Information Sharing vs. Collusive Coordination: Algorithmic Pricing in US Multifamily Rentals*

Sophie Calder-Wang† Mitsuru Igami‡ Gi Heung Kim§ Takuo Sugaya¶

November 6, 2026

Abstract

Does pricing software simply inform firms about current market conditions, or does it also facilitate collusive conduct? This is the question at the center of the US Department of Justice's suit against RealPage, which the existing conduct tests cannot answer because they assume complete information and leave no room for demand uncertainty. We propose a framework to study such situations by building on the Bayes correlated equilibrium of Bergemann and Morris (2016), and empirically implement it to the data and demand estimates of Calder-Wang and Kim (2026). For each product, we identify and estimate the set of marginal costs consistent with the observed prices and the demand system under any information structure. Our preliminary estimates suggest that information sharing alone can rationalize the software-driven rent increase under the least restrictive informational assumption. However, once we impose an additional assumption that the affected neighborhoods achieve complete information, our conduct test indicates potential departure from unilateral profit maximization.

Keywords: algorithmic pricing; algorithmic collusion; antitrust; Bayes correlated equilibrium; information sharing; multifamily housing; set identification

JEL Classification: L13, L41, C57, D83

*We thank Stephen Morris and Paulo Somaini for helpful discussions, Jean-François Houde and Samuel Goldberg for discussing earlier versions of the paper, and Gaston Illanes, Lorenzo Magnolfi, and Daniel Ackerberg for their comments. We also thank the participants at the seminars and conferences where earlier versions were presented, including the 2026 International Industrial Organization Conference (IIOC), the 2026 MIO conference at Yale, and the 2026 Banff IO conference.

†The Wharton School, University of Pennsylvania, and the National Bureau of Economic Research (NBER).

‡Department of Economics, University of Toronto.

§Department of Economics, Boston College.

¶Graduate School of Business, Stanford University.

1 Introduction

Information sharing among competing landlords has become the subject of antitrust policy. In August 2024, the US Department of Justice (DOJ) and eight states sued RealPage, whose revenue-management software collects subscribers' nonpublic lease data and recommends a rent for each apartment (USDOJ, 2024).[1] The use of such software has become increasingly common: by 2019, the property-management companies that used one operated 25% of the apartment buildings and 34% of the units in a commercial database that covers about half of the country's market-rate apartments (Calder-Wang and Kim, 2026). In November 2025, the DOJ and RealPage proposed a settlement that restricts the data that the software may use (USDOJ, 2025b).

Do these algorithms simply aggregate market data and reduce demand uncertainty—with a potentially efficiency-enhancing "information effect"—or do they also facilitate collusion by recommending rents at supracompetitive levels? The answer matters to antitrust policy, the regulation of algorithms, and the empirical analysis of market competition. For antitrust, it sets the boundary between lawful market intelligence and an illegal restraint of trade. Whether a remedy should restrict the data that the software uses, or the conduct of its users, or both is an open question. For the empirical analysis of competition more broadly, the issue exposes a gap in the standard toolkit. The tests of competitive conduct in differentiated-product markets (Bresnahan, 1982; Berry and Haile, 2014; Backus et al., 2021) assume firms know the demand when they set prices. A test that assumes complete information leaves information sharing no role in pricing and cannot separately identify the information effect from collusion.

This paper proposes a test of competitive conduct that is robust to what firms know, and applies it to US apartments before and after the adoption of algorithmic pricing (the "event"). The idea is simple. If the software purely serves market-intelligence purposes, the observed rents are a Bayes Nash equilibrium (BNE) under *some* information structure, both before and after the event, and the firms' marginal costs (MCs) implied by the rents should stay unchanged. If, however, no information structure can rationalize the rents after the event without a rise in the implied MC, the rents are too high for unilateral profit maximization. The set of outcomes that unilateral (i.e., competitive) conduct can produce under some information structure is the set of Bayes correlated equilibria (BCE) of Bergemann and Morris

[1] Tenants' class actions against RealPage and its subscribers, filed since 2022, are consolidated as *In re: RealPage, Inc., Rental Software Antitrust Litigation (No. II)*, MDL No. 3071 (M.D. Tenn.). The DOJ amended its complaint in January 2025 to add several large landlords as defendants (USDOJ, 2025a).

(2016). Their main result is that every BCE is a BNE under some information structure, and that every BNE is a BCE, which allows us to identify the set of MCs consistent with the data without specifying any particular information structure. Thus, we first determine the set of MC parameter values that rationalize the observed rents under some admissible information structure. We then ask whether the widespread adoption of algorithmic pricing shifts the set of implied MCs upward in the focal (or "treated") neighborhoods, relative to their matched-control neighborhoods that did not experience the surge of software. As such, the main statistic of interest is the treated-control difference in the shift of the implied MC, which we call the "gap" for simplicity. The null hypothesis has two parts: (i) firms price unilaterally under some admissible information structure, and (ii) the treatment did not change any firm's true MC relative to its controls'. The gap is zero under the null because whatever the software did to a firm's information is already inside the set of information structures that we permit in the estimation of the implied MC. Conversely, a positive gap means that the post-treatment rents in the treated neighborhoods are too high to be optimal for a unilateral profit-maximizer under any information structure.

Three theoretical insights underpin our BCE approach. First, under incomplete information, a firm's first-order condition (FOC) does not hold exactly in every period but instead only on average (i.e., in expectation) across periods. Hence, the FOC equates the rent with the sum of MC and a "pooled" markup. This markup consists of the ratio of the firm's interim expectations of its market share and the slope of its demand (section 3). Second, the joint distribution of those interim expectations that an admissible information structure induces is a mean-preserving contraction (MPC) of the distribution of the market state, one that respects what the firm knows. Conversely, every such contraction is induced by some admissible information structure (Proposition 1 in section 4). Hence, the search over information structures is a search over a convex set of distributions without specifying any signal. Third, the moment conditions that identify the MC parameters are linear in those distributions. Hence, the generalized method of moments (GMM) estimates of a product's smallest and largest implied-MC shift are the solutions of two linear programs (section 5). We call the interval between them the product's range of GMM estimates, which contains the estimate under complete information. We aggregate the ranges across products, firms, and matched pairs of neighborhoods into the range of the gap and report the confidence interval (CI) of Imbens and Manski (2004) at its end points.

The US multifamily rental industry is a natural setting for our test, for three reasons. First, the stakes are high. Rent is the largest item in most renters' budgets, and the DOJ's

case against RealPage highlights the importance of distinguishing between information and coordination. Second, the treatment is well documented. Calder-Wang and Kim (2026) hand-collected data on the adoption of software by the property-management companies. Third, the setting cleanly maps to the model. A building is a firm, a room type is its product, and a census tract is a neighborhood in which nearby buildings compete. The software reached different neighborhoods in different years, which gives treated neighborhoods sufficiently rich pools of comparable controls. The demand model and its estimates by Calder-Wang and Kim (2026) are already available. We use a slightly extended version of their data set, which covers 34,217 buildings in 13,240 census tracts of 50 metropolitan areas ("metros") from 2005 to 2019.

Our empirical analysis starts with the measurement of the effect of algorithmic pricing on rents, where we use a propensity-score-matched difference-in-differences (PSM-DID) design at the level of the neighborhood in all 50 metros. A neighborhood is treated in the first year in which the buildings using algorithmic pricing reach 20% of its incumbent room capacity, starting from zero in all previous years. Its controls are never-treated neighborhoods within the same metro, matched on a propensity score (i.e., the predicted probability of being treated based on a logistic regression of the indicator of treatment on the neighborhood's characteristics). We then proceed to the implementation of our BCE approach to these matched pairs of neighborhoods (and the buildings and products therein). Currently, the demand estimates are available only for the 11 metros of Calder-Wang and Kim (2026).[2]

Three findings emerge. First, algorithmic pricing raises rents. In the 395 treated census tracts that experienced the introduction of pricing software during the stable market environment,[3] the rents of incumbent buildings rise by 1.53% with standard error (SE) of 0.53 relative to matched controls. The effect grows from +1.21% (SE 0.49) in the event year, measured between the treatment year and the previous year, to +2.34% (SE 0.77) four years later. Sixteen alternative specifications of the PSM-DID procedure generate estimates between +1.13% and +2.14%, with 14 of them significantly larger than zero at the 5% level. Second, the implied MC rises by more in the treated neighborhoods than in their controls. On the 56 matched pairs (1,701 products) for which demand estimates are available, the gap is +\$38.9 ($p = 0.047$) under complete information and +\$41.0 ($p = 0.050$) at the information structures that fit each product's moment conditions best. Third, whether some other

[2]We are planning to extend the demand-estimation coverage to all 50 metros.

[3]We focus on the middle tercile of metro-wide rent growth (averaged across never-adopter buildings' products), for two reasons. First, our BCE framework assumes stationarity. Second, the rent effect is most precisely estimated in this relatively flat phase of housing business cycles compared with either boom or bust phase.

admissible information structure can rationalize the gap depends on how much uncertainty the firms are assumed to face. Under the least restrictive informational assumption, in which a firm knows only its own rent (up to its quartiles in the baseline specification) and whether the treatment has happened, the range of the gap estimates is $[-\$2.8, +\$81.5]$. The null of unilateral conduct cannot be rejected ($p = 0.83$). If we further assume that the treated firms obtain complete information after the treatment, the range is $[+\$13.3, +\$61.9]$, even though the Imbens–Manski CI at its end points, $[-\$3.1, +\$75.9]$, includes zero. With five rent bins instead of four, the null is rejected at 5% under the latter, more restrictive informational assumption.

These BCE estimates are preliminary and subject to change in the future versions of the paper when we expand the data set, enrich the model, and improve the econometric methods. That said, the results are promising in the sense that, even within the current sample with limited size, our conduct test exhibits statistical power under certain configurations.

We have organized the rest of the paper as follows. Section 2 briefly summarizes the policy context. We explain the model of price competition under uncertainty in section 3, the identification of the MC parameters in section 4, and the estimator and the inference in section 5. Section 6 describes the data. Section 7 introduces the PSM-DID design and reports the reduced-form analysis of rents. We explain the demand model and estimate it in section 8. Section 9 empirically implements our BCE framework and reports the gap estimates under a variety of configurations.

Related Literature. This work relates to four strands of literature. The first studies pricing algorithms and collusion from mostly theoretical perspectives. Learning algorithms can sustain supracompetitive prices without communication (Calvano et al., 2020; Klein, 2021; Asker et al., 2022; Johnson et al., 2023) and a third party that sells the same pricing software to competing firms raises antitrust concerns of its own (Harrington, 2018, 2022, 2025a). Whether better information raises prices is an older question. Sharing demand information among oligopolists can raise or lower prices, depending on the nature of the uncertainty and of the competition (Clarke, 1983; Vives, 1984; Gal-Or, 1985; Raith, 1996; Kühn and Vives, 1995), and better demand forecasts can make collusion either easier or harder to sustain (Miklós-Thal and Tucker, 2019). In the recent theory of third-party pricing algorithms, the effect of information sharing on average prices depends on the third party's objective. A third party that extracts its fee from each seller's individual gain makes prices more responsive to demand without raising their average, whereas one that maximizes the

sellers' joint profit raises it (Aleksenko and Miklós-Thal, 2026).[4]

Second, the empirical evidence is scarce. Assad et al. (2024) find that prices rose in German gasoline duopolies after both stations adopted algorithmic pricing. Brown and MacKay (2023) document the pricing technologies of online retailers and their consequences for competition. Clark and Houde (2026) show that the members of a gasoline cartel stopped sharing their sales data through a commercial platform when the cartel collapsed. Calder-Wang and Kim (2026) find that adopters' rents respond more to demand conditions and that a conduct test favors joint-profit maximization among the users of the same software. We contribute a framework and test that formally treat the firms' information as a latent object.

The third strand is conduct testing, which starts from the FOC of an oligopoly pricing model (Bresnahan, 1982; Berry and Haile, 2014). Recent tests compare models of conduct by the cost-side exclusion restrictions that each model implies, under the maintained assumption of complete information (Backus et al., 2021; Magnolfi and Sullivan, 2022; Duarte et al., 2024). By contrast, we maintain competitive conduct as a null hypothesis and relax information structure. These two classes of tests operate in different dimensions and complement each other.

The fourth strand is the econometrics of games with weak assumptions on information, built on the BCE of Bergemann and Morris (2013, 2016). Magnolfi and Roncoroni (2023) set-identify the payoffs of static entry games under any information structure. Syrgkanis et al. (2021) do so for auctions, and Magnolfi and Lomys (2024) for games played by no-regret learning algorithms. By contrast, we work with the canonical oligopoly model with differentiated products and price-setting game. This setting differs in that the action is a continuous price and the equilibrium condition is an equality (the FOC), which allows us to characterize the admissible information structures through the MPCs of the state and to compute the range of estimates by linear programs. Our treatment of information is similar to Pakes (2010) and Pakes et al. (2015), whose best-response condition holds at the agent's information without restricting its content. Similarly to their optimality conditions, we assume rational expectations in the sense that the FOC holds on average. The main difference is that, by characterizing every information structure via BCE, we turn their necessary condition into an equivalence, which means any point in our identified set is guaranteed to be an equilibrium. For inference on the set-identified gap, we draw on Imbens

[4]Harrington (2025b) argues on the same ground that shared data are neither necessary nor sufficient for the harm, which the remedies of 2025 nonetheless target. Our test takes the data flows as given and asks what rents they can rationalize under any information structure.

and Manski (2004), Stoye (2009), and Canay et al. (2023).

2 Industry and Policy Context

This section describes the pricing software, the litigation it has provoked, and the two interpretations of its effect.

Multifamily housing (buildings with five or more units) holds about 20 million of the 44 million occupied rental units in the United States (Calder-Wang and Kim, 2026). Owners, institutional investors in particular, commonly outsource the operation of their buildings to a management company, which sets the rents, manages the leases, and maintains the property. Since the early 2000s, a few software firms have been selling revenue-management software and services that recommend a rent for each unit. Two programs dominated the industry during our sample period. RealPage acquired YieldStar in 2002. The Rainmaker Group developed LRO (Lease Rent Options) around 2005 and sold it to RealPage in December 2017.[5]

RealPage's system works as follows (USDOJ, 2024). A subscribing company inputs its lease data, such as effective rents, lease terms, and expected move-outs, into the system. The software combines these data with the data from other subscribers in the same area and with its own forecasts of local demand, and recommends a rent for each unit.[6] RealPage described YieldStar as using "statistically derived supply/demand forecasts and price elasticity models" to compute optimal prices. Its dashboard offers the property manager two buttons: to accept the recommended rents or to override them (Calder-Wang and Kim, 2026, section 2.1). The complaint alleges that RealPage encouraged the users' compliance with its recommendations through automatic acceptance settings, the monitoring of override rates, and the procedural rule that obligates the users who override its recommendations to explain their reasons to senior managers.

The litigation began with tenants' class actions against RealPage and its subscribers in 2022, alleging a price-fixing conspiracy under Section 1 of the Sherman Act.[7] On August 23, 2024, the DOJ and eight state attorneys general sued RealPage (USDOJ, 2024). The

[5] The two programs were kept separate after the acquisition, and Calder-Wang and Kim (2026) find no effect of the acquisition on rents or occupancy. We treat the users of either program alike.

[6] The complaint describes YieldStar and its successor. LRO, by the complaint's account, relied mainly on competitor data that its users entered from public sources.

[7] *In re: RealPage, Inc., Rental Software Antitrust Litigation (No. II)*, MDL No. 3071 (M.D. Tenn.). Related suits name Yardi Systems, a smaller software firm, and its subscribers (*In re Yardi Revenue Management Antitrust Litigation*, No. 2:23-cv-01391 (W.D. Wash.)).

complaint alleges that RealPage and its subscribers agreed to share competitively sensitive information and to use the software to align rents, in violation of Section 1, and that RealPage monopolized the market for revenue-management software, in violation of Section 2. An amended complaint of January 2025 added several large landlords as defendants (USDOJ, 2025a). In November 2025, the DOJ and RealPage filed a proposed final judgment (USDOJ, 2025b). Its remedy is behavioral and has three main provisions. RealPage (i) may not use competitors' nonpublic data in its pricing recommendations, (ii) must remove or redesign the features that limit rent decreases, and (iii) may not host pricing discussions at its user-group meetings. Several states have gone further and restricted the use of such software for setting rents by statute.[8]

The same facts admit two economic interpretations, which the complaint does not clearly separate. The first interpretation is purely informational. Software allows a landlord to observe its competitors' current leases and the current state of demand. Consequently, its rent becomes closer to the one it would set under complete information. The recommendations are correlated across landlords because their inputs are shared, but each landlord's rent remains its own best response and serves unilateral profit maximization. In this scenario, the rents are a BNE under richer information than before, and their average may rise or fall with the demand conditions of the year. Under the second interpretation, the third-party software facilitates coordination. The recommendations maximize the subscribers' joint profit rather than each subscriber's individual profit. The compliance-inducing features of the software and organizational procedures mitigate the temptation to undercut and help sustain tacit collusion. The remedy of 2025 restricts both the data that the first channel uses and the features that serve the second channel.[9] Our analysis asks how far the first channel alone can go to rationalize the observed rents. We let each landlord know (or not know) anything as long as it is consistent with the data. We then test whether the rents after the spread of algorithmic pricing are too high under the null hypothesis of unilateral profit maximization.

[8]New York's law of October 2025 imposes liability for facilitating an agreement among rental owners or managers to set prices through software or algorithmic devices that perform a coordinating function, whether or not a horizontal agreement is proven separately (New York State Senate, 2025).

[9]Harrington (2025b) argues that shared data are neither necessary nor sufficient for the harm, because a third party can produce supracompetitive recommendations from its subscribers' own data. Our test does not depend on which data the software used.

3 Model

This section lays out the pricing model underpinning our BCE framework. Section 3.1 describes the firm's problem, section 3.2 its information and the equilibrium, and section 3.3 what the model implies for the rent, the markup, and the implied marginal cost.

3.1 Firms, Products, and Costs

A firm owns a building and rents its rooms. We index products (buildings × room types) by j and years by t. In year t the firm sets the rent p_{jt} of each of its products. The demand model determines the product's market share s_{jt} as a function of the rents and characteristics of all products on the market and the renters' preferences. We use $B_{jt} = \partial s_{jt}/\partial p_{jt} < 0$ to denote the price derivative evaluated at the actual rents. We focus on the pricing of each room type and abstract from the intra-building substitution across different room types.[10]

The marginal cost (MC) of renting out a room of type j in year t is

$$c_{jt} = \theta_j + \delta_j \, d_{jt} + \omega_{jt}, \tag{1}$$

where $d_{jt} = \mathbf{1}[t \geq \tau_j]$ indicates the years since the first substantial use of algorithmic pricing either by the owner of j or its competitors in the same local neighborhood (the "event"), with τ_j as the event year in that market. θ_j is the cost level before the event, and δ_j is the potential shift at the event. The firm knows all components of its own cost, including the shock ω_{jt}. The firm's profit from product j is $\pi_{jt} = (p_{jt} - c_{jt})\, s_{jt}$.

3.2 Information and Equilibrium

Within a year, the game proceeds in three steps. First, the year's fundamentals are realized. They are the demand state, which comprises the preferences and the demographic mix of the renters, and the cost shocks of every product. Second, each firm observes a signal about the fundamentals and about its rivals, which we call its information, $\mathcal{I}_{jt}$. Third, every firm simultaneously sets its rents, after which renters choose rooms. The information structure consists of the joint distribution of the signals and the fundamentals. It may differ across firms and may change at the event, which is how we model the information-sharing role of algorithmic pricing. We impose three restrictions on what a firm knows when it sets a rent.

[10] We provide the empirical support for this approximation in section 8.2.

Assumption 1 (Information)**.** For every product j and year t, (i) the rent p_{jt}, (ii) the cost shock ω_{jt}, and (iii) the regime indicator d_{jt} are measurable with respect to $\mathcal{I}_{jt}$.

Part (i) says that a firm knows its own rent. It is without loss because a rent that is a function of a signal is itself part of the firm's information. Part (ii) says that the firm knows its own cost shock and hence the exact MC of its product, c_{jt}. Without it, the first-order condition below cannot be solved for the cost.[11] Part (iii) says that the firm knows whether the event has happened. We call an information structure *admissible* when it satisfies all three. Everything else is unrestricted, including what the firm knows of the demand state and of its rivals' rents.

We study a Bayes Nash equilibrium (BNE) of this pricing game, stated in outcome form. It is a joint distribution over the fundamentals, the firms' information, and the rents with three properties. First, its marginal distribution over the fundamentals is the true one. Second, each rent is measurable with respect to the firm's information. Third, for every realization of its information, each firm's rent maximizes its expected profit given that information, with its rivals' rents set by the same rule. The first property is a consistency restriction; with a second one, it means that firms share a correctly specified prior over the fundamentals and the signal-generating process. Each firm forms its beliefs by conditioning this distribution on its own information. Thus, expectations are taken under the true conditional distribution, although firms may have different information. The expectation in (2) below is taken under the true joint distribution of the fundamentals and the signal.[12] The first-order condition (FOC) of the third property (expected-profit maximization) for product j is

$$\mathbb{E}\big[\, s_{jt} + B_{jt}\,(p_{jt} - c_{jt}) \,\big|\, \mathcal{I}_{jt}\big] = 0. \tag{2}$$

At the rent set by the firm, the expected gain from a marginal rent increase on its occupied apartments equals the expected loss from the apartments it vacates. As is standard in FOC approaches, we do not impose second-order conditions separately.

[11] Without part (ii), the cost shock cannot be taken outside the conditional expectation in equation (2) below; hence, the cost inversion and the resulting moment conditions we use will become unavailable. An alternative pure-BCE approach allows firms not to observe their own cost shocks and combines the first-order conditions with instrument orthogonality and an upper bound on cost-shock second moments to obtain bounds using a Cauchy–Schwarz argument. The implementation in this paper maintains own-cost observability.

[12] A common prior means that the firms' beliefs derive from the same prior; agreement alone does not make the prior correct. The argument here needs each firm's beliefs to coincide with the true conditional distribution given its information, which is the rational-expectations restriction. If all firms condition the same true joint distribution on their respective information, their beliefs also have a common-prior representation, which makes the two readings compatible. Section 4.3 compares this restriction with the best-response condition of Pakes (2010) and Pakes et al. (2015).

The set of outcome distributions that satisfy these conditions under some admissible information structure is the set of Bayes correlated equilibria (BCE) of Bergemann and Morris (2016) with Assumption 1 as an additional restriction.[13] Characterizing every information structure, rather than a parametric family of signals, is what allows us to describe the set of MCs consistent with the rents under any of them (Bergemann and Morris, 2013) in section 4. It is the sense in which our BCE framework maintains own-profit maximization under any information structure consistent with the data.

3.3 Rents, Markups, and the Implied Marginal Cost

Under complete information, the firm observes all of the fundamentals, and the expectation in (2) is replaced by their realizations. Hence, the FOC holds exactly year by year:

$$p_{jt} - c_{jt} = -\frac{s_{jt}}{B_{jt}} =: \mu_{jt}. \tag{3}$$

The markup μ_{jt} is a function of the demand estimates alone, allowing us to identify the implied MC for each j and t as $c_{jt} = p_{jt} - \mu_{jt}$.

Under incomplete information (with an admissible information structure), Assumption 1 makes c_{jt} and p_{jt} measurable with respect to $\mathcal{I}_{jt}$, and (2) rearranges to

$$p_{jt} - c_{jt} = -\frac{\bar{s}_{jt}}{\bar{B}_{jt}}, \tag{4}$$

where an overbar denotes what we call an "interim mean," that is, the firm's expectation given its information: $\bar{s}_{jt} := \mathbb{E}[s_{jt} \mid \mathcal{I}_{jt}]$ and $\bar{B}_{jt} := \mathbb{E}[B_{jt} \mid \mathcal{I}_{jt}]$. We call $-\bar{s}_{jt}/\bar{B}_{jt}$ the *pooled markup*, the ratio of the interim mean of the market share to the interim mean of the demand slope, which pools the states that the firm cannot tell apart given its information. It differs from the *expected markup*, $\mathbb{E}[\mu_{jt} \mid \mathcal{I}_{jt}]$, the interim mean of the complete-information markup μ_{jt} of (3), because the ratio of two interim means is not the interim mean of the ratio. The implied MC is $\bar{c}_{jt} := p_{jt} + \bar{s}_{jt}/\bar{B}_{jt}$.

Three aspects of these objects are worth highlighting before we proceed to identification (section 4) and estimation (section 5). First, a firm that knows less does not misprice on average. Under correctly specified beliefs, the law of iterated expectations implies that the interim means are correct on average. Hence, its rent is the cost plus a markup that is correct on average. Second, the pooled markup can be either higher or lower than the expected

[13] Equivalently, a mediator who observes the fundamentals recommends a rent to each firm; an outcome distribution is an equilibrium when every firm follows its recommendation.

markup.[14] Third, a departure from own-profit maximization, such as price coordination with rivals, adds the rivals' terms to (2) and raises the rent above what any admissible information structure can rationalize at the same MC. Under the maintained model, such a departure appears as a rise in the implied MC. Our empirical analysis asks whether the implied MC (i.e., the residual of the unilateral pricing model) changes differently at the event in treated neighborhoods than in their matched control neighborhoods.

4 Identification

This section explains what the data identify under the model of section 3. Section 4.1 states what is observed and the benchmark case under complete information, and section 4.2 derives the moment conditions when the firm's information is latent. Section 4.3 characterizes the admissible information structures, and section 4.4 the identified set of the cost parameters.

4.1 Observables and the Complete-Information Benchmark

For each product j and year t we observe the rent p_{jt}, the regime indicator d_{jt}, the share s_{jt} and its own-rent derivative B_{jt} from the demand estimates, and a vector of K instruments z_{jt}. We call $x_{jt} := (p_{jt}, s_{jt}, B_{jt}, z_{jt}, d_{jt})$ the state of the year. We focus on one product at a time, and regard the product's years within a regime as draws from one distribution. The instruments satisfy the exclusion restriction $\mathbb{E}[z_{jt}\,\omega_{jt}] = 0$; that is, they are uncorrelated with the cost shock ω_{jt}.[15]

Under complete information, (3) and (1) give the cost shock as a function of the state and the two cost parameters, $\omega_{jt} = p_{jt} + s_{jt}/B_{jt} - \theta_j - \delta_j d_{jt}$. The exclusion restriction becomes K linear equations in (θ_j, δ_j):

$$\mathbb{E}\big[\, z_{jt}\,(p_{jt} + s_{jt}/B_{jt})\big] - \theta_j\,\mathbb{E}[z_{jt}] - \delta_j\,\mathbb{E}[z_{jt}\,d_{jt}] = 0. \tag{6}$$

The two parameters are point identified when $\mathbb{E}[z_{jt}]$ and $\mathbb{E}[z_{jt}d_{jt}]$ are linearly independent,

[14] Using $D_{jt} := -B_{jt}$ to denote the absolute demand slope and $\mu_{jt} = s_{jt}/D_{jt}$ the complete-information markup of (3), the pooled markup minus the expected markup is

$$\frac{\bar{s}_{jt}}{\bar{D}_{jt}} - \mathbb{E}[\mu_{jt} \mid \mathcal{I}_{jt}] = \frac{\mathrm{Cov}(D_{jt}, \mu_{jt} \mid \mathcal{I}_{jt})}{\bar{D}_{jt}}, \tag{5}$$

which can take either sign.

[15] The instruments need not be known to the firm, because the restriction binds the joint distribution of the state rather than the firm's information. Section 4.2 shows that such instruments enter the moment conditions through their interim means, by the law of iterated expectations, as in equation (8).

which we call the rank condition. It requires at least two instruments, and their means should differ between the regimes; the constant and the regime indicator are such a pair whenever the product is observed in both regimes. The parameters are overidentified when $K > 2$. The level θ_j is identified by the instrument-weighted mean of the implied cost $p_{jt} + s_{jt}/B_{jt}$ before the event, and the shift δ_j by the difference between the two regimes' instrument-weighted means of the implied cost.

4.2 Moment Conditions under Latent Information

Now consider the incomplete-information case. Under an admissible information structure, (1) implies

$$\omega_{jt} = \bar{c}_{jt} - \theta_j - \delta_j \, d_{jt}, \tag{7}$$

where $\bar{c}_{jt} = (\overline{Bp}_{jt} + \bar{s}_{jt})/\bar{B}_{jt}$ is the implied MC of (4) and $\overline{Bp}_{jt} := \mathbb{E}[B_{jt}p_{jt} \mid \mathcal{I}_{jt}]$. The product $B_{jt}p_{jt}$ enters as its own interim mean. Let ν_j denote the joint distribution, induced by the information structure, of the vector of interim means $(p_{jt}, \bar{s}_{jt}, \bar{B}_{jt}, \overline{Bp}_{jt}, \bar{z}_{jt}, d_{jt})$, with $\bar{z}_{jt} := \mathbb{E}[z_{jt} \mid \mathcal{I}_{jt}]$. Because ω_{jt} and d_{jt} are known to the firm, the exclusion restriction can be written in terms of interim means,

$$\mathbb{E}[\, z_{jt} \, \omega_{jt} \,] = \mathbb{E}\big[\, \bar{z}_{jt} \, (\bar{c}_{jt} - \theta_j - \delta_j \, d_{jt}) \big] =: \eta_j(\theta_j, \delta_j; \nu_j), \tag{8}$$

where the first equality is the law of iterated expectations; we call the K-vector η_j the moment vector. The distribution ν_j is unobserved because the information structure is. Different information structures induce different interim means from the same distribution of the state. Hence, the identification question is which pairs of an information structure and cost parameters set (8) equal to zero.

4.3 Admissible Information Structures

The moment vector (8) depends on the distribution of the interim means rather than on the signal itself. Hence, we characterize the distributions that admissible information structures can induce. Even though the firm's information is unobserved, its interim means must be matched. That is, conditional on any vector of interim means, the actual state must average to that vector. The restriction follows from the law of iterated expectations, as the interim means are a function of the firm's information. The expectation of the state conditional on the interim means equals the interim means themselves. This restriction allows us to

characterize the feasible distributions of the interim means without specifying the signals themselves.

A distribution with this property is a mean-preserving contraction (MPC) of the distribution of the state. By Assumption 1, the price and regime entries of the vector of interim means are the state's own price and regime indicator, because the firm knows both; we call a contraction with this property price- and regime-respecting. Conversely, every price- and regime-respecting contraction is induced by some admissible information structure, because the interim means themselves can serve as the signal.

Proposition 1 (Cost Parameters Consistent with the Model)**.** *Fix a product j and cost parameters (θ_j, δ_j), and suppose that $\bar{B}_{jt}$ is bounded away from zero on the support of the interim means. The following are equivalent.*

(a) There exists a joint probability model for the observed state, the cost shock ω_{jt}, and the firm's information $\mathcal{I}_{jt}$, with the given marginal distribution of the observed state, such that Assumption 1 holds, the first-order condition (2) *holds almost surely, and the exclusion restriction $\mathbb{E}[z_{jt}\,\omega_{jt}] = 0$ holds.*

(b) There exists a price- and regime-respecting mean-preserving contraction (MPC) ν_j of the distribution of the state such that $\eta_j(\theta_j, \delta_j; \nu_j) = 0$.

The proof is in Appendix A. The identified set of the cost parameters is the set for which (b) holds:

$$\Theta_j := \big\{(\theta_j, \delta_j) : \eta_j(\theta_j, \delta_j; \nu_j) = 0 \text{ for some admissible } \nu_j\big\}. \tag{9}$$

It is the projection onto the cost parameters of the joint set of triples $(\theta_j, \delta_j, \nu_j)$ that set (8) equal to zero. The other projection, onto the information structures, is the set of information structures under which some cost parameters set (8) equal to zero.

Proposition 1 is an equivalence rather than a necessary condition, which matters in three ways. First, every point of Θ_j is rationalized by a BNE under some admissible information structure; hence, a point outside it is rejected by the model. Second, particular information structures, complete information among them, can be tested as members of the joint set. Third, instruments that the firm need not observe enter (8) through their interim means. All three are what the characterization of every information structure adds to the FOC approach of Pakes (2010) and Pakes et al. (2015).[16]

[16] Their best-response condition is (2) with the information set unrestricted. Their identified set is not sharp in general, because their inequalities are necessary conditions for a best response at the agent's information,

4.4 The Identified Set of the Cost Parameters

The affine form of the moment conditions makes Θ_j tractable. The moment is affine in the cost parameters, with coefficients that do not depend on the information structure. An MPC preserves means, $\mathbb{E}[\bar{z}_{jt}] = \mathbb{E}[z_{jt}] =: b_j$. Because the regime indicator is known to the firm, $\mathbb{E}[\bar{z}_{jt}\, d_{jt}] = \mathbb{E}[z_{jt}\, d_{jt}] =: c_j$. Substituting both into (8) gives

$$\eta_j(\theta_j, \delta_j; \nu_j) = a_j(\nu_j) - \theta_j\, b_j - \delta_j\, c_j, \tag{10}$$

where $a_j(\nu_j) := \mathbb{E}_{\nu_j}[\,\bar{z}_{jt}\,\bar{c}_{jt}\,]$, which is the instrument-weighted mean of the implied cost at the signal. This $a_j(\nu_j)$ is the only object that changes with the information structure. Under the rank condition in section 4.1 (i.e., that b_j and c_j are linearly independent), this affine form yields a unique, closed-form solution by the generalized method of moments (GMM) for each candidate information structure,

$$\big(\theta_j(\nu_j),\, \delta_j(\nu_j)\big)^\top = (M_j^\top W M_j)^{-1} M_j^\top W\, a_j(\nu_j), \tag{11}$$

where $M_j := [\,b_j \;\; c_j\,]$, for any positive definite weighting matrix W (the solution does not depend on W when $\eta_j(\theta_j, \delta_j; \nu_j) = 0$). The moment vector at that solution is the residual. Let $r_j(\nu_j)$ denote the Euclidean norm of the residual, which is zero if and only if every moment condition holds exactly at the solution. The identified set is the image of the admissible information structures under (11), restricted to those with a zero residual:

$$\Theta_j = \big\{\big(\theta_j(\nu_j), \delta_j(\nu_j)\big) : \nu_j \text{ admissible},\; r_j(\nu_j) = 0\big\}. \tag{12}$$

It is set-valued because the information structure is latent. Complete information is the admissible information structure whose interim means equal the actual state. In that case, (11) becomes identical to the complete-information benchmark in section 4.1. In the exactly identified case, $K = 2$, every admissible information structure has a zero residual; hence, the set is the image of all of them.

not sufficient ones. A parameter that satisfies them need not be rationalized by any information structure and equilibrium. The similarities and differences between theirs and ours are as follows. Their expectational error is mean zero given the agent's information, as our correctly specified beliefs make ours. Their structural error, known to the agent and not to the econometrician, is our cost shock. Their instruments must lie in the agent's information set, whereas ours enter through their interim means. Where their inequalities set-identify payoffs in discrete choices, our FOCs are equalities in a continuous rent, with the identified set arising from the latent information alone.

5 Estimation and Inference

This section explains the estimator and the inference. Section 5.1 constructs the information structures in the sample, section 5.2 the moment conditions and the range of GMM estimates of a unit's cost shift, section 5.3 the aggregation across units and markets, and section 5.4 the inference under set identification. The empirical specification and the operational details are deferred until section 9.2.

5.1 Information Structures in the Sample

An estimation unit is a product j observed for T_j years, with state $x_{jt} = (p_{jt}, s_{jt}, B_{jt}, z_{jt}, d_{jt})$. The sample distribution of the state puts mass $1/T_j$ on each year. Since the firm may not have complete information, its FOC (2) holds only *on average* across multiple years rather than exactly in each year. In our BCE framework, an information structure is a rule that can pool multiple years of a unit into a single "signal," with weights assigned to those years, which determine how the weighted-average FOC is computed. A firm sets each year's rent based on *some* knowledge about the state x_{jt} in that year. A signal represents the firm's knowledge when it sets a rent. It is a set of years that look alike to the firm, among which it cannot tell the current one apart. The weights are the probabilities it attaches to the years within the signal. Complete information is the special case in which every year constitutes its own signal.

Formally, let r index the rent bins defined below, and let L_r denote the number of the unit's years in bin r. A signal h of bin r is a probability vector β_h over those L_r years. The weights are its elements, representing the share of the signal's mass that comes from each year. A firm that receives the signal sees the weighted average of the states in those years, $y_h := \sum_{t \in r} \beta_{ht}\, \tilde{x}_{jt}$, where $\tilde{x}_{jt} := (x_{jt}, B_{jt} p_{jt})$ is an augmented state vector. The unit's implied MC at the signal is computed from that average. An information structure assigns a nonnegative mass α_{rh} to every such signal, subject to the mass-preservation constraint

$$\sum_h \alpha_{rh}\, \beta_{ht} = \frac{1}{L_r} \quad \text{for every year } t \text{ of bin } r. \tag{13}$$

In words, the masses α_{rh} of the signals that contain a given year (each multiplied by the year's weight β_{ht} in that signal) must add up to the mass of that year, $1/L_r$. Under this condition, the information structure is "a transport plan" that moves each year's mass to signals. Each year's state equals the weighted average at its signal plus a deviation, the mean

of which within the signal is zero. Thus, the distribution of the weighted averages of the signals is an MPC of the sample distribution of the state. It is the sample version of the MPC in section 4.3. The admissible information structures of a unit are the vectors of masses, $\alpha := (\alpha_{rh})$, that satisfy the nonnegativity constraint $\alpha_{rh} \geq 0$ and the mass-preservation constraint (13). We denote the set of such vectors by $\mathcal{A}_j$. Both constraints are linear in α.

Two basic assumptions—parts (i) and (iii) of Assumption 1—discipline the admissible information structures. First, a firm knows its own rent p_{jt}. In a sample with as many rent levels as years, this knowledge would allow no pooling at all. We operationalize it at the level of a rent bin. We cut the unit's rent history into several bins, and a signal may pool years inside the same bin but not across bins. The bin-level restriction is a conservative relaxation because the FOC holds conditional on a rent bin whenever it holds conditional on the exact level of rent. Hence, the set that this restriction produces contains the set that the exact restriction would produce (Lemma 1 in Appendix A). The number of bins determines how coarsely the firm is allowed to know its own rent. Second, a firm knows whether the event has happened. A signal may pool observations within the pre-event regime or within the post-event regime, but may not pool observations across the two regimes. Thus, each rent bin is split at the event, so that every bin holds years of one rent range and one regime. Within a bin, any pooling with any weights is admissible.

5.2 The Moment Conditions and the Range of GMM Estimates

At signal h of bin r, the implied MC is $\bar{c}_h = (\overline{Bp}_h + \bar{s}_h)/\bar{B}_h$, computed from the weighted average y_h. Using $D := -B$ to denote the absolute demand slope, the implied MC at a signal is the mean rent across the years pooled in that signal, minus the pooled markup, plus a covariance term,

$$\bar{c}_h = \bar{p}_h - \frac{\bar{s}_h}{\bar{D}_h} + \frac{\mathrm{Cov}_h(D,p)}{\bar{D}_h}, \tag{14}$$

where an overbar with the subscript h denotes the mean over the years in that signal with the signal's weights (e.g., $\bar{p}_h := \sum_{t\in r} \beta_{ht}\, p_{jt}$), and $\mathrm{Cov}_h(D,p) := \sum_{t\in r} \beta_{ht}\,(D_{jt}-\bar{D}_h)(p_{jt}-\bar{p}_h)$ is the covariance with the same weights. The covariance term is zero when the pooled years share only one rent. The sample moment at cost parameters (θ,δ) and information structure α is

$$\eta_j(\theta,\delta;\alpha) = \sum_r \lambda_r \sum_h \alpha_{rh}\,\bar{z}_h\,(\bar{c}_h - \theta - \delta\, d_r) = a_j(\alpha) - \theta\, b_j - \delta\, c_j, \tag{15}$$

where $\lambda_r := L_r/T_j$ is the bin's share of the unit's years, $d_r \in \{0,1\}$ is the regime of bin r, $a_j(\alpha) := \sum_r \lambda_r \sum_h \alpha_{rh}\,\bar{z}_h\,\bar{c}_h$, $b_j := T_j^{-1}\sum_t z_{jt}$, and $c_j := T_j^{-1}\sum_t z_{jt}\,d_{jt}$. The second

equality uses mass preservation, which fixes b_j and c_j at their sample values regardless of the information structure. This is the sample analog of (10).

With $K > 2$ instruments and only two cost parameters (θ_j and δ_j), the model is overidentified. At a given information structure, the K sample moments cannot all be zero in general. The GMM estimator minimizes the quadratic form of the sample moments with the unit's weighting matrix W_j. The matrix is the inverse of the covariance, across the unit's years, of the K moment contributions $z_{jt}\,\omega_{jt}$ under complete information, where $\omega_{jt} = p_{jt} + s_{jt}/B_{jt} - \hat{k}_j$ and $\hat{k}_j$ is the unit's mean of $p_{jt} + s_{jt}/B_{jt}$.[17] The matrix is computed once per unit and held fixed across information structures, so that the moments under every information structure are weighted on the same scale.

Given an information structure, the moment conditions are linear in the two cost parameters, which allows the GMM estimator to have a closed form,

$$\big(\hat{\theta}_j(\alpha),\, \hat{\delta}_j(\alpha)\big)^\top = (M_j^\top W_j M_j)^{-1} M_j^\top W_j\, a_j(\alpha), \tag{16}$$

where $M_j := [\,b_j \;\; c_j\,]$. This is the sample analog of (11).

The unit's estimated MC shift $\hat{\delta}_j(\alpha)$ is a linear function of the masses that constitute the information structure, as follows. The sample moment is a weighted average (over bins and signals) of the instruments times the implied cost shock at the signal, $\bar{c}_h - \theta - \delta\, d_r$, which is linear in the masses on the signals. With W_j fixed, the GMM estimator is a linear transformation of the sample moment. Let $\ell_j^\top$ denote the second row of the $2 \times K$ matrix $(M_j^\top W_j M_j)^{-1} M_j^\top W_j$ in (16), the row that produces $\hat{\delta}_j(\alpha)$. The second row of (16) is

$$\hat{\delta}_j(\alpha) = \ell_j^\top a_j(\alpha) = \sum_r \sum_h \alpha_{rh}\, \lambda_r\, \ell_j^\top \bar{z}_h\, \bar{c}_h. \tag{17}$$

Consequently, the range of $\hat{\delta}_j(\alpha)$ over all admissible information structures is an interval, $[\hat{\delta}_j^{\text{lo}}, \hat{\delta}_j^{\text{hi}}]$, the two ends of which can be computed by solving two linear programs over $\mathcal{A}_j$. That is, we search for the smallest and the largest shifts that some pooling with some weights within the unit's bins can rationalize.[18] We call this interval the unit's range of estimates,

[17] A ridge of 5% of the mean diagonal element is added to the diagonal before the inversion, which makes the matrix invertible for every unit.

[18] Operationally, the signals' probability vectors are discretized to a grid of the simplex; section 9.2 states the step. Every year on its own is a vector of the grid; hence, complete information is always available. A conic version of Carathéodory's theorem (Rockafellar, 1970) shows that at most $L_r + K$ signals per bin are ever needed to reach any value of the sample moment. Hence what the discretization loses is confined to probability vectors between the points of the grid. The bound concerns the possible signal realizations in a candidate information structure, not the signals realized in the observed sample. Each year has one realized

or its "range" for short, that is, the smallest and the largest GMM estimates of the shift over the admissible information structures. The range contains the estimate under complete information, the information structure in which every year constitutes its own signal, because that information structure is admissible.

This range is not the identified set of section 4.4. That set keeps the shifts at which some admissible information structure sets all K moments to zero, whereas the range also contains the shift that the best-fitting information structure delivers when it leaves some of them away from zero. The range contains the sample analog of the identified set's projection onto the shift, and the two coincide when every information structure fits exactly, as under exact identification ($K = 2$).[19]

As a further restriction, we also consider keeping only the information structures whose GMM objective lies within a slack c of the unit's best fit, on the scale of the overidentification test statistic. The slack is a reference point on that scale, not a critical value with an established coverage for a criterion profiled over information structures. Each end of the fit-restricted range is found by the same linear program with one added quadratic constraint.

The unit's best fit is the smallest value of the GMM objective over the cost parameters and the admissible information structures together. For given MC parameters, the smallest value over the information structures is the profiled criterion

$$S_j(\theta, \delta) := \min_{\alpha \in \mathcal{A}_j} \; \eta_j(\theta, \delta; \alpha)^\top W_j \, \eta_j(\theta, \delta; \alpha), \tag{18}$$

a quadratic program at every (θ, δ), because the moment is affine in α and the constraints are linear. Minimizing S_j over (θ, δ) is itself one convex quadratic program once the information structure's masses are included as unknowns (Appendix A.5). Its solution is the unit's best fit, the global minimum of the objective on the grid of probability vectors. The population counterpart of (18) is zero when some bin-respecting information structure sets the moment to zero, a set that contains Θ_j in (9) by Lemma 1.

signal, and L_r years contain at most L_r distinct realized signals, whereas an information structure specifies a distribution over signals conditional on each state and may split a state's mass across several possible signals. The bound preserves the state masses and the moment vector; it does not reconstruct an empirical distribution of the realized signals.

[19] The set of shifts with an exact fit, the sample analog of (12), is empty for a unit whose overidentifying restrictions are rejected at every information structure.

5.3 Aggregation to the Pairwise Gap

We aggregate the unit-level intervals $[\hat{\delta}_j^{\text{lo}}, \hat{\delta}_j^{\text{hi}}]$ first at the level of local markets, or "neighborhoods" in our empirical context (section 6), and then across their matched pairs in our propensity-score-matching design (section 7). For a neighborhood n with its set of qualified units J_n, let

$$\hat{\delta}_n^{\text{lo}} = \frac{1}{|J_n|} \sum_{j \in J_n} \hat{\delta}_j^{\text{lo}} \quad \text{and} \quad \hat{\delta}_n^{\text{hi}} = \frac{1}{|J_n|} \sum_{j \in J_n} \hat{\delta}_j^{\text{hi}}. \tag{19}$$

A matched pair $q \in \{1, 2, \ldots, Q\}$ consists of a treated neighborhood $trtd(q)$ and its control neighborhoods $n \in Ctrl(q)$, where $Ctrl(q)$ is the set of control neighborhoods in pair q, with the matching weights w_{qn} (section 7.1). The pair's difference ("gap") is the treated neighborhood's mean shift minus the matching-weighted mean of its controls', and the overall gap is its mean across all Q pairs. It is a difference-in-differences on the matched pairs, with the shift in the implied MC as the "outcome" variable. The gap is linear in the units' shifts, with positive weights on the treated units and negative weights on the controls. Hence, its lowest value is reached when every treated unit is at its lowest shift ($\hat{\delta}_{trtd(q)}^{\text{lo}}$) and every control unit at its highest ($\hat{\delta}_n^{\text{hi}}$), and its largest value in the reverse ($\hat{\delta}_{trtd(q)}^{\text{hi}}$ and $\hat{\delta}_n^{\text{lo}}$):

$$\begin{aligned} \hat{\Delta}^{\text{lo}} &= \frac{1}{Q} \sum_{q=1}^{Q} \Big[\hat{\delta}_{trtd(q)}^{\text{lo}} - \sum_{n \in Ctrl(q)} w_{qn}\, \hat{\delta}_n^{\text{hi}} \Big] \quad \text{and} \\ \hat{\Delta}^{\text{hi}} &= \frac{1}{Q} \sum_{q=1}^{Q} \Big[\hat{\delta}_{trtd(q)}^{\text{hi}} - \sum_{n \in Ctrl(q)} w_{qn}\, \hat{\delta}_n^{\text{lo}} \Big]. \end{aligned} \tag{20}$$

The interval $[\hat{\Delta}^{\text{lo}}, \hat{\Delta}^{\text{hi}}]$ is the pairwise gap's range of estimates, the gap's range for short.

5.4 Inference under Set Identification

We consider two inferential targets. The first target is a unit's identified set. A confidence set for it profiles the criterion over the information structures at every candidate value of the cost parameters and takes its critical value from the bootstrap of Fang and Santos (2019). Its asymptotics run in the number of years in the unit's history, which is small in our application, and its extension to the pairwise gap requires a joint sampling and bootstrap construction that we have not developed yet. Hence, this approach is currently under development.

The second approach targets a gap built from one selected estimate per unit. A selection rule picks one admissible information structure for every unit, and the gap is aggregated from

the unit-level estimates as in section 5.3. With each pairwise gap as one observation, the mean over the Q pairs has a cluster-robust standard error (SE), and a wild cluster bootstrap-t supplies p-values when the clusters are few (as in our empirical application). The variation of the pairwise gaps across pairs contains the estimation error of the units' shifts; hence, no separate standard error is needed for them. This inference runs in the number of pairs and their clusters. It is valid for the gap that the selection rule induces. However, whether that gap is centered on the true gap is unknown.

In the current version of the paper, we report estimates based on the second approach. We focus on the two end points of the range of GMM estimates in section 5.3. Each end point is the gap under one selection rule. For every unit, the lower end selects the information structure at which its MC shift is lowest if the unit is treated (and highest if it is a control), and the upper end does the opposite. At those information structures, the mean of the pairwise gaps has a cluster-robust SE, $\mathrm{SE}(\hat{\Delta}^{\mathrm{lo}})$ and $\mathrm{SE}(\hat{\Delta}^{\mathrm{hi}})$. The confidence interval is that of Imbens and Manski (2004) on the two ends,

$$\big[\,\hat{\Delta}^{\mathrm{lo}} - \kappa\,\mathrm{SE}(\hat{\Delta}^{\mathrm{lo}}),\ \hat{\Delta}^{\mathrm{hi}} + \kappa\,\mathrm{SE}(\hat{\Delta}^{\mathrm{hi}})\,\big]. \tag{21}$$

The critical value κ solves $\Phi(\kappa + \varrho) - \Phi(-\kappa) = 0.95$, where Φ is the standard normal cumulative distribution function and ϱ is the width of the range divided by the larger of the two SEs. This critical value lies between 1.645 and 1.96 and approaches the former as the range widens relative to its SEs.[20]

6 Data

This section describes the data. Section 6.1 introduces the building panel and its algorithmic-pricing indicators, section 6.2 the local neighborhoods and their demographics, and section 6.3 the estimation samples.

[20] As the range widens, a point of the range cannot be close to both ends at once, and the interval can miss it at only one end; the one-sided critical value of 1.645 gives 95% coverage in that case. In the number of pairs, the interval covers every point of the range that the two selections induce with probability at least 95% (Appendix A.4). Stoye (2009) gives the version that remains valid uniformly when the width is close to zero. An end that ties across several information structures is a minimum over several estimators and is biased outward, which makes the interval conservative; the correction of Chernozhukov et al. (2013) would remove that bias.

6.1 The Building Panel and the Adoption of Pricing Software

We use the building-level panel of Calder-Wang and Kim (2026). Its source is Real Estate Information Services (REIS) by Moody's Analytics, a commercial database of market-rate apartment buildings with at least 20 units in 50 US metropolitan areas ("metros"). REIS surveys owners and managers regularly. The database covers about half of the market-rate units in the country and about 80% of those in large buildings.[21] The panel is an annual snapshot from 2005 to 2019. Its unit of observation is a room type within a building (studio, one bedroom, two bedrooms, three bedrooms, and four bedrooms). We define this unit (building $\times$ room type) as a *product*. For each product and year, we observe the asking rent, the occupancy rate, and the number of units of that type. Each building also reports its year of construction, its class (A against B and C), its number of stories, and the mean size of its units. Rents are monthly asking rents, adjusted for inflation. The 50-metro panel is unbalanced, with 34,217 buildings and 1,123,127 product-years.

Calder-Wang and Kim (2026) hand-collected the use of pricing software at the level of the management company, from the archived websites of industry conferences, the marketing announcements of the software firms, and web searches. Two programs dominate the industry: YieldStar, sold by RealPage, and LRO, sold by Rainmaker until RealPage acquired it in December 2017. A third, Yardi's, had a market share of about 1% by the end of the sample. A company tagged as an adopter in a year is treated as an adopter in every later year, and the flag is attached to every building the company manages.[22] The result is a building-year indicator of algorithmic pricing. By 2019, 25% of the buildings and 34% of the units in the database were managed by adopting companies.

6.2 Neighborhoods, Submarkets, and Demographics

We focus on competition among nearby buildings at the level of a *neighborhood*, either a census tract (our primary unit) or a ZIP code area (for a sensitivity check). There are 13,240 census tracts and 3,710 ZIP code areas in our panel. A tract is designed to hold between 1,200 and 8,000 residents, with 4,000 as the target (U.S. Census Bureau, 1994); a ZIP code area holds several times as many. Even though we use these geographical units for practical relevance, they need not coincide with the exact set of buildings that a firm

[21] As of 2019 the database holds 37,216 buildings with 7.2 million units; the average building has 194 units (Calder-Wang and Kim, 2026, section 3).

[22] The management company is observed as of 2019 and assigned to every earlier year. A building that changed managers before 2019 is dated by the adoption of its 2019 manager.

competes against, or with the choice set of each individual renter. We do not call them "markets" for this reason. Between the neighborhood and the metro sits the *submarket*, the data provider's own partition of each metro into about a dozen areas (649 across the 50 metros).

Neighborhood demographics come from the five-year estimates of the American Community Survey (ACS) for census tracts. We use all 11 variables: population, median household income, median gross rent, the shares of residents who are white, Black, Asian, and Hispanic, the share with a bachelor's degree, median age, the share of renter-occupied units, and the vacancy share. The earliest usable vintage is 2010, because earlier ones are tabulated on the 2000 tract boundaries. (For ZIP code areas, we have no demographics.)

6.3 Estimation Samples

We use three nested samples. The 50-metro sample ("all 50") contains every neighborhood in the panel. The 11-metro sample ("core 11") is the part with demand estimates from Calder-Wang and Kim (2026). Our final sample is the set of matched pairs of neighborhoods in the core-11 sample whose buildings satisfy the requirements of the structural analysis. The reduced-form analysis of section 7 uses all three of them, whereas the structural analysis of section 9 focuses on the final sample. Table 1 reports the counts at the neighborhood level.

Two restrictions apply to every sample. First, adoption events (to be defined at the neighborhood level in section 7.1) are dated from 2007 to 2017, and we exclude the cohorts of 2008, 2009, and 2010 from the analysis. Those years are the housing slump that followed the subprime mortgage crisis. Second, every estimate is reported by the "phase" of the metro's rents at the event, defined as follows. For every metro and year, we compute the mean rent growth of the metro's never-adopting buildings, so that adopters' own rents do not enter the classification. We split the metro-years of the sample into terciles. A neighborhood inherits the tercile of its metro in its event year. The middle tercile, where the metro was neither slumping nor booming ("flat" phase), is the primary estimation sample, for two reasons. First, the effect of algorithmic pricing on rents is largest and most precisely estimated there (see section 7.3). Second, our BCE framework assumes stationarity. In our final sample, 46 of the 56 pairs are in Seattle, Minneapolis, and the District of Columbia, and the modal event year is 2011 (Figure B1 in Appendix B).

Table 1: Samples

	All-50 metros	Core-11 metros
Metropolitan areas	50	11
Treated neighborhood-cohorts		
All cohorts (2007–2017)	2,865	903
Baseline cohort sample (2007 and 2011–2017)	1,435	405
of which in the middle tercile of rent growth	395	101
Matched control neighborhood-cohorts		
All cohorts (2007–2017)	7,989	2,419
Baseline cohort sample (2007 and 2011–2017)	4,360	1,262
of which in the middle tercile of rent growth	1,187	279
Final sample for structural estimation		
Treated neighborhoods with at least one qualified unit	—	313
Qualified pairs	—	306
of which in the middle tercile of rent growth	—	56
Qualified units in those 56 pairs (treated / control)	—	1,701 (376 / 1,325)
Unit-years	—	24,412

Note: Neighborhoods are census tracts. The treatment and the matching are defined in section 7.1. The all-cohort rows are reported for reference; the analysis uses the baseline cohort sample. A neighborhood-cohort is one neighborhood in one event year. A control neighborhood that serves several treated neighborhoods of the same cohort is counted once, and one that serves several cohorts is counted once per cohort. A qualified unit is a product (building × room type) observed at least four years before the event and four from the event year on, in a building built before the event; a pair is qualified when its treated neighborhood and at least one of its controls hold such units. The final sample holds the cohorts of 2011 to 2016, because no unit of the 2007 or the 2017 cohort can qualify under the four-year rule.

7 Reduced-Form Analysis of Rents

This section estimates the effect of algorithmic pricing on the rents in the neighborhood. Section 7.1 explains the design, section 7.2 the estimates, and section 7.3 their sensitivity and heterogeneity.

7.1 Design

We use the propensity-score-matched difference-in-differences (PSM-DID) approach to account for the observed heterogeneity across neighborhoods (Rosenbaum and Rubin, 1983; Heckman et al., 1997).

Treatment. A neighborhood is treated in the first year in which an incumbent building adopts the software and the adopting buildings' share of the neighborhood's incumbent room capacity (number of units) reaches a threshold. No building observed in the neighborhood may have used algorithmic pricing in any earlier observed year; any prior adoption disqualifies

the neighborhood, even when the adopting buildings' capacity share was below the threshold. An incumbent building is one built before the event year and observed in the panel before it; an adopting incumbent is one that operated at least two years without the software. The threshold is 20% at the census tract and 5% at the ZIP code area.[23] Section 7.3 shows that the results are not sensitive to other thresholds. The event year is the neighborhood's cohort. Cohorts from 2007 to 2017 are stacked, with 2008 to 2010 excluded (section 6.3).

Controls. Controls are neighborhoods with no algorithmic pricing in any observed year. Within each cohort, we use a logistic regression of treatment on pre-event covariates to calculate a propensity score (i.e., the predicted probability of treatment) for every neighborhood. Each treated neighborhood is matched with its five nearest neighbors on that score among the controls in its own metro, with replacement. Each control carries one fifth of the pair's weight. The estimand is the average effect on treated neighborhoods.

Covariates. The propensity score uses 21 pre-event characteristics of the neighborhood, automatically selected from 32 candidates. Following the covariate-selection idea of Belloni et al. (2014), we first run a penalized logistic regression of treatment on the 32 candidates and a penalized linear regression of rent growth (i.e., the outcome variable) on the same candidates among the controls, and retain the union of the covariates that the two regressions select. The 21 selected covariates then enter the unpenalized logistic regression of each cohort to estimate the propensity scores.

The 21 consist of (a) nine summary statistics of the neighborhood's buildings, (b) four neighborhood-level summaries of the gap between incumbent buildings' rents and those of comparable buildings, and (c) eight census demographics. Table B1 in Appendix B lists all 32 candidates. Groups (a) and (b) are measured in the building panel before the event. At the ZIP code area, where no demographics are available, the same procedure selects 8 of 21 candidates.

Estimation. We compare neighborhoods rather than individual buildings because software use may affect the rents of both adopting buildings and their non-adopting neighbors. Our outcome averages rents across all incumbent buildings in each neighborhood, including those that do not adopt. For each cohort, the treated and the matched control neighborhoods are observed from four years before the event to four after. We stack the cohorts and regress the log of inflation-adjusted rent, averaged over the neighborhood's incumbent buildings, on an

[23] The same adoption is a smaller share of a ZIP code area, which is several times larger.

indicator for a treated neighborhood in the years from the event on, with neighborhood-by-stack and cohort-by-year fixed effects and the matching weights.[24] The coefficient on that indicator is the PSM-DID estimate, with a standard error (SE) clustered on the submarket. The same regression with event-time indicators in place of the single post-event indicator (and with the year before the event omitted) gives the event-study estimates.

Pre-trend, Balance, and Common Support. We test for a pre-trend, jointly on the coefficients at four, three, and two years before the event, and report its p-value beside the estimate. (A low value indicates a problem.) The event-time window is unbalanced. A neighborhood enters with at least two of its four pre-event years observed, and each coefficient is estimated from the neighborhoods observed at that event time. After matching, the mean absolute standardized difference over the 21 covariates is 0.09 and the largest is 0.23, with 9 of the 21 beyond 0.10. Before matching, the treated neighborhoods have higher revenues, newer buildings, larger stocks, more buildings in their submarkets, higher rents, more educated and higher-income residents, and slower rent growth in the preceding years (Table B1). The treated and the control neighborhoods share a common support of the propensity score (Table B2). No treated neighborhood has a score above 0.9. In the middle tercile, 387 of the 395 treated neighborhoods have a score inside the range of their matched controls' scores in the same cohort.

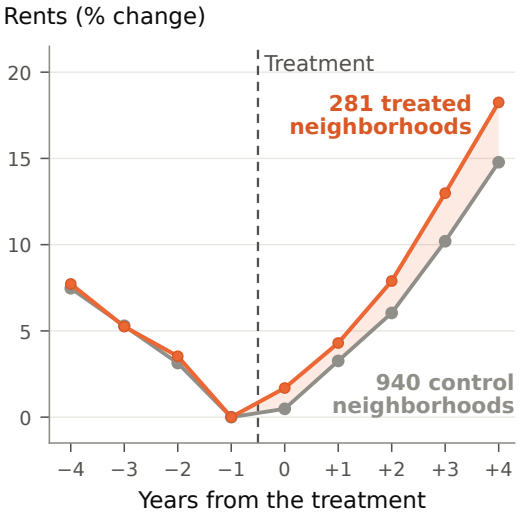
7.2 Estimates

Figure 1 shows the time evolution of rents surrounding the treatment year in the all-50-metro sample. Panel (a) plots the mean rent relative to time -1 by treatment group. The two groups move down together before the event and diverge in the pace of recovery afterward. The common decline before the event is driven by the subprime mortgage crisis in 2008–2010, because the modal cohort is 2011. Panel (b) plots the event-study estimates, the change in the treated-control gap in rents relative to the year before the event. No pre-trend is detected ($p = 0.72$). The post-event coefficients rise from +1.21% (SE 0.49) in the event year to +2.34% (SE 0.77) four years later.

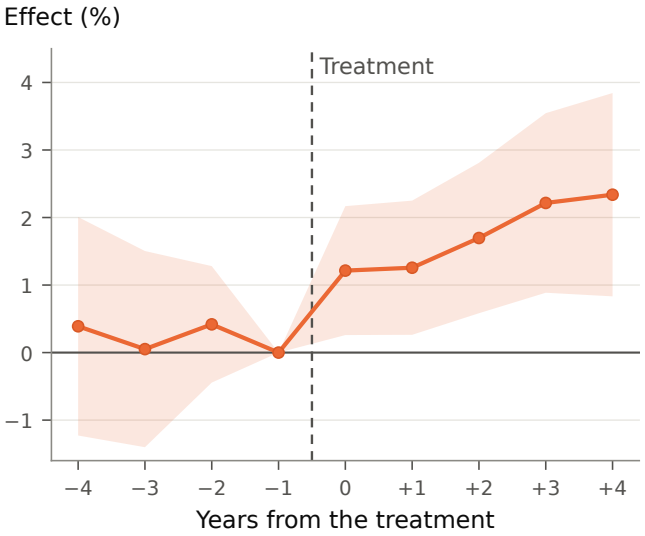
[24] The stacked data set has one entry per neighborhood, cohort, and calendar year. A treated neighborhood appears in the stack of its own cohort, for the nine years from four before the event to four after. A control neighborhood appears in the stack of every cohort in which it is matched, each time over that cohort's nine years. The neighborhood-by-stack fixed effect is a separate intercept for each neighborhood in each stack, under which every comparison is of a neighborhood with itself over time. The cohort-by-year fixed effect is a separate intercept for each cohort in each calendar year, under which a treated neighborhood is compared with the controls of its own cohort in the same year and never with another cohort.

Figure 1: Evolution of Rents around the Treatment Year

(a) Mean Rents Relative to Time −1



(b) Event-Study Estimates Relative to Matched Controls



Note: Both panels are based on census tracts in all 50 metros (middle tercile of rent growth). Panel (a) averages each neighborhood's mean incumbent rent relative to the year before the event, over a balanced-panel subset of the all-50-metro sample in which all nine years of the event-time window are available (281 treated and 940 control neighborhoods). Panel (b) plots the estimated change in the treated-control gap in log rent relative to the year before the event, in %, from the regression of section 7.1 with event-time indicators, on the full all-50-metro sample (395 treated, 1,187 controls), with 95% confidence intervals; zero indicates no change in the treated-control gap relative to the reference year. Standard errors are clustered on the submarket. The joint test of the three pre-event coefficients gives $p = 0.72$; hence no pre-trend is detected.

Table 2 reports the PSM-DID estimates by sample (all 50, core 11, and final) and geographical unit (census tract and ZIP code area), in the middle rent-growth tercile. The effect is more precisely estimated in the larger, all-50-metro sample than in the core-11 sample. Census tracts are smaller and more numerous than ZIP code areas, and feature larger estimates and smaller SEs. Given these patterns, we focus on the all-50-metro sample at the census-tract level as our baseline data set for the sensitivity and heterogeneity analysis.

Table 2: PSM-DID Estimates of the Effect of Algorithmic Pricing on Rents

Sample	Neighborhood	Effect	Treated	No-pre-trend p
All 50 metros	Census tract	+1.53 (0.53)	395	0.72
All 50 metros	ZIP code area	+1.34 (0.75)	186	0.92
Core 11 metros	Census tract	+1.57 (0.66)	101	0.58
Core 11 metros	ZIP code area	+0.89 (1.10)	49	0.41
Final sample	Census tract	+1.90 (0.90)	56	0.20

Note: All samples are in the middle rent-growth tercile, defined at the metro level in the event year. The outcome variable is incumbent-only log rent. Standard errors in parentheses are clustered on the submarket. The treatment threshold is 20% of incumbent capacity at the census tract and 5% at the ZIP code area. The tract specifications match on 21 covariates and the ZIP specifications on 8, selected by the same procedure from 32 and 21 candidates. The no-pre-trend p is the joint test that the three pre-event coefficients are zero; a high value means the null (no pre-trend) is not rejected. The final-sample row applies the same specification to the 56 treated neighborhoods of the final sample (section 6.3) and their matched controls, with the matching weights renormalized within the retained pairs.

7.3 Sensitivity and Heterogeneity

We assess the sensitivity and heterogeneity of the PSM-DID estimates in four ways. First, we alter the treatment threshold. Second, we alter the matching procedures. Third, we compare results across the three rent-growth terciles. Fourth, we distinguish neighborhoods by the number of incumbent buildings.

Treatment Threshold. The baseline threshold to define the treatment is 20% of incumbent capacity. By this definition, 395 neighborhoods (census tracts) are treated, and the estimated treatment effect is +1.53% (SE 0.53). When we relax it to 10%, 5%, and any adoption by an incumbent building, the number of treated neighborhoods rises to at most 442, and the estimated effect stays between +1.51% and +1.57% (panel A of Table B3 in Appendix B). Thus, most of the neighborhoods that experienced the use of pricing software did so at a capacity share of at least 20% in its first year.[25]

Matching Procedures. We change the details of the PSM procedures in five different aspects, for a total of 16 alternative specifications, and report the results in Table 3. Our baseline specification entails the estimate of +1.53% (SE 0.53); the 16 other specifications result in the range between +1.13% and +2.14%, with 14 of the 16 statistically significant at 5%. First, we alter the set of first-stage covariates from the 21 auto-selected ones to (1) all 32 candidates, (2) a cluster-robust selection of 17,[26] and (3) the 10 auto-selected in

[25]Meanwhile, raising the threshold from 20% to 30%, 40%, and 50% leads to substantial reductions in the number of treated neighborhoods and the precision of estimates.

[26]A control neighborhood is eligible in every cohort. The stacked data contain the same control neighborhoods several times over. The cluster-robust selection computes the penalty on clusters of the same

the core-11-metro sample. Second, we alter the controls and their weights from five nearest neighbors (NNs) to (4) one NN, (5) 10 NNs, (6) five NNs within a caliper of 0.2 standard deviations of the propensity score, (7) kernel matching, and (8) entropy balancing.[27] Third, we tighten the exact-match criteria from the same metro to also include (9) terciles of the rent pre-trend and the number-of-occupied-units pre-trend, (10) terciles of the rent pre-trend and the occupancy-rate pre-trend, (11) the same submarket,[28] and (12) an indicator for having only a single incumbent building. Fourth, we change the eligibility criteria for controls from the never-treated neighborhoods in the same metro (13) to exclude ones that share a submarket with a treated one,[29] and (14) to include the not-yet-treated ones that stay untreated through the fourth year after the event year of the treated one.[30] Fifth, we tighten or relax other sampling criteria: (15) we strengthen the requirement that at least two pre-event years are observed to at least three years, and (16) we add back the cohorts of 2008 to 2010 that we excluded in section 6.3.

Rent-Growth Terciles. The estimates are positive across all metro-rent-growth terciles, with the middle tercile exhibiting the largest and most significant treatment effect (+1.53%, SE 0.53). The top and bottom terciles entail smaller estimates, as does the pooled sample (panel B of Table B3).

Number of Incumbent Buildings. We find some heterogeneity when we split the sample between neighborhoods with only one incumbent building and those with more than one (panel C of Table B3). The estimate is larger but less precise for the former. It is tempting to conjecture about the number of competitors and its interaction with algorithmic pricing; we refrain from over-interpreting this mild heterogeneity, because census tracts (and ZIP code areas) do not necessarily coincide with "markets" in the sense of market definition. Even when a building is the only one in a census tract, it could still be competing with other buildings in adjacent tracts and any rental units outside the REIS data.

neighborhood, so that the repeated entries count as one, and selects 17 of the 32 candidates.

[27] A caliper is the maximum distance allowed between the propensity scores of a treated neighborhood and its controls; it guards against matching a treated neighborhood with controls that are not close to it. A treated neighborhood with no close control is dropped, which keeps 349 of the 395. Kernel matching weights the controls in the metro by their distance in the propensity score, with no weight beyond a bandwidth, and entropy balancing reweights every control so that the weighted means of the covariates among the controls equal their means among the treated. Both draw on the entire pool of eligible controls rather than the five nearest.

[28] The purpose is to ensure that each matched pair shares unobserved economic conditions.

[29] The purpose is to mitigate a concern over potential contamination (or spillovers) across neighborhoods.

[30] The purpose is to partially mitigate a concern that never-treated neighborhoods are fundamentally different from the treated ones.

Table 3: PSM-DID Estimates of Rent Effect under Alternative Specifications

Specification	Effect	Treated	No-pre-trend p
(0) Baseline	+1.53 (0.53)	395	0.72
A. Covariates (Baseline: 21 auto-selected)			
(1) All 32 candidates	+1.50 (0.55)	395	0.57
(2) A cluster-robust selection of 17	+1.52 (0.54)	395	0.58
(3) The 10 auto-selected in the core-11-metro sample	+2.14 (0.57)	395	0.65
B. Controls (Baseline: five NNs)			
(4) One NN	+1.32 (0.68)	395	0.29
(5) Ten NNs	+1.38 (0.53)	395	0.57
(6) Five NNs within a caliper	+1.58 (0.58)	349	0.83
(7) Kernel matching	+1.24 (0.56)	374	0.61
(8) Entropy balancing	+1.71 (0.54)	388	0.98
C. Exact match (Baseline: same metro)			
(9) Metro & pre-trend terciles (rent & occupied units)	+1.69 (0.57)	393	0.46
(10) Metro & pre-trend terciles (rent & occupancy rate)	+1.80 (0.57)	395	0.35
(11) Metro & submarket	+1.73 (0.48)	391	0.25
(12) Metro & single-building indicator	+1.37 (0.52)	395	0.84
D. Eligible controls (Baseline: never-treated)			
(13) Exclude those in the same submarket	+1.13 (0.77)	355	0.86
(14) Include not-yet-treated	+1.41 (0.56)	395	0.52
E. Sample (Baseline: two pre-event years; exclude 2008–2010)			
(15) Require three pre-event years	+1.48 (0.64)	299	0.90
(16) Include 2008–2010 cohorts	+1.28 (0.45)	602	0.94

Note: All specifications use census tracts, all 50 metros, and the middle rent-growth tercile. Each row changes one element of the baseline specification (see the main text for details). The outcome variable is incumbent-only log rent. Standard errors in parentheses are clustered on the submarket. The no-pre-trend p is the joint test that the three pre-event coefficients are zero; a high value means the null (no pre-trend) is not rejected.

8 Demand

This section explains the demand model and estimates from Calder-Wang and Kim (2026) that our BCE analysis uses. Section 8.1 summarizes the demand model and its estimates based on our longer sample. Section 8.2 reports the markups and elasticities it implies in our final sample.

8.1 The Demand Model and Its Estimation

We use the demand model of Calder-Wang and Kim (2026), re-estimated on the 2005–2019 data. A renter in submarket m and year t chooses among the products present there, or the outside option. A product is a building by bedroom count. The outside option is renting a unit that the REIS data do not cover, which accounts for about 82% of a submarket's renter households on average. Renter i's utility from product j in year t is

$$u_{ijt} = \alpha_i p_{jt} + x_{jt}\beta_i + \gamma_j + \rho_{mt} + \xi_{jt} + \epsilon_{ijt}, \tag{22}$$

where p_{jt} is the rent and x_{jt} contains four attributes: the bedroom count, an indicator for a building less than 10 years old, an indicator for class A, and an indicator for the inside good. The terms γ_j and ρ_{mt} are product and submarket-year fixed effects; ξ_{jt} is a demand shock, and ϵ_{ijt} is an idiosyncratic taste with the type I extreme value distribution. Preferences vary with the renter's demographics. The price coefficient is $\alpha_i = \alpha_0 + \alpha_y \log y_i$, where y_i is household income. Each element of β_i is linear in one or more of four demographics (income, household size, the presence of children, and whether the household head is over 35) with the interactions listed in Table 4.

The estimation uses two sources of variation, as in Petrin (2002) and Berry et al. (2004). First, the observed shares are inverted into mean utilities (Berry, 1994; Berry et al., 1995), and the rent is instrumented. Second, micro moments fit the model to two statistics of household data. One is the covariance between a household's demographics and the attributes of the product it chooses. The other is the income of the renters of class-A buildings.[31] The macro moments of the first source and the micro moments of the second are stacked in a two-step GMM estimated with the `pyblp` package (Conlon and Gortmaker, 2020). The model is estimated twice, with two sets of instruments for the rent. The first estimation uses the

[31] The covariances come from the microdata of the American Community Survey. The income of class-A renters comes from a commercial sample of households with their address histories (Data Axle), which Calder-Wang and Kim (2026) merge with the REIS buildings to identify the households that live in them.

differentiation instruments of Gandhi and Houde (2019), which measure how far a product's characteristics are from its rivals' in the same (m, t). The second estimation replaces them with approximations to the optimal instruments (Chamberlain, 1987), one per parameter, computed from the first estimates. We report the second.

In Table 4, panel (a) shows the estimates. The patterns are the same as in Calder-Wang and Kim (2026). Renters with higher income are less sensitive to the rent; larger households and households with children prefer more bedrooms; higher-income households prefer newer, class-A buildings. Those who choose an inside good tend to be richer, smaller, and younger households.

8.2 The Markups the BCE Analysis Uses

The key inputs to our BCE analysis are the share s_{jt} and the derivative of the demand with respect to its own rent, $B_{jt} = \partial s_{jt}/\partial p_{jt}$, both evaluated at the estimates. Under complete information, a firm's optimal price for a product satisfies $s_{jt} + (p_{jt} - c_{jt})B_{jt} = 0$; hence, the markup is $\mu_{jt} = -s_{jt}/B_{jt}$ and the implied marginal cost is $p_{jt} - \mu_{jt}$.[32]

Panel (b) of Table 4 reports these key objects on the final sample for our structural analysis in section 9. The mean complete-information markup is \$468, with a Lerner index (markup/rent) of 0.42. The own-price elasticity of a treated unit before treatment is -2.43; it is -2.42 for a control unit.

[32]This FOC omits the cross-price terms among the room types of one building, as an approximation. In the demand model, a product that raises its rent loses renters to all other products roughly in proportion to their market shares. The outside option (any rental outside the REIS data) accounts for about 82% of a submarket's renter households, and the median product's market share is 0.0007 (Table 4); hence, the other room types of the same building receive a negligible part of the diverted renters.

Table 4: Demand Estimates and the Implied Markups

(a) Demand Parameters	
Rent, α_0	−8.34 (0.25)
Rent × log income, α_y	0.73 (0.02)
Bedrooms × household size	0.82 (0.01)
Bedrooms × children	0.21 (0.01)
New building × income	0.03 (0.00)
New building × head over 35	−0.17 (0.01)
Class A × income	0.28 (0.01)
Inside good × income	0.60 (0.00)
Inside good × household size	−0.80 (0.01)
Inside good × children	−0.02 (0.01)
Inside good × head over 35	−0.25 (0.00)
Product-years	397,007
Submarket-years	2,849
(b) Markups and Elasticities on the Final Sample	
Qualified units	1,701
Unit-years	24,412
Mean rent	\$1,183
Mean markup under complete information	\$468
Standard deviation of the markup across units	\$57
Lerner index (markup/rent)	0.42
Median market share	0.0007
Own-price elasticity, treated units before treatment	−2.43
Own-price elasticity, control units before treatment	−2.42
Within-unit standard deviation of the rent	\$125
Within-unit standard deviation of the markup	\$18
Within-unit standard deviation of the implied cost	\$115
Within-unit slope of the markup on the rent	0.09
Within-unit slope of the implied cost on the rent	0.91

Note: In panel (a), standard errors are in parentheses. The sample is the 2005–2019 data. The rent is in thousands of dollars per month. "Children" is an indicator for a household with children and "head over 35" an indicator for a household head older than 35; the demographic variables other than log income are standardized. Panel (b) reports the main objects on the final sample for our BCE analysis in section 9. The complete-information markup is $\mu_{jt} = -s_{jt}/B_{jt}$ and the implied cost is $p_{jt} - \mu_{jt}$. Each statistic is computed within a unit over its years and averaged across the 1,701 units, with three exceptions. The standard deviation across units is that of the units' mean markups. The median market share is the median across units of the unit's mean share. The own-price elasticities average the years before treatment, over the 376 treated units and over the 1,325 control units. The two slopes are from regressions on the unit-years with unit fixed effects.

9 Empirical Implementation of the BCE Approach

This section asks whether the rent increase in section 7 can be rationalized by information sharing alone. Section 9.1 explains what we test and how it differs from Calder-Wang and Kim (2026). Section 9.2 describes the implementation details, and section 9.3 the results, limitations, and next steps.

9.1 What We Test

Our maintained hypothesis is that each firm unilaterally maximizes its own profit by setting the rent of each product. The BCE framework allows any information structure consistent with the data and estimates the range of the implied marginal costs (MCs) that rationalize the observed rent. In this sense, our implied-MC estimates are a residual of the pricing model (under the assumption of unilateral profit maximization), and not a robust measurement of cost, a primitive that applies to all possible conducts. Our test statistic is the pairwise MC-shift gap in section 5.3, that is, the change in this residual (implied MC) at the time of treatment, in treated neighborhoods relative to their matched controls. We call it the "gap" for simplicity.

To be precise, the null hypothesis has two parts: (i) firms price unilaterally under some admissible information structure, and (ii) treatment changed no firm's true MC relative to its controls'. Under the null, the gap is zero. Whatever the software did to a firm's information is already inside the set of admissible information structures, and the PSM procedure cancels out the cost movements common to the matched pair of treated and control neighborhoods. Hence, a positive gap means that post-treatment rents in the treated neighborhoods are too high to be optimal for a unilateral profit-maximizer under any admissible information structure. That is, rationalizing the observed rents under unilateral pricing requires a larger increase in the implied MC in treated neighborhoods than in their matched controls. The possibility is that either the true MC increased specifically in the treated neighborhoods, or the treated firms no longer price unilaterally.

The test is a conduct test in the general sense of Berry and Haile (2014). The expansion of algorithmic pricing is potentially the excluded markup shifter, and the markup is computed under every admissible information structure. That said, the test differs from that of Calder-Wang and Kim (2026), and hence of Backus et al. (2021), in the dimension along which the firm's FOC is relaxed. They maintain complete information and compare conducts; we maintain unilateral conduct and relax informational assumptions. Thus, the two frameworks do not nest each other, and are complementary.

9.2 Implementation Details

Here we describe the operational details of the estimator in section 5.

Estimation Unit and Sample. An estimation unit is a product (a room type of a building). We estimate the model unit by unit. Each product's FOC is evaluated separately, with

the own-price demand derivative B_{jt} from section 8.2. Our final sample consists of the 56 qualified pairs of census tracts from the core-11 metros in the middle tercile of rent growth (see Table 1 in section 6). They contain 1,701 qualified units (376 in treated and 1,325 in control neighborhoods), observed 9 to 15 years per unit, for the total of 24,412 unit-years. A control unit takes the event year of its treated pair as its own.

Cost. A unit's MC is constant within each regime (the years before the event or the years since) and is allowed to shift at the event. Our empirical specification of MC is the same as equation (1) in section 3.1, where θ_j is the pre-event level, δ_j is the shift at the event, and ω_{jt} is a cost shock.

Information. An information structure can pool multiple years of a unit into a single signal, with weights (section 5.1). Two basic assumptions discipline the admissible information structures. First, a firm knows its own rent p_{jt}, which we operationalize as the quartiles of the unit's rent in the entire sample period (we consider terciles and quintiles as well). A signal may pool years inside the same quartile bin but not across bins. Second, a firm knows whether the treatment has happened (i.e., whether a substantial fraction of incumbent capacity in the neighborhood has started using pricing software), which means a signal may not pool years from both the pre- and post-treatment periods. Thus, a unit has up to eight information bins, four rent quartiles in each of the two regimes. Within a bin, any pooling with any weights is admissible. We discretize the signals' probability vectors to multiples of 0.05. Complete information is a special case in which every year constitutes its own signal (i.e., no pooling).

Instruments and Overidentification. The moment conditions are $\mathbb{E}[z_{jt}\,\omega_{jt}] = 0$, where z_{jt} contains nine of the 11 optimal instruments in the demand estimation.[33] With nine instruments and only two cost parameters (θ_j and δ_j), the model is overidentified. The overidentifying restrictions (i.e., the seven moment conditions beyond the two needed to solve for θ_j and δ_j) are rejected at the best-fitting information structure in 59% of the units. The set of shifts with an exact fit would be empty for most units, whereas the range of GMM

[33] The two instruments (the expected rent and its interaction with income) are dropped because, within a unit and regime, they are correlated with the cost residual (-0.55 and -0.61), whereas the other nine's correlation coefficients lie within ± 0.18. The cost residual in this assessment is computed under the assumption of complete information to generate yearly variation. The two IVs' correlation with the cost residual is not a problem for demand estimation because cost shifters are valid IVs for demand estimation. Instead, this finding suggests the MC function in (1) should be augmented to include cost shifters. With no cost shifter in (1), any yearly change in cost is absorbed by ω_{jt}, and an IV built from the fitted rent moves with it.

estimates in section 5.2 is not. For the fit restriction (also in section 5.2), we use the slacks: (i) $c = 3.84$ (the 95% point of a χ^2 with one degree of freedom), under which the information structures fit the moments nearly as well as the best one, and (ii) $c = 14.07$ (the 95% point with seven degrees of freedom), under which they fit no worse than the best by the margin of the overidentification test. Appendix C reports what these fit restrictions achieve.

Aggregation and Inference in Practice. The gap is aggregated over the 56 pairs as in section 5.3, with the second approach of section 5.4. Some control neighborhoods serve treated neighborhoods in different submarkets. Because those pairs share control observations, their gaps can be correlated across submarkets. We merge the submarkets linked by shared controls and cluster at that level, which gives nine groups. Appendix C reports the same p-values with the clusters at the 29 submarkets of the treated neighborhoods, a finer clustering that ignores the sharing of controls across submarkets, which represents a less conservative inference. With few clusters, we take the p-values from the wild cluster bootstrap-t with the null imposed (Cameron et al., 2008; Webb, 2023), with 9,999 draws and Webb's six-point weights. The p-value of the lower end against zero, from the bootstrap, is the conservative test of the conduct null, because a rejection says that the gap is positive at every information structure the range admits. The demand estimates from section 8 are held fixed throughout, and every range and confidence interval is conditional on them.[34]

Restrictions on Information. The range of estimates depends on what the information structures are allowed to be. Our baseline is the least restrictive version, in which any admissible information structure is allowed in each regime under the two basic assumptions alone (i.e., a firm knows its own rent up to the quartiles and whether the treatment has happened). We add two restrictions in turn, each an assumption beyond the main model. The first is the information-sharing mechanism's strongest form, under which a treated firm's information is complete after treatment. That is, every post-treatment year constitutes its own signal, whereas any admissible information structure is allowed before treatment and for the control firms. The second, added to the first, is that a control firm's information does not change at its matched treated neighborhood's event year. The rationale for this specification is that the control firms are in the PSM-DID design to cancel out the cost movements common to the pair, and equation (1) allows a control's MC to shift at that

[34] Carrying the sampling error of a first-stage demand system into the second stage requires re-estimating the first stage inside the bootstrap, which Canay et al. (2023) discuss for set-identified models. We have not implemented this re-estimation due to high computational costs.

date. This restriction says that nothing happens to its information on that date. It is the information-structure counterpart of the parallel-trends assumption in the PSM-DID analysis. Operationally, the share of a control unit's years that its information structure pools into signals with other years must be the same before and after that date.

9.3 Results

Table 5 reports the range of estimates for the gap across the admissible information structures and the confidence interval (CI) at the end points of the range. The three panels discretize the unit's rent into three, four, and five bins. Within each panel, the three rows start from the baseline, in which every admissible information structure is allowed to every unit, and move down to increasingly more restrictive assumptions. The second row makes the treated firms' information complete after treatment but leaves the control firms' unconstrained. The third further requires that the control firms' information be unchanged at the event.

In the first row of panel (b) (all admissible information structures with four rent bins), the range of the gap is [−\$2.8, +\$81.5]. The unilateral-conduct null cannot be rejected when the firms are allowed to be sufficiently ignorant, and the lower end's p-value against zero is 0.83. The benchmark of complete information after treatment (the second row) tightens the range to [+\$13.3, +\$61.9], which excludes zero, whereas the CI at its end points, [−\$3.1, +\$75.9], still includes it. The invariance of the control firms' information (the third row) does not change the range by more than a dollar at either end.

The number of rent bins governs the extent of information pooling, thereby affecting the width of the ranges as well. With three bins, a signal may pool more years, and every row of panel (a) is wider than its counterpart in panel (b). With five bins, firms are assumed to know the market state more precisely; consequently, every row of panel (c) is narrower, with its two bottom rows rejecting zero at 5%. The overall pattern is the same in the three panels. The range excludes zero whenever the treated firms' information is assumed complete after treatment (i.e., in every row from the second on), whereas the CI excludes zero only under more restrictive assumptions.

Table 5: The Range of Estimates of the Treated-Control Gap

Information structure	Range ($)	CI at the end points ($)	p
(a) Three Rent Bins			
Baseline: any admissible structure	$[-22.0, +98.5]$	$[-43.0, +113.4]$	0.41
Treated: complete ex-post; controls: any	$[+1.3, +73.7]$	$[-16.3, +86.5]$	0.91
Treated: complete ex-post; controls: unchanged at the event	$[+1.6, +73.3]$	$[-16.0, +86.2]$	0.89
(b) Four Rent Bins (Baseline)			
Baseline: any admissible structure	$[-2.8, +81.5]$	$[-21.1, +96.3]$	0.83
Treated: complete ex-post; controls: any	$[+13.3, +61.9]$	$[-3.1, +75.9]$	0.25
Treated: complete ex-post; controls: unchanged at the event	$[+13.9, +61.5]$	$[-2.7, +75.5]$	0.23
(c) Five Rent Bins			
Baseline: any admissible structure	$[+9.6, +73.1]$	$[-6.7, +87.2]$	0.49
Treated: complete ex-post; controls: any	$[+20.4, +56.6]$	$[+4.2, +70.9]$	0.04
Treated: complete ex-post; controls: unchanged at the event	$[+21.0, +55.9]$	$[+4.9, +70.3]$	0.04

Note: The range of estimates is the set of values that the GMM estimate of the gap takes across the admissible information structures (section 5.3). The CI at the end points of the range is that of Imbens and Manski (2004), namely each end minus or plus 1.645 times its standard error, clustered on the nine merged groups of the 56 pairs (section 5.4). The critical value of that section equals 1.645 in every row, because every range is wide relative to the standard errors of its ends. "p" is the bootstrap p-value of the lower end against zero on the same clusters, the conservative test of the conduct null. Appendix C reports the p-values with the clusters at the 29 submarkets.

Selected Points. Table 6 reports the gap at two selected information structures. The inference is the second approach of section 5.4, with one estimate per unit aggregated along the pairs as in section 5.3 and the wild cluster bootstrap-t on the pairwise gaps. The first row selects, for every unit, the best-fitting information structure that the grid search of Appendix A.5 found.[35] The gap is \$120.2 − \$79.2 = \$41.0, with $p = 0.032$ on the 29 submarkets and 0.050 on the nine merged groups. The second row selects complete information, the information structure in which every year of every unit constitutes its own signal and the one that Calder-Wang and Kim (2026) maintain. The gap is +\$38.9, with $p = 0.040$ and 0.047. With terciles or quintiles in place of quartiles, the gap at the best fit is +\$37.4 or +\$43.0. Each p-value tests the gap that its selection rule induces; the range of Table 5 shows how far the other admissible information structures move it.

[35]The unit's best fit of section 5.2 is the solution of one convex quadratic program, whose value lies below the grid search's in every unit (Appendix A.5). The fit restriction of Table C4 and the 59% of section 9.2 are computed at it. With every unit at the program's best fit, the gap is +\$40.3, with $p = 0.042$ on the submarkets and 0.049 on the merged groups.

Table 6: The Gap at Two Selected Information Structures ($)

				p-value	
Information structure	Treated	Control	Gap	Submarkets	Merged
Best fit	+120.2	+79.2	+41.0	0.032	0.050
Complete information	+114.8	+75.9	+38.9	0.040	0.047

Note: The 56 qualified pairs, 1,701 units, four rent bins. "Best fit" selects for every unit the information structure at the best fit that the grid search of Appendix A.5 found; "Complete information" makes every year of every unit its own signal. "Treated" and "Control" are the mean shifts in the implied MC of the treated neighborhoods and of their matched controls, aggregated as in section 5.3, and the gap is their difference. The p-values are from the wild cluster bootstrap-t with the null imposed and 9,999 draws, with each pairwise gap as one observation, and the clusters at the 29 submarkets of the treated neighborhoods (Rademacher weights) or at the nine groups of submarkets merged through shared controls (Webb's six-point weights).

What the Width of the Range Measures: Diagnosis. The width of the baseline range, \$84.3 ($= 81.5 - (-2.8)$) at four bins, is not a measure of how much information could change the markup that unilateral pricing implies. Appendix C.5 decomposes it by re-solving the two linear programs under two diagnostic changes of the estimator, with everything else held fixed. Replacing the slope-weighted rent at a signal by the plain mean rent, which removes the covariance term of (14), changes the width by under a dollar. Replacing the nine instruments by the constant and the regime indicator, under which the moments are exactly identified and the shift at an information structure is the mass-weighted before-after mean shift of the implied MC, cuts the width from \$84.3 to \$7.4; with both changes it is \$2.2, the width that the pooled markup $\bar{s}_h/\bar{B}_h$ alone produces. Thus, the width comes from the re-pairing of the nine instruments' year-to-year values with the implied costs inside bins. The shift is identified through the instruments' differences between the regimes, which are small within a unit. Pooling swings an overidentified GMM estimate by tens of dollars, whereas it moves the implied markup by a dollar or two, as the nearly constant markup within a unit (Table 4) implies. Under the exactly identified moments the range over all admissible information structures is [+\$23.9, +\$31.3], around the before-after mean shift of the implied MC of +\$27.7 (Table C5). The two restrictions move the range by under \$3 at either end. The standard errors of its ends are about \$14, and the lower end's p-value is 0.22. The nine instruments carry the overidentifying content that Appendix C.1 tests, and the two do not; the choice between them is a choice of moment conditions, not of information.

Limitations and Next Steps. Three limitations qualify the results, and each points to a next step. The first is sampling uncertainty. The final sample has 56 pairs in nine clusters, and the standard errors of the ends of the range are about \$14. No restriction on information

reduces them. A larger sample does, and the extension to the metros outside the core 11 requires demand estimates for them. The second is the moment conditions. The Hansen test rejects the overidentifying restrictions of the constant-cost specification (Table C1 in Appendix C), and the width of the baseline range comes from the instruments rather than from the markup. A cost equation with submarket-by-year effects, under which the asymptotic test's rejection weakens to $p = 0.07$, is a candidate revision of the estimator. The third is inference. The CI at the end points is valid for the gap that each end's selection rule induces (section 5.4), whereas a confidence set for a unit's identified set, the first approach of that section, is under development. Every range and CI is also conditional on the demand estimates, which are held fixed.

References

Aleksenko, A. and J. Miklós-Thal (2026): "Third-Party Pricing Algorithms and Information Sharing," Working paper, May 2026.

Asker, J., C. Fershtman, and A. Pakes (2022): "Artificial Intelligence, Algorithm Design, and Pricing," *AEA Papers and Proceedings*, 112, 452–456.

Assad, S., R. Clark, D. Ershov, and L. Xu (2024): "Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market," *Journal of Political Economy*, 132, 723–771.

Backus, M., C. Conlon, and M. Sinkinson (2021): "Common Ownership and Competition in the Ready-to-Eat Cereal Industry," NBER Working Paper 28350.

Belloni, A., V. Chernozhukov, and C. Hansen (2014): "Inference on Treatment Effects after Selection among High-Dimensional Controls," *Review of Economic Studies*, 81, 608–650.

Bergemann, D. and S. Morris (2013): "Robust Predictions in Games with Incomplete Information," *Econometrica*, 81, 1251–1308.

——— (2016): "Bayes Correlated Equilibrium and the Comparison of Information Structures in Games," *Theoretical Economics*, 11, 487–522.

Berry, S., J. Levinsohn, and A. Pakes (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 63, 841–890.

——— (2004): "Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market," *Journal of Political Economy*, 112, 68–105.

Berry, S. T. (1994): "Estimating Discrete-Choice Models of Product Differentiation," *RAND Journal of Economics*, 25, 242–262.

Berry, S. T. and P. A. Haile (2014): "Identification in Differentiated Products Markets Using Market Level Data," *Econometrica*, 82, 1749–1797.

Bresnahan, T. F. (1982): "The Oligopoly Solution Concept Is Identified," *Economics Letters*, 10, 87–92.

Brown, Z. Y. and A. MacKay (2023): "Competition in Pricing Algorithms," *American Economic Journal: Microeconomics*, 15, 109–156.

Calder-Wang, S. and G. H. Kim (2026): "Algorithmic Pricing in Multifamily Rentals: Efficiency Gains or Price Coordination?" Working paper, The Wharton School and Boston College, May 2026.

Calvano, E., G. Calzolari, V. Denicolò, and S. Pastorello (2020): "Artificial Intelligence, Algorithmic Pricing, and Collusion," *American Economic Review*, 110, 3267–3297.

Cameron, A. C., J. B. Gelbach, and D. L. Miller (2008): "Bootstrap-Based Improvements for Inference with Clustered Errors," *Review of Economics and Statistics*, 90, 414–427.

Canay, I. A., G. Illanes, and A. Velez (2023): "A User's Guide for Inference in Models Defined by Moment Inequalities," *Journal of Econometrics*, 235, 1636–1660.

Chamberlain, G. (1987): "Asymptotic Efficiency in Estimation with Conditional Moment Restrictions," *Journal of Econometrics*, 34, 305–334.

Chernozhukov, V., S. Lee, and A. M. Rosen (2013): "Intersection Bounds: Estimation and Inference," *Econometrica*, 81, 667–737.

Clark, R. and J.-F. Houde (2026): "Do Information Sharing Platforms Facilitate Coordination? Evidence from the Collapse of a Cartel," Working paper, September 2026.

Clarke, R. N. (1983): "Collusion and the Incentives for Information Sharing," *Bell Journal of Economics*, 14, 383–394.

Conlon, C. and J. Gortmaker (2020): "Best Practices for Differentiated Products Demand Estimation with PyBLP," *RAND Journal of Economics*, 51, 1108–1161.

Duarte, M., L. Magnolfi, M. Sølvsten, and C. Sullivan (2024): "Testing Firm Conduct," *Quantitative Economics*, 15, 571–606.

Eichenbaum, M. S., L. P. Hansen, and K. J. Singleton (1988): "A Time Series Analysis of Representative Agent Models of Consumption and Leisure Choice under Uncertainty," *Quarterly Journal of Economics*, 103, 51–78.

Fang, Z. and A. Santos (2019): "Inference on Directionally Differentiable Functions," *Review of Economic Studies*, 86, 377–412.

Gal-Or, E. (1985): "Information Sharing in Oligopoly," *Econometrica*, 53, 329–343.

Gandhi, A. and J.-F. Houde (2019): "Measuring Substitution Patterns in Differentiated-Products Industries," NBER Working Paper 26375.

Hansen, L. P. (1982): "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, 50, 1029–1054.

Harrington, Jr., J. E. (2018): "Developing Competition Law for Collusion by Autonomous Artificial Agents," *Journal of Competition Law & Economics*, 14, 331–363.

——— (2022): "The Effect of Outsourcing Pricing Algorithms on Market Competition," *Management Science*, 68, 6889–6906.

——— (2025a): "The Challenges of Third Party Pricing Algorithms for Competition Law," *Theoretical Inquiries in Law*, forthcoming.

——— (2025b): "A Critique of Recent Remedies for Third-Party Pricing Algorithms and Why the Solution Is Not Restrictions on Data Sharing," *Journal of Competition Law & Economics*, nhaf022.

Heckman, J. J., H. Ichimura, and P. E. Todd (1997): "Matching as an Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Programme," *Review of Economic Studies*, 64, 605–654.

Imbens, G. W. and C. F. Manski (2004): "Confidence Intervals for Partially Identified Parameters," *Econometrica*, 72, 1845–1857.

Johnson, J. P., A. Rhodes, and M. Wildenbeest (2023): "Platform Design When Sellers Use Pricing Algorithms," *Econometrica*, 91, 1841–1879.

Klein, T. (2021): "Autonomous Algorithmic Collusion: Q-Learning under Sequential Pricing," *RAND Journal of Economics*, 52, 538–558.

Kühn, K.-U. and X. Vives (1995): "Information Exchanges Among Firms and Their Impact on Competition," Tech. rep., Office for Official Publications of the European Communities, Luxembourg.

Magnolfi, L. and N. Lomys (2024): "Estimation of Games under No Regret," Working paper.

Magnolfi, L. and C. Roncoroni (2023): "Estimation of Discrete Games with Weak Assumptions on Information," *Review of Economic Studies*, 90, 1413–1448.

Magnolfi, L. and C. Sullivan (2022): "A Comparison of Testing and Estimation of Firm Conduct," *Economics Letters*, 212, 110316.

Miklós-Thal, J. and C. Tucker (2019): "Collusion by Algorithm: Does Better Demand Prediction Facilitate Coordination Between Sellers?" *Management Science*, 65, 1552–1561.

New York State Senate (2025): "Senate Bill S7882," Enacted legislation.

Pakes, A. (2010): "Alternative Models for Moment Inequalities," *Econometrica*, 78, 1783–1822.

Pakes, A., J. Porter, K. Ho, and J. Ishii (2015): "Moment Inequalities and Their Application," *Econometrica*, 83, 315–334.

Petrin, A. (2002): "Quantifying the Benefits of New Products: The Case of the Minivan," *Journal of Political Economy*, 110, 705–729.

Raith, M. (1996): "A General Model of Information Sharing in Oligopoly," *Journal of Economic Theory*, 71, 260–288.

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

Rosenbaum, P. R. and D. B. Rubin (1983): "The Central Role of the Propensity Score in Observational Studies for Causal Effects," *Biometrika*, 70, 41–55.

Stoye, J. (2009): "More on Confidence Intervals for Partially Identified Parameters," *Econometrica*, 77, 1299–1315.

Syrgkanis, V., E. Tamer, and J. Ziani (2021): "Inference on Auctions with Weak Assumptions on Information," Working paper.

U.S. Census Bureau (1994): "Geographic Areas Reference Manual," U.S. Department of Commerce, Washington, DC; chapter 10, Census Tracts and Block Numbering Areas.

USDOJ (2024): "United States et al. v. RealPage, Inc." U.S. Department of Justice, Complaint, Case No. 1:24-cv-00710, U.S. District Court for the Middle District of North Carolina.

——— (2025a): "United States et al. v. RealPage, Inc., et al." U.S. Department of Justice, Amended Complaint, Case No. 1:24-cv-00710, U.S. District Court for the Middle District of North Carolina.

——— (2025b): "United States et al. v. RealPage, Inc., et al." U.S. Department of Justice, Proposed Final Judgment and Competitive Impact Statement, Case No. 1:24-cv-00710, U.S. District Court for the Middle District of North Carolina; 90 Fed. Reg. 56286 (December 5, 2025).

Vives, X. (1984): "Duopoly Information Equilibrium: Cournot and Bertrand," *Journal of Economic Theory*, 34, 71–94.

Webb, M. D. (2023): "Reworking Wild Bootstrap-Based Inference for Clustered Errors," *Canadian Journal of Economics*, 56, 839–858.

Appendix

Contents

A		**The Model and the Estimator: Proofs and Derivations**	**45**
	A.1	Proof of Proposition 1	45
	A.2	The Affine Moment and the Linearity of the Shift	46
	A.3	The Bin-Level Relaxation Is Conservative	46
	A.4	The Confidence Interval at the End Points of the Range	47
	A.5	The Computation of a Unit's Best Fit	48
B		**Reduced-Form Analysis: Additional Exhibits**	**49**
	Figure B1	Where the 56 Qualified Pairs Are	49
	Table B1	The 32 Candidate Covariates: Selection and Balance	50
	Table B2	Common Support of the Propensity Score, by Cohort	51
	Table B3	PSM-DID Estimates by Treatment Threshold, Rent-Growth Tercile, and Number of Incumbent Buildings	52
C		**Structural (BCE) Analysis: Additional Exhibits**	**53**
	C.1	The Instruments and the Cost Moment	53
	Table C1	Overidentification Tests of the Cost Moment	54
	C.2	The Restriction Sets on Information	54
	Table C2	The Range of Estimates of the Gap under the Restriction Sets at Four Rent Bins ($)	55
	Table C3	The Test of the Lower End with the Clusters at the Submarkets	56
	C.3	The Fit Restriction	56
	Table C4	The Fit Restriction at Four Rent Bins over All Admissible Structures	57
	C.4	The Rent Change and Its Decomposition under Complete Information	57
	Table C5	The Rent and Its Complete-Information Decomposition, Before-After Mean Shifts ($)	57
	C.5	The Width of the Range, Decomposed	58
	Table C6	The Width of the Range of the Gap under Two Diagnostic Changes of the Estimator ($)	58
	Table C7	The Range of the Gap under the Exactly Identified Moments ($)	59

A The Model and the Estimator: Proofs and Derivations

This appendix gives the proof of Proposition 1 (section A.1), the derivation of the affine form of the moment and of the linearity of the shift in the information structure's masses (section A.2), the argument that the bin-level relaxation of section 5.1 is conservative (section A.3), the argument for the confidence interval at the end points of the range in section 5.4 and its limits (section A.4), and the computation of a unit's best fit (section A.5). The product index j is dropped throughout.

A.1 Proof of Proposition 1

Write $\tilde{x}_t := (p_t, s_t, B_t, B_t p_t, z_t, d_t)$ for the state extended by the product $B_t p_t$. A distribution ν is a mean-preserving contraction of the distribution of $\tilde{x}_t$ if there is a joint distribution of $(\tilde{x}_t, Y_t)$ under which Y_t has distribution ν and $\mathbb{E}[\tilde{x}_t \mid Y_t] = Y_t$ almost surely. It is price- and regime-respecting if, under that joint distribution, the rent and regime entries of Y_t equal p_t and d_t almost surely. We write the entries of Y_t as $(p_t, \bar{s}_t, \bar{B}_t, \overline{Bp}_t, \bar{z}_t, d_t)$ and set $\bar{c}_t := (\overline{Bp}_t + \bar{s}_t)/\bar{B}_t$.

(a) implies (b). Let $\mathcal{I}_t$ and ω_t be as in (a), and set $Y_t := \mathbb{E}[\tilde{x}_t \mid \mathcal{I}_t]$. Since p_t and d_t are measurable with respect to $\mathcal{I}_t$, their entries of Y_t are p_t and d_t. Since Y_t is measurable with respect to $\mathcal{I}_t$, the law of iterated expectations gives $\mathbb{E}[\tilde{x}_t \mid Y_t] = \mathbb{E}[\mathbb{E}[\tilde{x}_t \mid \mathcal{I}_t] \mid Y_t] = Y_t$. Hence the distribution ν of Y_t is a price- and regime-respecting mean-preserving contraction. By (1) and the measurability of $c_t = \theta + \delta d_t + \omega_t$, the first-order condition (2) reads $\bar{s}_t + \overline{Bp}_t - \bar{B}_t(\theta + \delta d_t + \omega_t) = 0$. Since $\bar{B}_t \neq 0$, it gives $\omega_t = \bar{c}_t - \theta - \delta d_t$. Then

$$\mathbb{E}[z_t\, \omega_t] = \mathbb{E}\big[\mathbb{E}[z_t \mid \mathcal{I}_t]\, \omega_t\big] = \mathbb{E}\big[\bar{z}_t\,(\bar{c}_t - \theta - \delta d_t)\big] = \eta(\theta, \delta; \nu),$$

where the first equality uses the measurability of ω_t. Under the exclusion restriction, the left-hand side is zero, which gives (b).

(b) implies (a). Let ν be as in (b). Take a joint distribution of $(\tilde{x}_t, Y_t)$ under which Y_t has distribution ν, $\mathbb{E}[\tilde{x}_t \mid Y_t] = Y_t$, and the rent and regime entries of Y_t equal p_t and d_t. Let the signal be the vector of interim means itself, $\mathcal{I}_t := Y_t$, so that p_t and d_t are measurable with respect to $\mathcal{I}_t$. Define $\omega_t := \bar{c}_t - \theta - \delta d_t$, which is a function of Y_t and measurable with respect to $\mathcal{I}_t$. Assumption 1 holds. Conditioning the integrand of (2) on $\mathcal{I}_t$ and using $\mathbb{E}[\tilde{x}_t \mid Y_t] = Y_t$ entry by entry,

$$\mathbb{E}\big[s_t + B_t\,(p_t - \theta - \delta d_t - \omega_t)\,\big|\,\mathcal{I}_t\big] = \bar{s}_t + \overline{Bp}_t - \bar{B}_t\,(\theta + \delta d_t + \omega_t) = \bar{s}_t + \overline{Bp}_t - \bar{B}_t\,\bar{c}_t = 0,$$

by the definition of $\bar{c}_t$. Thus (2) holds almost surely. Finally, $\mathbb{E}[z_t\omega_t] = \mathbb{E}[\mathbb{E}[z_t \mid Y_t]\,\omega_t] = \mathbb{E}[\bar{z}_t\,\omega_t] = \eta(\theta,\delta;\nu) = 0$, which is the exclusion restriction. Thus (a) holds. ■

A.2 The Affine Moment and the Linearity of the Shift

Population. Under a mean-preserving contraction, $\mathbb{E}[\bar{z}_t] = \mathbb{E}[\mathbb{E}[z_t \mid Y_t]] = \mathbb{E}[z_t] = b$. Because d_t is an entry of Y_t, $\mathbb{E}[\bar{z}_t\,d_t] = \mathbb{E}[\mathbb{E}[z_t \mid Y_t]\,d_t] = \mathbb{E}[\mathbb{E}[z_t d_t \mid Y_t]] = \mathbb{E}[z_t d_t] = c$. Substituting both into (8) gives (10). The information structure enters only through $a(\nu) = \mathbb{E}_\nu[\bar{z}_t\,\bar{c}_t]$.

Sample. Fix a unit with T years, bins r with L_r years and shares $\lambda_r = L_r/T$, probability vectors β_h, and masses α_{rh} satisfying (13). A signal's interim means are $y_h = \sum_{t\in r}\beta_{ht}\,\tilde{x}_t$, with $\bar{z}_h = \sum_{t\in r}\beta_{ht}\,z_t$ as the instrument entry. Mass balance gives

$$\sum_r \lambda_r \sum_h \alpha_{rh}\,\bar{z}_h = \sum_r \frac{L_r}{T}\sum_{t\in r} z_t \sum_h \alpha_{rh}\,\beta_{ht} = \sum_r \frac{L_r}{T}\sum_{t\in r}\frac{z_t}{L_r} = \frac{1}{T}\sum_t z_t = b.$$

The same computation with $d_r = d_t$ for every year t of bin r gives $\sum_r \lambda_r \sum_h \alpha_{rh}\,\bar{z}_h\,d_r = c$. Hence (15) holds, with $a(\alpha) = \sum_r \lambda_r \sum_h \alpha_{rh}\,\bar{z}_h\,\bar{c}_h$ linear in α. At a fixed α the moment is affine in (θ,δ). The minimizer of $\eta^\top W\eta$ is the GMM solution (16), whose second row gives (17). The profiled criterion (18) minimizes a quadratic form in an affine function of α over a polytope, which is a quadratic program. The ends of the unit's interval minimize and maximize a linear function of α over the same polytope, which are two linear programs. Every vertex of the simplex is a probability vector of the grid; hence, the information structure in which every year is its own signal is admissible, and the complete-information estimate lies in the range.

A.3 The Bin-Level Relaxation Is Conservative

Lemma 1. *Let the bins of a unit be a partition of its rents, each bin split at the event, and call an information structure bin-respecting if every signal pools years of one bin only. Every admissible information structure is bin-respecting. Consequently, at every (θ,δ) the profiled criterion over the bin-respecting information structures is no larger than the criterion over the admissible ones, and the set of cost parameters at which the former is zero contains Θ_j.*

Proof. The bin of a year is a function of its rent and its regime. If the rent and the regime are measurable with respect to the information set, the bin is measurable as well. Hence, given a signal, the state lies in one bin almost surely, and the signal's probability vector puts mass on that bin alone. Thus the admissible information structures form a subset of

the bin-respecting ones. A minimum over a larger set is no larger; hence, the zero set of the minimum can only grow. □

The lemma also has the reading of the law of iterated expectations. When the first-order condition holds conditional on an information set that contains the rent, it holds conditional on the bin, which the information set also contains.

A.4 The Confidence Interval at the End Points of the Range

Fix the sample's years of every unit, so that each unit's admissible set $\mathcal{A}_j$ and its two ends $\hat{\delta}_j^{\text{lo}}$ and $\hat{\delta}_j^{\text{hi}}$ are statistics of the unit's own data. An end is the unit's GMM estimate at the information structure that the linear program selects. The pair's lower-end gap, $\hat{G}_q^{\text{lo}} := \hat{\delta}_{trtd(q)}^{\text{lo}} - \sum_{n \in Ctrl(q)} w_{qn}\, \hat{\delta}_n^{\text{hi}}$, is one observation per pair, as a selected point's gap is, and $\hat{\Delta}^{\text{lo}}$ of (20) is its mean over the Q pairs; likewise $\hat{G}_q^{\text{hi}}$ and $\hat{\Delta}^{\text{hi}}$. Let $\Delta_*^{\text{lo}} := \mathbb{E}[\hat{G}_q^{\text{lo}}]$ and $\Delta_*^{\text{hi}} := \mathbb{E}[\hat{G}_q^{\text{hi}}]$ denote the expectations of the pairs' end gaps at the sample's years. As the number of pairs grows with the years per unit fixed, with the pairs' end gaps independent across the clusters, the two means are jointly asymptotically normal around $(\Delta_*^{\text{lo}}, \Delta_*^{\text{hi}})$, with variances that the cluster-robust standard errors estimate consistently. The argument is the one that justifies the inference on a selected point, applied to the two selections. Under these conditions, Imbens and Manski (2004) show that the interval (21), with κ chosen by the rule stated after it, covers every point of $[\Delta_*^{\text{lo}}, \Delta_*^{\text{hi}}]$ with asymptotic probability at least 0.95; Stoye (2009) gives the version that is uniform in the interval's width. If the information structure that attains an end is not unique, the sample end is the minimum over several asymptotically normal estimators and is biased outward, which makes the interval conservative.

Two things the argument does not establish. First, the relation between $[\Delta_*^{\text{lo}}, \Delta_*^{\text{hi}}]$ and the population range of section 5.3, and with it the true gap. The sample's admissible set is the set of transport plans over a unit's observed years, whereas the true information structure is a distribution over signals given the state; the former approximates the latter only as the unit's years grow, and each end is a GMM estimate on a few years per regime (in the final sample, the median unit has 6 years before the event, with a range of 4 to 9, and 9 from the event year on, with a range of 5 to 9). The selected point under complete information faces the same question, and we make no claim about it for either. Second, sharpness. The range admits information structures whose moments fit badly, and the interval inherits their shifts. A profiled construction that treats all units at once, with a critical value from the bootstrap of Fang and Santos (2019), would address both, at the cost of re-solving every unit's quadratic programs in every bootstrap draw.

A.5 The Computation of a Unit's Best Fit

Section 5.2 defines a unit's best fit as the smallest value of the GMM objective over the cost parameters and the admissible information structures together. On the grid of probability vectors, the objective is a convex quadratic function of the information structure's masses and the two cost parameters jointly. The sample moment is affine in the masses and in (θ, δ) (Appendix A.2), the constraints on the masses, nonnegativity and the mass preservation of (13), are linear, and W_j is fixed. The best fit is the solution of one convex quadratic program,

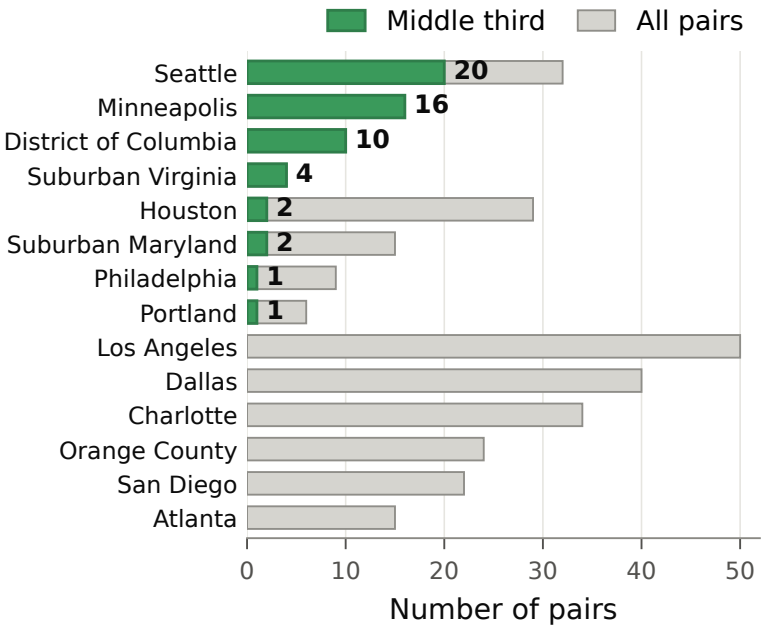
$$\min_{\alpha \geq 0,\, \theta,\, \delta} \; \eta_j(\theta, \delta; \alpha)^\top W_j \, \eta_j(\theta, \delta; \alpha) \quad \text{subject to (13)}, \tag{23}$$

and it is the global minimum of the objective on the grid of probability vectors. We solve the equivalent program in the masses alone, obtained by substituting the closed form (16) for (θ, δ) at each information structure. The cost parameters at the best fit are that closed form at the minimizing information structure. No bound on the cost level or on the shift is imposed. The probability vectors themselves stay on the grid, because moving them off it is a nonconvex problem, which we do not attempt.

The best fits reported in Table 6 were found by a two-pass grid search over the cost level and the shift, with (18) re-solved at every point. The coarse pass evaluates 11 levels by 11 shifts, the shift over ±\$750 per month and the level over a window around the complete-information implied costs of the units in the same neighborhood and cohort. The fine pass evaluates 11 by 13 around the coarse minimizer. The program's value lies below the search's in every one of the 1,701 units, and the program's minimizer lies inside the search's window for all but ten of them. The program's best fit is the reference point of the fit restriction of Table C4 and of the share of units that reject the overidentifying restrictions in section 9.2.

B Reduced-Form Analysis: Additional Exhibits

Figure B1: Where the 56 Qualified Pairs Are



Note: The 306 qualified pairs of the final sample by metro (gray), and the 56 in the middle tercile (green). The tercile is cut on metro-level rent growth, which puts a whole metro-year into one tercile at once; 46 of the 56 pairs are in Seattle, Minneapolis, and the District of Columbia, and 50 of the 56 adopted in 2011. Washington's Maryland and Virginia suburbs and Orange County are separate labels in the data provider's geography.

Table B1: The 32 Candidate Covariates: Selection and Balance

Variable	Selected	Standardized difference	
		Before	After
A. The neighborhood's buildings			
Revenue (log)	T	+0.71	+0.18
Capacity (units)	O	+0.45	+0.17
Capacity (log)	O	+0.60	+0.15
Number of products (buildings × room types)	—	+0.23	+0.14
Share of products in class-A buildings	O	+0.35	−0.03
Mean building age	O	−0.69	−0.03
Number of buildings in the submarket	O	+0.31	+0.23
Mean unit size (square feet)	—	+0.27	−0.10
Number of stories	O	+0.00	+0.06
Rent, log of the mean rent	—	+0.38	+0.09
Rent, mean of the log rents	O	+0.38	+0.09
Occupancy rate	—	−0.15	−0.05
Rent trend, log of the mean rent	—	−0.26	−0.02
Rent trend, mean of the log rents	O	−0.27	−0.02
Occupied-unit trend (log)	—	+0.10	+0.07
Occupancy-rate trend	—	+0.01	+0.02
B. The gap between a building's rent and comparable buildings'			
Gap, level	O	+0.24	+0.04
Gap, trend	O	−0.03	−0.11
Gap, transitory part	O	+0.06	−0.07
Gap, dispersion across the neighborhood's buildings	O	+0.28	+0.17
Share of buildings more than 10 log points under benchmark	—	−0.23	−0.05
C. Census demographics of the tract			
Population (log)	—	+0.21	+0.05
Median household income (log)	—	+0.34	−0.04
Median gross rent (log)	O	+0.27	+0.04
Share white (non-Hispanic)	—	+0.16	−0.10
Share Black (non-Hispanic)	O	−0.23	+0.02
Share Asian (non-Hispanic)	O	+0.35	+0.11
Share Hispanic	O	−0.15	+0.05
Share with a bachelor's degree or higher, ages 25 and over	O	+0.36	−0.04
Median age	O	−0.11	−0.11
Share of renter-occupied units	O	−0.13	+0.11
Vacancy share	O	−0.10	−0.08

Note: All results use census tracts, all 50 metros, the middle tercile, and the baseline specification. "Selected" names the equation of the post-double-selection procedure that selected the variable: T the treatment equation, O the outcome equation; "—" is a candidate that neither selected. The standardized difference is the treated-minus-control difference in means divided by the square root of the average of the two groups' variances, before matching and after matching. Groups A and B are measured in the building panel before the event, where a level is the mean over the four pre-event years (fewer if some are unobserved) and a trend the least-squares slope over them; group C is the level recorded in the census vintage labeled with the year before the event (the 2010 vintage for the 2007 cohort). The two rent levels, and the two rent trends, differ only in the form of the average over the neighborhood's buildings, either the log of the mean rent or the mean of the buildings' log rents. In group B, the gap is a product's rent residual from a hedonic regression estimated year by year on the buildings that never adopt the software, relative to its submarket-year benchmark, averaged over the neighborhood's incumbent buildings; its transitory part is the last pre-event year's gap minus the building's own mean over the earlier years.

Table B2: Common Support of the Propensity Score, by Cohort

A. All terciles of rent growth

	Treated				
Cohort	Count	Has propensity score ≥ 0.9	Matched with a control within caliper	Control pool	Pseudo-R^2
2007	206	0	146	7,747	0.30
2011	854	0	781	7,949	0.15
2012	133	0	123	7,969	0.07
2013	96	0	88	8,008	0.09
2014	23	0	21	8,035	0.07
2015	64	0	52	8,066	0.18
2016	34	0	30	8,139	0.17
2017	25	0	18	8,219	0.31
All	1,435	0	1,259	—	—

B. Middle tercile of rent growth only

	Treated					Controls			
	Count	Score			Inside range	Count	Score		
Cohort		Min	Median	Max			Min	Median	Max
2007	91	0.003	0.094	0.825	89	167	0.003	0.055	0.650
2011	148	0.020	0.140	0.873	146	362	0.013	0.121	0.599
2012	67	0.003	0.028	0.109	66	278	0.002	0.026	0.104
2013	72	0.001	0.017	0.186	70	293	0.001	0.016	0.136
2014	10	0.001	0.004	0.032	10	50	0.001	0.004	0.033
2016	5	0.003	0.017	0.050	4	25	0.002	0.014	0.047
2017	2	0.109	0.129	0.149	2	10	0.056	0.065	0.154
All	395	0.001	0.062	0.873	387	1,187	0.001	0.036	0.650

Note: All results use census tracts, all 50 metros, and the baseline specification. Panel A uses all terciles of rent growth, because the score is calculated and the matching is done before the sample is split by rent-growth tercile. One logistic regression is run per cohort; it converged in every cohort, and its pseudo-R^2 is reported. The control pool is the number of control neighborhoods scored in that cohort. The pools do not add up to a total because the same neighborhood is eligible in several cohorts, and neither column has an All row because each cohort has its own regression. "Has propensity score ≥ 0.9" counts the treated neighborhoods with a propensity score at or above 0.9, and "Matched with a control within caliper" those that find at least one control within a caliper of 0.2 standard deviations of the score and are matched to up to five such controls. Panel B describes the scores of the 395 treated neighborhoods of the middle tercile and of their 1,187 matched controls. "Inside range" counts the treated neighborhoods whose score lies inside the range of their own cohort's matched controls, a lower bound on the count inside the range of the whole control pool. The 2015 cohort has no treated neighborhood in the middle tercile; its two matched controls there enter only the All row.

Table B3: PSM-DID Estimates by Treatment Threshold, Rent-Growth Tercile, and Number of Incumbent Buildings

	Effect	Treated	No-pre-trend p
A. Treatment threshold (share of incumbent capacity)			
20% (baseline)	+1.53 (0.53)	395	0.72
10%	+1.51 (0.50)	435	0.74
5%	+1.57 (0.50)	442	0.80
Any adoption by an incumbent building	+1.56 (0.50)	442	0.78
B. Rent-growth tercile of the metro at the event			
Bottom tercile (slumping)	+0.98 (0.43)	739	0.99
Middle tercile (flat; baseline)	+1.53 (0.53)	395	0.72
Top tercile (booming)	+0.79 (0.63)	301	0.42
All terciles pooled	+1.10 (0.34)	1,435	0.67
C. Number of incumbent buildings in the neighborhood			
One	+1.82 (1.03)	136	0.72
More than one	+1.35 (0.65)	255	0.64

Note: All rows use census tracts, all 50 metros, and the baseline specification of section 7.1, with one element changed at a time; panels A and C use the middle rent-growth tercile. In panel A, the threshold is the adopting buildings' share of the neighborhood's incumbent room capacity in the event year, and "any adoption" counts a neighborhood as treated at the first adoption by an incumbent building. In panel B, the tercile is that of the metro's rent growth among its never-adopting buildings in the event year (section 6.3). In panel C, the split is by the number of incumbent buildings in the neighborhood at the event, applied to the treated and the control neighborhoods alike, which is why the two rows do not add up to the 395 treated neighborhoods of the baseline. The outcome variable is incumbent-only log rent. Standard errors in parentheses are clustered on the submarket. The no-pre-trend p is the joint test that the three pre-event coefficients are zero; a high value means the null (no pre-trend) is not rejected.

C Structural (BCE) Analysis: Additional Exhibits

This appendix carries the exhibits that section 9 summarizes in the text: (i) the overidentification tests of the cost moment (Table C1), (ii) the restriction sets on information under which the range of the gap was computed at four rent bins (Table C2) and the rows of Table 5 with the clusters at the submarkets (Table C3), (iii) the fit restriction of section 5.2 and what it achieves (Table C4), (iv) the before-after shift of the rent and its decomposition under complete information (Table C5), and (v) the width of the range decomposed by two diagnostic changes of the estimator (Tables C6 and C7).

C.1 The Instruments and the Cost Moment

Table C1 tests the overidentifying restrictions of the complete-information cost moment on the final sample. The cost levels of (1), one per unit and regime, are removed by a within transformation. With no parameter left in the moment $\mathbb{E}[z_{jt}\,\omega_{jt}] = 0$, the Hansen J statistic (Hansen, 1982) tests it directly. The covariance of the moment vector is clustered on the neighborhood-cohort (249 clusters). A cluster bootstrap with the moments recentered at the sample supplies a second p-value. The C statistic (Eichenbaum et al., 1988) tests the two expected-rent instruments given the other nine. Three cost specifications are compared. First, the constant cost of (1). Second, the same with a metro-by-year cost effect. Third, the same with a submarket-by-year cost effect. Each effect is partialled out of the instruments and the implied cost. The fourth block of Table C1 tests the difference of the moment vectors after and before the event, which is the part of a violation that moves a unit's shift. The last block tests the treated-minus-control difference of that difference, which is the part the gap inherits.

Table C1: Overidentification Tests of the Cost Moment

Cost specification and sample	Test	Statistic	df	p	Bootstrap p
Constant cost, every unit-year	J, eleven	217.8	11	<0.001	<0.001
	J, nine	143.8	9	<0.001	<0.001
	C, two given nine	74.0	2	<0.001	<0.001
Plus metro-by-year cost effects	J, eleven	79.0	11	<0.001	<0.001
	J, nine	36.5	9	<0.001	<0.001
	C, two given nine	42.5	2	<0.001	<0.001
Plus submarket-by-year cost effects	J, eleven	52.1	11	<0.001	<0.001
	J, nine	15.9	9	0.070	0.163
	C, two given nine	36.3	2	<0.001	<0.001
Constant cost, after minus before the event	J, eleven	160.9	11	<0.001	<0.001
	J, nine	135.5	9	<0.001	<0.001
	C, two given nine	25.4	2	<0.001	0.001
Constant cost, treated minus control of after minus before	J, eleven	17.3	11	0.098	0.256
	J, nine	17.0	9	0.049	0.142
	C, two given nine	0.3	2	0.851	0.883

Note: The final sample's 24,412 unit-years (1,701 units, 56 pairs). Each J is a Hansen test of the overidentifying restrictions $\mathbb{E}[z_{jt}\,\omega_{jt}] = 0$ for the stated instrument set. The unit-by-regime cost levels are removed by a within transformation, and the moment covariance is clustered on the neighborhood-cohort. C is the difference between the J of the eleven instruments and the J of the nine, computed with the covariance of the eleven. It tests the two expected-rent instruments given the other nine. The asymptotic p-value uses the χ^2 reference with the stated degrees of freedom. The bootstrap p-value draws clusters with replacement, 9,999 times, with the moments recentered so that the resampled data satisfy the null. A cost effect is partialled out of the instruments and the implied cost by alternating demeaning. "After minus before" tests the difference of the two regimes' moment vectors. The last block tests the treated-minus-control difference of that difference.

C.2 The Restriction Sets on Information

Table C2 lists the restriction sets under which the range of the gap of section 9.3 was computed at four rent bins, fifteen of the 22. The seven omitted, the treated caps with the controls free, the ordering restriction with the controls unchanged, and the bin-by-bin invariance with the treated free, differ from a row shown by under a dollar at either end. The table carries the CI at the end points, the p-value at the lower end on the submarkets and on the merged groups, and the p-value at the upper end on the merged groups. The sets were fixed before the computation, and none was chosen on the result. A restriction set is a pair, the information structures allowed for the treated units and those allowed for the controls. "Pooling after treatment at most x" caps the share of a treated unit's post-treatment years that its information structure may pool, with any admissible information structure before treatment. It runs from the pin of the benchmark ($x = 0$) to the baseline ($x = 1$). The

ordering restriction, that a firm is at least as informed after treatment as before, is reported here and not in Table 5, for two reasons. A smaller share of pooled mass after treatment does not establish that the firm is more informed when prices are endogenous, and the restriction moves the range by under a dollar. Table C3 repeats the test of the lower end for the nine rows of Table 5 with the clusters at the 29 submarkets of the treated neighborhoods.

Table C2: The Range of Estimates of the Gap under the Restriction Sets at Four Rent Bins ($)

			Submarkets	Merged groups	
Treated units / control units	Range	CI at the end points	p, lower	p, lower	p, upper
Any / any	$[-2.8, +81.5]$	$[-21.1, +96.3]$	0.866	0.831	0.019
At least as informed after treatment / the same	$[-2.0, +80.6]$	$[-20.3, +95.8]$	0.905	0.879	0.020
Complete information after treatment / the same	$[+28.7, +49.1]$	$[+14.0, +63.9]$	0.114	0.046	0.042
Complete information after treatment / any	$[+13.3, +61.9]$	$[-3.1, +75.9]$	0.446	0.254	0.030
Complete information after treatment / unchanged at the event (the third row of Table 5)	$[+13.9, +61.5]$	$[-2.7, +75.5]$	0.430	0.231	0.030
Complete information after treatment / unchanged, bin by bin	$[+13.5, +61.8]$	$[-3.0, +75.7]$	0.441	0.245	0.030
Any / unchanged at the event	$[-2.3, +81.1]$	$[-20.7, +95.9]$	0.890	0.862	0.019
Pooling after treatment at most 0.05 / unchanged	$[+11.2, +64.7]$	$[-5.7, +78.8]$	0.515	0.426	0.028
Pooling after treatment at most 0.1 / unchanged	$[+8.6, +67.9]$	$[-8.6, +82.1]$	0.611	0.550	0.026
Pooling after treatment at most 0.2 / unchanged	$[+3.4, +74.4]$	$[-14.5, +88.8]$	0.841	0.781	0.021
Pooling after treatment at most 0.3 / unchanged	$[+0.8, +77.6]$	$[-17.4, +92.4]$	0.960	0.946	0.020
Pooling after treatment at most 0.5 / unchanged	$[-1.6, +80.3]$	$[-20.1, +95.2]$	0.923	0.903	0.019
Complete information after treatment / pooling at most 0.1 in either regime	$[+27.6, +49.4]$	$[+11.9, +63.9]$	0.129	0.044	0.040
Complete information after treatment / at most 0.2 in either regime	$[+21.3, +54.8]$	$[+5.2, +69.0]$	0.234	0.042	0.036
Complete information after treatment / at most 0.5 in either regime	$[+14.5, +60.9]$	$[-1.9, +75.0]$	0.411	0.188	0.031

Note: The 56 qualified pairs, four rent bins. Each row describes the information structures allowed for the treated and the control units. The range, the CI at the end points, and the p-values are constructed as in Table 5. The critical value of Imbens and Manski (2004) is 1.645 in every row. "Unchanged at the event" requires the share of a control unit's years that its information structure pools to be the same before and after its pair's event year, as one equality per unit. "Bin by bin" imposes the same equality within each rent bin that holds two or more years in both regimes; it has content in 44% of the units. "At least as informed after treatment" requires the share of pooled years after treatment to be no larger than before. "Submarkets" clusters on the 29 submarkets of the treated neighborhoods (Rademacher weights) and "Merged groups" on the nine groups of submarkets merged through shared controls (Webb's six-point weights).

Table C3: The Test of the Lower End with the Clusters at the Submarkets

	p at the lower end, submarkets		
Information structure	Three bins	Four bins	Five bins
Baseline: any admissible structure	0.20	0.87	0.56
Treated: complete ex-post; controls: any	0.94	0.45	0.25
Treated: complete ex-post; controls: unchanged at the event	0.92	0.43	0.24

Note: The bootstrap p-value of the lower end of the range against zero for the nine rows of Table 5, with the clusters at the 29 submarkets of the treated neighborhoods (Rademacher weights, 9,999 draws, the null imposed). It is a finer clustering than the nine merged groups of Table 5, and it ignores that some controls serve pairs in different submarkets. The rows and the columns are those of Table 5.

C.3 The Fit Restriction

Section 5.2 defines a restriction on fit rather than on information, which keeps only the information structures whose GMM objective lies within a slack c of the unit's best fit, with $c = 3.84$ or 14.07 on the scale of the overidentification test statistic (section 9.2). Table C4 reports what it achieves at four rent bins when every admissible information structure is allowed to every unit, the first row of panel (b) of Table 5. The slack of 3.84 narrows the range at both ends and lowers the p-value of the lower end; the slack of 14.07 leaves the range nearly where it was. The restriction bites because the ends of the baseline range are reached by information structures that fit worse than the unit's best fit. The GMM objective at the lower end exceeds the unit's best fit by 6.0 at the median unit and by more than 3.84 in 67% of the units, and at the upper end by 4.5 and in 55% of the units. Even so, the restriction narrows the range less than the restrictions on information of Table 5 do, and three cautions apply. The slacks are reference points on the scale of the statistic, not critical values with an established coverage for a criterion profiled over information structures, which makes the fit-restricted ranges descriptive. The slack is relative to the unit's best fit, which itself rejects the overidentifying restrictions in 59% of the units, and the restriction is not an absolute test of fit. And it is applied unit by unit, which makes the result a fit-restricted range, not a confidence set with joint coverage across the 1,701 units.

Table C4: The Fit Restriction at Four Rent Bins over All Admissible Structures

Information structures kept	Range ($)	CI at the end points ($)	p
Baseline: any admissible structure	$[-2.8, +81.5]$	$[-21.1, +96.3]$	0.83
GMM objective within 3.84 of the unit's best fit	$[+7.1, +72.2]$	$[-10.8, +87.3]$	0.61
GMM objective within 14.07 of the unit's best fit	$[-1.5, +80.4]$	$[-19.9, +95.3]$	0.90

Note: The 56 qualified pairs, four rent bins, every admissible information structure allowed to every unit (the first row of panel (b) of Table 5), and the same with the fit restriction of section 5.2. The range, the CI at the end points, and the p-value are constructed as in Table 5; the critical value of Imbens and Manski (2004) is 1.645 in every row. The slack is on the scale of the overidentification test statistic: 3.84 is the 95% point of a χ^2 with one degree of freedom, and 14.07 is the 95% point with seven degrees of freedom.

C.4 The Rent Change and Its Decomposition under Complete Information

Table C5 reads the gap against the rent change on the same pairs. Under complete information the rent is the implied MC plus the markup by identity; hence, the before-after shifts decompose. For each unit and each of the three quantities, the shift is the plain difference of the unit's mean after and before the event; the shifts are aggregated along the pairs as in section 5.3, and the gap is the treated mean minus the control mean. On the 56 pairs the rent gap is +\$33.7. Of it, +\$27.7 is implied MC and +\$6.1 is markup. The rent row is not the regression of Table 2, which uses every incumbent building of a neighborhood within the event window. The markup absorbs under a fifth of the rent increase. Section 9.3 draws on this table.

Table C5: The Rent and Its Complete-Information Decomposition, Before-After Mean Shifts ($)

				p-value	
	Treated	Control	Gap	Submarkets	Merged
Rent	+109.9	+76.1	+33.7	0.010	0.031
Implied MC	+94.7	+67.0	+27.7	0.011	0.102
Implied markup	+15.2	+9.1	+6.1	0.125	0.065

Note: The 56 qualified pairs, 1,701 units. "Treated" and "Control" are the mean before-after shifts of the treated neighborhoods and of their matched controls, each averaged within a neighborhood over its units and across the pairs with the matching weights, and the gap is their difference. The rows apply the plain before-after difference of a unit's mean, with no instruments, to the rent, to the complete-information implied MC $p_{jt} - \mu_{jt}$, and to the markup μ_{jt} of section 8.2. The rent row is the sum of the other two by construction. The p-values are from the wild cluster bootstrap-t of section 5.4, with the clusters at the 29 submarkets of the treated neighborhoods (Rademacher weights) or at the nine merged groups (Webb's six-point weights).

C.5 The Width of the Range, Decomposed

Section 9.3 reads the width of the range as the re-pairing of the instruments with the implied costs inside bins. Tables C6 and C7 are the computation behind that reading. Two diagnostic changes of the estimator are made, each alone and both together, with the information structure of record, the restriction sets, the aggregation, and the inference held fixed. The first replaces the slope-weighted rent at a signal by the plain mean rent of the signal's years, which removes the covariance term of (14) and nothing else. The second replaces the nine instruments by the constant and the regime indicator, so that the moments are exactly identified, the weighting matrix is immaterial, and the shift at an information structure is the mass-weighted mean of the implied MC over the post-treatment signals minus that over the pre-treatment signals. Under both changes at once, mass preservation fixes the plain mean rent of each regime, which leaves the range to the pooled markup $\bar{s}_h/\bar{B}_h$ alone. Table C6 reports the width of the gap's range under each of the four estimators and four restriction sets at four rent bins; its last row pins every unit, treated and control, at complete information after the event. Table C7 reports the range itself under the exactly identified moments. At complete information the plain mean rent equals the slope-weighted rent, which leaves the complete-information point unchanged under the first change; under the second the point is the before-after mean shift of the implied MC of Table C5, +\$27.7, unit by unit. With the record's estimator every row reproduces Table 5 and Table C2.

Table C6: The Width of the Range of the Gap under Two Diagnostic Changes of the Estimator (\$)

	Nine instruments		Constant and regime	
Information structure	Slope-weighted	Plain mean	Slope-weighted	Plain mean
Baseline: any admissible structure	84.3	84.9	7.4	2.2
Treated: complete ex-post; controls: any	48.6	48.6	4.8	1.5
Treated: complete ex-post; controls: unchanged at the event	47.6	47.6	4.8	1.5
Complete information after the event, every unit	20.3	20.3	2.6	0.9

Note: The 56 qualified pairs, four rent bins. Each entry is the width of the range of the gap, the upper end minus the lower end, under the stated estimator and information structures. "Nine instruments, slope-weighted" is the estimator of Table 5, with the slope-weighted rent at a signal. "Plain mean" replaces it by the plain mean rent of the signal's years. "Constant and regime" replaces the nine instruments by the constant and the regime indicator, under which the moments are exactly identified. The first three rows are the rows of panel (b) of Table 5; the last pins every unit, treated and control, at complete information after the event.

Table C7: The Range of the Gap under the Exactly Identified Moments ($)

Information structure	Range, slope-weighted rent	Range, plain-mean rent	p
Baseline: any admissible structure	[+23.9, +31.3]	[+26.6, +28.9]	0.221
Treated: complete ex-post; controls: any	[+25.4, +30.2]	[+26.8, +28.3]	0.184
Treated: complete ex-post; controls: unchanged at the event	[+25.4, +30.2]	[+26.8, +28.3]	0.183
Complete information after the event, every unit	[+26.4, +29.0]	[+27.2, +28.1]	0.145

Note: The 56 qualified pairs, four rent bins. The moments are $\mathbb{E}[\omega_{jt}] = 0$ and $\mathbb{E}[d_{jt}\,\omega_{jt}] = 0$ in place of the nine of section 9.2; every admissible information structure fits them exactly, and the shift at an information structure is the before-after difference of the mass-weighted mean implied MC. The range is constructed as in Table 5, with the slope-weighted rent at a signal (the record) or the plain-mean rent. "p" is the bootstrap p-value of the lower end against zero on the nine merged groups, under the slope-weighted rent; the standard errors of the two ends of the first row are \$14.1 and \$13.6. The complete-information point under these moments is the before-after mean shift of the implied MC of Table C5, +\$27.7.