbol	Meaning / Units
$i \in [0,1]$	Tenant index (continuum).
$\ell \in \mathcal{L}$	Landlord index; $\mathcal{L}$ set of landlords.
$h \in \mathcal{H}$	Housing unit; $\mathcal{H}$ set of units.
$j \in \mathcal{J}$	Jurisdiction/neighbourhood (e.g., A, B); $\mathcal{J}$ set of jurisdictions.
$t = 0, 1, \ldots$	Time (discrete).
s_t	Aggregate state at t (portfolio microstates, occupancy, qualities, policies).
$T_{h,t} \in \{0,1\}$	Occupancy indicator for unit h at t.
$j(h)$	Jurisdiction of unit h.
Prices, choices, qualities	
$r_{h,t}$	Rent posted/charged for unit h at t (generic).
$r_{inc,h,t}$	Incumbent regulated rent for unit h at t (subject to cap).
$r_{mkt,j,t}$	Market rent benchmark for new leases in j at t (when used without j, read as local).
$m_{h,t} \geq 0$	Maintenance/upgrade effort.
$e_{h,t} \in \{0,1\}$	Eviction attempt/decision (1 if attempted).
$q_{h,t} \in \mathbb{R}_+$	Quality of unit h at t.
$\theta_{i,h,t}$	Match-specific utility draw for tenant i and unit h.
η_j	Jurisdiction amenity shifter.
$\mathrm{dist}(h)$	Location disutility proxy (e.g., distance/commute).
$\kappa > 0$	Per-unit disutility weight on distance: utility includes $-\kappa\,\mathrm{dist}(h)$.
Regulatory/policy parameters	
$\mathcal{R}_j^{\mathrm{pol}} = \{\Delta\bar{r}_j,\ \phi_j^{\mathrm{evict}},\ \psi_j^{\mathrm{permit}}\}$	Policy vector in j: cap on rent increases, eviction frictions, permit lags.
$\Delta\bar{r}_j$	Cap on allowable rent increases for incumbents in j (pp or %).
ϕ_j^{evict}	Cost/delay of eviction in j.

Continued on next page

Symbol	Meaning / Units
ψ_j^{permit}	Permit/approval lag/cost for renovation in j.
$c^{\text{move}} > 0$	Tenant switching/moving cost.
Technology, dynamics, demand	
$\delta \in (0,1)$	Discount factor (agents/planner).
$\delta_q \in (0,1)$	Physical depreciation rate of quality.
$\chi > 0$	Effectiveness of maintenance in the quality law of motion.
$C(m)$	Maintenance/upgrade cost function; convex, increasing.
$\pi_{h,t} = r_{h,t} T_{h,t}$	Rental revenue for unit h at t.
$q_{h,t+1} = (1-\delta_q) q_{h,t} + \chi m_{h,t} - \mathbb{1}_{\text{permit}} \psi_j^{\text{permit}}$	Quality transition (permit friction if $\mathbb{1}_{\text{permit}} = 1$).
$u_{i,t}(h) = \theta_{i,h,t} + \eta_j - \kappa \cdot \text{dist}(h)$	Flow utility (net of rent when used inside value functions).
$D(r)$	(Residual) demand schedule faced by a unit/landlord (used in proofs/simulations).
MC	Marginal cost proxy (maintenance/opportunity cost) used in markup heuristics.
Equilibrium objects	
$V_{i,t}(h)$	Tenant value of staying in h vs. moving (Eq. (1)).
$V_\ell(s_t)$	Landlord ℓ's value function (dynamic program).
$\sigma_\ell(s_t)$	Stationary Markov strategy/policy of landlord ℓ.
MPE	Markov Perfect Equilibrium of the dynamic game.
$f(s_t, \sigma, \mu)$	Law of motion for states under policies and institutions.
$\mathbb{1}_{\text{permit}}$	Indicator: renovation requires a permit.
$\mathbb{E}[\cdot]$	Expectations operator (w.r.t. shocks and transition).
Welfare and decomposition	
W_t^{tenant}, W_t^{landlord}	Tenant and landlord welfare components at t.
$V^{\text{planner}}(s_t)$	Planner's value function.

Continued on next page

Symbol	Meaning / Units
$\Delta W = V^{\text{planner}} - V^{\text{MPE}}$	Welfare gap (planner vs. decentralised equilibrium).
$\{\Delta W_k\}$	Decomposition components (investment, allocation, mobility, markups, spillovers).
Calibration / SMM	
Θ	Parameter vector (preferences, technology, frictions, policy).
$\widehat{m}$	Empirical target moments stacked to match the model.
$m(\Theta)$	Model-implied moments from stationary MPE at Θ.
W	SMM weighting matrix (typically diagonal; second pass optional).
$\hat{\Theta}$	SMM estimate: $\operatorname{argmin}_{\Theta}(\widehat{m} - m(\Theta))' W (\widehat{m} - m(\Theta))$.
R, burn-in	Simulation replications and burn-in periods for moments.
Spillovers / proofs (Appendix)	
ϕ_{reno}	Renovation cost used in renoviction condition (Appendix D).
$P_{\text{exit}}(q)$	Probability incumbent exits as a function of quality (Appendix D).
λ	Lagrange multiplier for effort budget in spillover lemma (Appendix C).
α^j, β^j	Reduced-form return-to-investment coefficients (Appendix C).
Discretisation grids (no overload with $\mathcal{R}_j$)	
$\mathcal{G}_j^r$	Finite grid for rents used in DP/VI in jurisdiction j.
$\mathcal{G}^m$	Finite grid for maintenance effort.
$\mathcal{G}^q$	Finite grid for quality levels.

Additional shorthands used in calibration and diagnostics

Symbol	Meaning
RD	Rent dispersion: new-lease mean minus incumbent mean (pp).
v	Vacancy rate (%).
τ	Turnover rate (%/yr): $\sum_t \text{churn}_t \big/ \sum_t \text{occ}_{t-1}$.
s_m	Share with positive maintenance: $\mathbb{E}[\, m_{h,t} \,]$.
g_{AB}	Jurisdictional gap (A vs. B) on new leases (pp).
Θ	Parameter vector (here $\Theta = (\theta, d, c_m)$ in the K-block).
$G(\Theta)$	Jacobian of simulated moments w.r.t. parameters G.
W	SMM weighting matrix (typically diagonal).
$\text{rank}(\cdot)$	Numerical rank (SVD-based).
$\text{cond}(\cdot)$	Condition number (ratio $\sigma_{\max}/\sigma_{\min}$).
ΔQ	Change in the objective when a moment is dropped (LOO).
$\lVert \cdot \rVert_\infty$	Max norm (used for LOO parameter deviations).
$\text{status}_{h,t}$	Lease status in $\{\text{new}, \text{inc}, \varnothing\}$ (used in K/G).